Short Papers

Towards a General Multi-View Registration Technique

Robert Bergevin, Marc Soucy, Hervé Gagnon, and Denis Laurendeau

Abstract—We present an algorithm that reduces significantly the level of the registration errors between all pairs in a set of range views. This algorithm refines initial estimates of the transformation matrices obtained from either the calibrated acquisition setup or a crude manual alignment. It is an instance of a category of registration algorithms known as iterated closest-point (ICP) algorithms. The algorithm considers the network of views as a whole and minimizes the registration errors of all views simultaneously. This leads to a well-balanced network of views in which the registration errors are equally distributed, an objective not met by previously published ICP algorithms which all process the views sequentially. Experimental results show that this refinement technique improves the calibrated registrations and the quality of the integrated model for complex multi-part objects. In the case of scenes comprising man-made objects of very simple shapes, the basic algorithm faces problems common to all ICP algorithms and must thus be extended.

Index Terms—Multiple view registration, model building, three-dimensional objects, range images, surface modeling, multiple view integration.

✦

1 Introduction

BUILDING integrated models of existing 3-D objects is a key requirement for object duplication or re-design in reverse engineering systems. It is also essential for object matching and identification in object recognition systems. Up to very recently, the lack of appropriate tools to accomplish the difficult and time consuming model building task has been a serious obstacle to the practical realization of such systems. Fortunately, it is now becoming conceivable to develop an automatic model builder, based on active vision, which could significantly improve the speed and flexibility of 3-D model acquisition as compared to interactive techniques such as computer-aided design (CAD) and coordinate-measurement machines (CMM).

A typical automatic 3-D model builder based on active vision has to go through three main steps [6]. Firstly, the surface of the object under study must be sampled. Existing laser range finders now allow the acquisition of dense and accurate range maps [12]. Since a range map only samples the surface of the object which is visible from a given viewpoint, the acquisition of several range views is mandatory in order to scan the entire object. Each range view having its own reference frame, a second step is needed in which all range views are transformed into a common reference frame. The operation of estimating the rigid inter-frame transformations between the range views is known as registration. The result of this estimation is called the registration transformation. Once the set of range views is registered, it must be integrated, in a third step, into a non-redundant surface model. A general integration technique, such as one based on the reparameterization of the canonic subsets of the Venn diagram, may be used in order to integrate the set of range views into a non-redundant 2.5-D surface triangulation [14], [15], [16].

The accuracy that may be obtained with an integration technique such as the one presented in [14] is mainly influenced by two kinds of errors: the acquisition and the registration errors. The former is caused by the limited accuracy of the laser range finder. Since this error can be efficiently modeled as Gaussian [1], one can expect to improve the quality of information in areas measured by more than one range views simply by fusing the different measurements using a weighted average [14], [16]. However, the combination of range measurements having different variances is to give rise to an improvement if and only if the Gaussian distributions are all centered around the same mean [16]. The registration error is thus critical in the integration process since the misalignment of the range views will cause the noise distributions of range measurements from different views to be centered around different means. Hence, the accurate registration of a set of range views is of major concern in the design and implementation of a 3-D model builder based on active vision. In this correspondence, we present an algorithm that reduces significantly the level of the registration errors between all pairs in a set of range views.

Two acquisition strategies are often used in order to obtain multiple views of a single object. In the first instance, the object is placed on a turntable which rotates in front of a fixed range finder, taking images at discrete steps. In the second approach, the range finder itself is mounted on a moving robot's arm, taking views around the fixed object. Both strategies are used in the experiments reported here. A first estimation of the inter-frame transformations may be obtained, in the first case, by measuring the center of rotation and the rotation angles applied to the turntable and, in the second case, by computing the direct kinematics equation of the robot.[1] We have observed through several experimentations that the accuracy of calibrated registrations are typically insufficient for the sake of accurate surface modeling. A refinement technique is thus needed in order to minimize the effects of registration errors. The algorithm presented in this correspondence refines initial estimates of the transformation matrices obtained from either the calibrated acquisition setup or a crude manual alignment.

A number of different methods have been proposed to tackle the refinement problem [11], [7], [9], [13], [6], [3], [17], [18]. Among these various techniques, the approach proposed by Chen and Medioni [6] proved to be the most appropriate starting point in our effort to accurately register a set of range views. Their technique may be seen as an instance of a category of registration algorithms known as iterated closest-point (ICP) algorithms [3], [17], [18]. It refines a nominal transformation by iteratively computing incremental transformations that minimize the distance between the transformed points from the first view and the surface from the second view. No fixed point-to-point match is needed since the minimized functional is expressed in terms of the distance between each control point in the first view and the tangent plane at an intersected point in the second view. This choice is appropriate since it tends not only to minimize the distance between the two surfaces but also to locally match their curvatures. By computing only small changes to the transformations, the incremental transformation matrix may be linearized and the six parameter increments, three angles and three translations, directly

1. The initial transformations could also be obtained directly from the pairs of views using a method such as the one described in [2].

• The authors are with the Computer Vision and Systems Laboratory, Department of Electrical and Computer Engineering, Laval University, Québec, Canada G1K 7P4. E-mail: bergevin@gel.ulaval.ca.

Manuscript received Jan. 25, 1995; revised Nov. 24, 1995. Recommended for acceptance by G. Medioni.
For information on obtaining reprints of this article, please send e-mail to: transactions@computer.org, and reference IEEECS Log Number P95177.

computed using a linear least-squares technique. The particularity of their technique is that it uses the raw range data without computing any surface representations or extracting features. Such a refinement technique is thus perfectly suited for refining the registration of views of complex objects even from computed coarse estimates [2]. However, because of the sequential nature of its registration process, the method of Chen and Medioni cannot produce the best possible registration for a given accuracy of the range data. The algorithm presented in this correspondence resolves this fundamental limitation by adapting the approach to the registration in parallel of the complete set of range views. Each range view is seen as a node in a network of views. A link between two nodes stores the inter-frame transformation between the two associated views. The approach presented here considers the network of views as a whole and minimizes the registration errors of all views simultaneously. This leads to a well-balanced network of views in which the registration errors are equally distributed, an objective not met by previously published ICP algorithms which all process the views sequentially. Experimental results show that this new refinement technique improves the calibrated registrations and the quality of the integrated model for complex multi-part objects. Still, in the case of scenes comprising man-made objects of very simple shapes, the basic algorithm faces other problems common to all ICP algorithms and must thus be extended.

The paper is organized as follows. A refinement algorithm to register a pair of range views is presented in Section 2. Improvements brought to the algorithm proposed by Chen and Medioni are described. First, a spatial neighborhood test and a surface visibility test are used to constrain the set of control points for the least-squares technique. Second, a sub-pixel interpolation process improves the computation of the distances between a point in one view and tangent planes in the other view. In Section 3, the new multi-view approach is presented, followed by experimental results in Section 4. The needed extensions to the algorithm are discussed in the final section.

2 Registration of a Pair of Range Views

The refinement technique described by Chen and Medioni [6] computes a rigid transformation matrix between two range views which minimizes the distance between points from the first view and tangent planes at intersected points in the second view. In order to associate a tangent plane Π_2 in the second view to a point p_1 in the first view, the intersection point p_2 between a line segment normal to point p_1 and the surface formed by the second view is computed. The plane Π_2 tangent to point p_2 is then associated to point p_1, as depicted in Fig. 1. The distance d between plane Π_2 and point p_1 is the quantity that is minimized by the refinement process. The algorithm iteratively computes ΔT, an incremental transformation matrix, using a linear least-squares technique. A number of improvements to the basic algorithm were implemented [8]. They are summarized in the following.

2.1 Choosing Control Points

For the refinement process, Chen and Medioni considered only points located in smooth areas. In order to verify the local smoothness of the surface, the residual standard deviation resulting from the least-squares fitting of a plane on the neighborhood of a pixel was considered. In our implementation, we explicitly extract and discard discontinuity points since these points tend to be rather noisy in range images. This way, one may confidently fit a local surface on the remaining points as needed by the algorithm. Tangent planes and normal vectors are then computed by fitting a least-squares plane on each candidate control point and its eight nearest neighbors. A more robust fitting method could have been used but it proved unnecessary given the relatively small amount of sensor noise in our images. A surface visibility test is also applied at each point p_1 in view 1 to verify that it is visible in view 2; when the dot product between the normal vector at p_1 and the orientation vector of view 2 is negative, point p_1 is not used as a control point in the refinement process.

2.2 Evaluating the Convergence

The convergence of the iterative algorithm is obtained when matrix ΔT tends toward an identity matrix. At each iteration, C, the sum of the squared elements of the difference matrix between ΔT and the identity matrix $C = \sum \left(\Delta T_{i,j} - I_{i,j}\right)^2$, is computed. The C measure is only used to monitor the convergence of the algorithm. It has no simple relationship with the registration error. The distance histograms (see Section 4) are more appropriate to evaluate the latter. The process has converged once C becomes smaller than a given threshold. In our experiments, we typically use a threshold value of 0.00000001 for C.

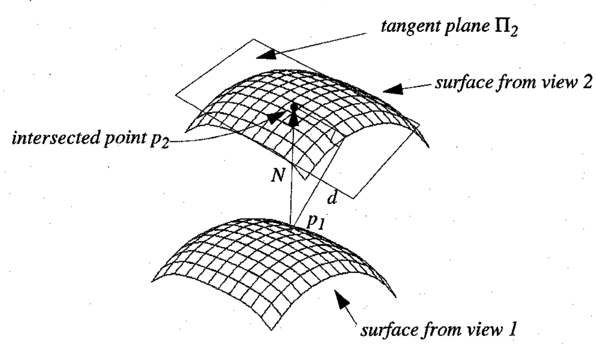

Fig. 1. Point p_2 corresponds to the intersection between a line segment normal to point p_1 and the surface of view 2. The refinement algorithm minimizes the distance d between the tangent plane at p_2 and point p_1.

2.3 Finding the Corresponding Tangent Planes

In order to find the tangent plane of view 2 corresponding to a point p_1 in view 1, Chen and Medioni select the discrete grid point p2 in view 2 nearest to the intersection between the surface of view 2 and a 3-D line segment N normal to p_1 (see Fig. 1). In order to improve the accuracy of the sought registration, it would be preferable to interpolate this intersection rather than computing an approximation. Computing the intersection between a line segment and a surface in three-dimensional space may be extremely slow. However, since both images are represented by square planar parametric grids, a simpler 2-D parametric approach was devised and implemented, as described in [8].

2.4 Computing the Incremental Transformation

The rigid transformation matrices used in this work are of the Euler form. Since the initial registration transformation is assumed to be a good approximation of the true registration, the rotation angles of each incremental transformation matrix ΔT is obtained using a small angles approximation [2]. The sought matrix ΔT is the one which minimizes the sum of the distances between a list of points from view 1 and their corresponding tangent planes in view 2. It is computed using a linear least-squares technique.

3 Registration of a Set of Range Views

This section describes a refinement technique to register a set of range views which is based on the above refinement technique for a pair of range views.

3.1 A Network of Range Views

Let us assume a set of N range views $\{V_1, V_2, ..., V_N\}$ of a given scene. In order to integrate the surface information provided by the N range views, it is required to define the reference frame of one of these views as being the world reference frame into which the integrated surface model will be described. To each range view is associated a node in a network of views. An arc between two nodes stores the inter-frame transformation between the two associated views. All range views are interconnected by this network of inter-frame transformations; that is, for any pair of nodes, at least one path joins the two nodes. A path of inter-frame transformations in a network of views consists in a sequence of matrix multiplications. If the transformation matrices between the views are inexact, such a sequence of multiplications may result in an accumulation of registration errors. It is thus of primary importance to use a network topology that minimizes the lengths of the paths between the nodes.

Fig. 2a shows a typical linear network of range views corresponding to an object rotated on a turntable in front of a fixed range finder. A star-shaped network is depicted in Fig. 2b. Four range views are linked to a common range view which constitutes a central reference frame. Finally, Fig. 2c represents a more general case where several paths may link two range views. The network of Fig. 2b has an interesting topology since any pair of views may be linked by multiplying at most two matrices. Furthermore, only $N - 1$ inter-frame transformation matrices need to be refined with such a network since all views are singly connected. This topology has been selected in designing our registration algorithm for a set of range views. Range view sets corresponding to other network topologies will be converted into the star-shaped topology before refining the registration transformations. The view corresponding to the central node in the network will be selected as the one for which all other views can be transformed using the smallest number of matrix multiplications. In Figs. 2a, 2b, and 2c, the third view would be used to define the world reference frame.

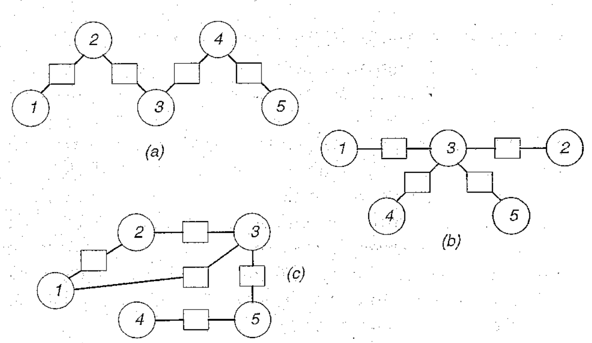

Fig. 2. Three different types of network topologies. Each link joining two nodes represents a transformation matrix between two range views.

A network of range views is well-balanced when

1) the registration error is similar for all transformation matrices, as to be expected, for instance, when a single moving camera is used for all views, and
2) the transformation matrix between any two views is uniquely defined regardless of the path chosen to link the views.

Any process to refine the registration of a set of views should lead to a well-balanced network of views. For the star-shaped topology, the second condition is necessarily met since there exists only one equivalent path between any two range views. Hence, minimizing the error on the $N - 1$ transformation matrices of a star-shaped network directly yields a well-balanced network of views. This result is conditional to having the same distribution of acquisition errors in all images, as is the case when a unique range finder is used to capture the images.

3.2 Sequential Registration of a Set of Views

The approach described in [6] to refine the registration of a set of range views combines registration and integration. It is proposed to process the images sequentially, that is to add images one by one to the integrated model. The registration of an image is refined just before adding it to the integrated model. For example, let us consider the processing of a set of four range views. The first two views are first registered, then integrated. Thereafter, the third view is registered with the integrated model built from two views and added to it. Finally, the fourth view is registered and added to the integrated model. In such an approach, the registration of an image does not change once it has been added to the integrated model. However, it is possible that a following view brings information that could have improved the registration of previously processed views. Such a sequential algorithm is very unlikely to result into a well-balanced network of views.

3.3 Obtaining a Well-Balanced Network of Views

In order to equally distribute the registration errors in all transformation matrices, the network of N views $\{V_1, V_2, ..., V_N\}$ must be considered as a whole and all registration errors must be minimized simultaneously. The developed algorithm is a generalization of the modified inter-frame registration method [8] which was summarized in Section 2.

The network defined by the initial transformations, typically of the type shown in Fig. 2a for a turntable imaging setup, is first converted to a star-shaped topology. The central view V_c defines the world reference frame and its transformation matrix remains unchanged throughout the refinement process. The algorithm computes $N - 1$ incremental transformation matrices at each iteration k, i.e., one for each of the non-central views. In doing so, the points of view V_i, for which the incremental transformation is computed, must be transformed into the reference frames of all other non-central views, each time using only two transformation matrices. The first matrix $M_{i,k-1}$ transforms the reference frame of V_i into the central reference frame of view V_c and the second matrix $M^{-1}_{j,k-1}$ is applied to reach the reference frame of the second non-central view V_j. Once the points of V_i are expressed in the reference frame of V_j, the tangent planes of V_j corresponding to points of V_i are obtained for this iteration. This is repeated for all non-central views and the complete set of points and corresponding tangent planes are used to compute an incremental transformation matrix using the same linear least-squares technique as in Section 2. The convergence of the process is attained when the incremental matrices of all N views are close to the identity matrix. The pseudo-code version of this algorithm appears in [8].

The time needed to refine a transformation matrix is proportional to the product of the number of surface points and the number of views. For a set of eight views comprising on average 10,000 surface points each, it takes about 30 seconds on a SPARC-10 workstation to go through one iteration of the refinement process. The process converged after 30 minutes of computation, the largest part being used for intersection computations. This figure is more than two orders of magnitude better than another recent global registration method [4]. Besides, a great amount of parallelism could be achieved in the refinement algorithm since all surface points may be processed simultaneously and the interactions are limited to local neighborhoods.

4 Experimental Results

The multi-view registration algorithm has been applied to several sets of actual range views of complex multi-part[2] objects with holes. Three typical results are reported here. The processed objects are a teapot, a toy soldier, and an elephant-shaped trinket. The first two objects were placed in turn on a calibrated turntable in front of a fixed synchronized scanners range finder [12]. An initial approximation of the registration was obtained by reading the rotation angle of the turntable and by calibrating its center of rotation. The views of the third object were obtained by moving the camera freely in space. The initial transformations were estimated using a coarse interactive software alignment tool.

Eight range views of the teapot are shown in Fig. 3. The number in the view name represents the angle of rotation on the turntable. These angles are at 0, 30, 60, 90, 120, 180, 240, and 300 degrees, respectively. The initial approximation of the registration is computed from these numbers and the calibrated center of rotation. The view $t120$ is chosen as the central reference frame of the star-shaped network. For each view V_i but $t120$, a pair of histograms of the distances between the points of V_i and their corresponding tangent planes in all other views are shown in Fig. 4. The left-hand histogram of each of the seven pairs is computed from the initial approximation of the registration, that is before refinement. One may observe that the different registration errors of each view generate several distinct peaks in the histogram. The right-hand histogram of each pair depicts the point-tangent plane distances after the refinement process. The distribution of the distances has clearly improved. The mean of the histogram is always centered near zero while the standard deviation has been dramatically reduced. Typical numbers are a few tenths of a millimeter for the absolute value of the average distance before refinement, and less than 10 micrometers after refinement. Typical standard deviation is about 1 mm before refinement with an order of magnitude improvement after refinement.

Fig. 5 clearly demonstrates that the minimization of the registration error is essential to remove the misalignment artefacts in the view integration process [14]. In Fig. 5a, the integrated model was built from the initial approximation of the registration while in Fig. 5b, it was computed from the refined registration. One can clearly observe the effect of the registration errors in the shaded rendering of the triangulated model presented in Fig. 5a. The misalignment of the range views introduces some noise in the integrated model. These artifacts are no more visible in the model of Fig. 5b.

Eight range views of a toy soldier and 24 views of an elephant-shaped trinket were also registered. A rendered view of each of the two integrated models built from the refined registrations is shown in Fig. 6a and Fig. 6b, respectively. Again, no registration error effects are visible in the resulting integrated models. The elephant result is particularly interesting given the high complexity of the object shape, the very large number of views taken from arbitrary true 6 d.o.f. positions around the object, the absence of visible errors in the integrated model, and the near-perfect shape of all 23 final distance histograms, as shown in Fig. 6c.

Tables 1 and 2 present the first two moments of the point-tangent plane distance distributions for the teapot and the toy soldier experiments, respectively. In all cases, the refinement process reduces the standard deviation of the point-tangent plane distance distributions. The effect is less dramatic with the toy soldier since the initial registration is already quite good. In the toy soldier experiment, one may notice that the mean of view $s300$ is nearer to zero before than after the refinement process has occurred. This is due to the fact that the refinement process minimizes the global registration error. The difference is negligible, though, given the precision of the data. Besides, the toy soldier network of views is much more well-balanced after the registration has been refined.

2. Even though no clear definition of a multi-part object is provided here, one may think of objects with significant possible self-occlusions; in particular, objects with one or more holes (non-zero genus) are multi-part. Articulated objects are also typically multi-part.

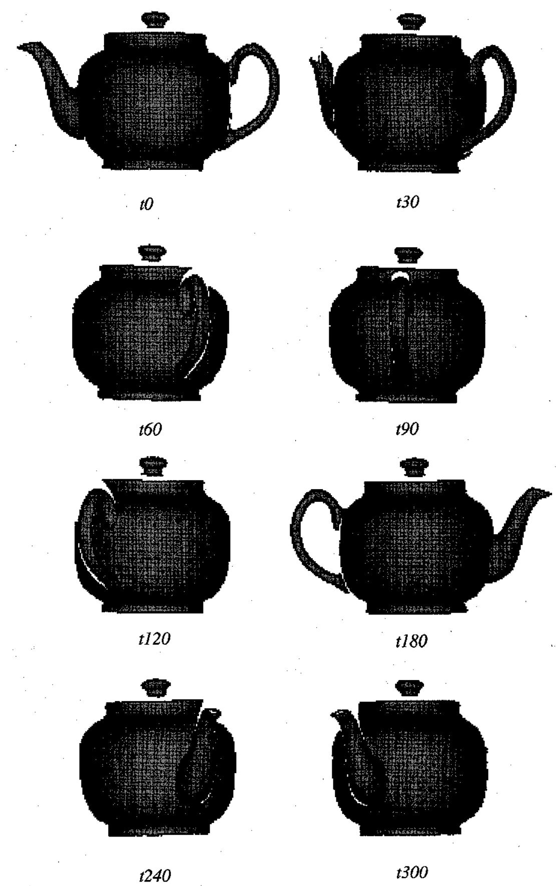

Fig. 3. Eight range views of a teapot. The number in the view name represents the angle of rotation in degrees on a turntable.

5 Discussion

A new method has been presented to refine the registration transformations for a set of range views. The refinement may be viewed as an optimization process whose goal is to minimize the registration error of a set of 3-D views of an object. Our method realizes this optimization by minimizing the distance between measured points of each view and the interpolated surfaces from all other views. This method decomposes the registration and integration steps. It minimizes the registration errors of all views simultaneously, leading to a well-balanced network of views in which the registration errors are all low and of the same order of magnitude, i.e., equally distributed. This approach is to be effective only if the surface itself offers enough constraints, i.e., only if there is a single way to paste the views together. Extensive experimentation, with some typical results shown in this correspondence, have demonstrated that the method performs well both quantitatively and qualitatively. It improves the calibrated registrations and the integrated model for a wide variety of complex 3-D objects with views

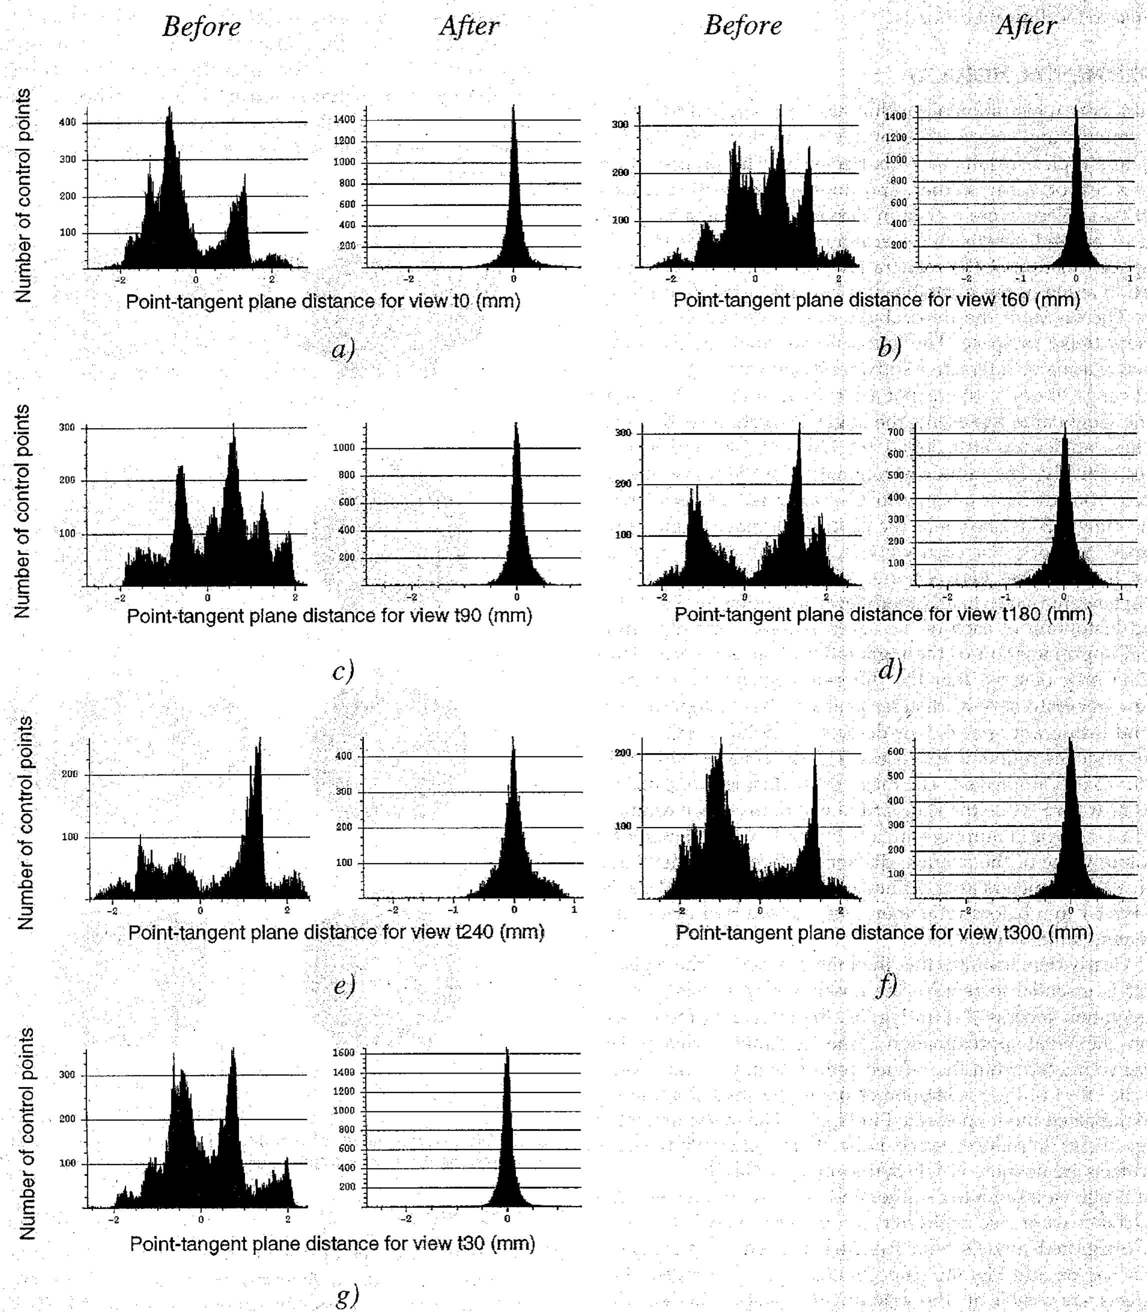

Fig. 4. Histograms of the distances between points from a given view of the teapot and their corresponding tangent planes in all other views before and after the refinement process. For all histograms, the Y-axis represents the number of control points. For all figure items, the left histogram is computed from the initial approximation of the registration while the right histogram shows the distribution of the distances after the refinement process has converged. View *t*120 is central reference frame of the star-shaped network. (a) view *t*0, (b) view *t*60, (c) view *t*90, (d) view *t*180, (e) view *t*240, (f) view *t*300, and (g) view *t*30.

related by unconstrained six-parameter transformations. For range views of complex objects, it was shown that the acquisition, or sensor, noise is the only remaining contribution to the final global error distributions, as visible in the final shape of the distance histograms. The elimination of the registration errors allows an improvement of the accuracy of the integrated model over the original data provided by the sensor. In all cases, the application of the method leads to a well-balanced network of views.

An obvious limitation of any ICP-like surface-based registration algorithm is that objects of enough structural complexity are required [3] and sufficient overlap must be present between pairs of views. More precisely, the common surfaces between any pair of views must take care of all degrees of freedom of the computed transformation. For instance, matching one or two planar patches common to a pair of views does not produce a unique transformation. This problem is generally not present when matching surfaces of complex objects with no global symmetry since all pairs of views are likely to share a large number of patches with various orientations.

A less obvious but important side-effect is that a large patch common to a pair of views may make difficult the accurate registration of small close-by patches. The reason is that the least-squares computed incremental parameter vector (and transformation matrix) is an average over all matches. The contribution of a small unregistered patch may thus end up being negligible with respect to the wide majority of contributions arising from a single partially constraining patch that appear to be registered. A possible solution to this second problem could be a final alignment of the surfaces based on their occluding contours [10]. Unfortunately, this is not possible with curved surfaces since the shape of the contours is viewpoint-dependent. Moreover, even planar objects present difficulties since the range data provided by the sensor are typically less reliable in the neighborhood of surface discontinuities.

TABLE 1
REGISTRATION OF A SET OF RANGE VIEWS OF A TEAPOT

View name	Mean and standard deviation of the distance in mm between points and tangent planes	
	Before refinement	After refinement
t0	$\mu = -0.253351$ $\sigma = 1.02047$	$\mu = 0.000046$ $\sigma = 0.241934$
t60	$\mu = 0.189274$ $\sigma = 0.896068$	$\mu = 0.003865$ $\sigma = 0.183611$
t90	$\mu = 0.201141$ $\sigma = 0.961355$	$\mu = 0.008138$ $\sigma = 0.207018$
t180	$\mu = 0.356614$ $\sigma = 1.22248$	$\mu = 0.006054$ $\sigma = 0.278336$
t240	$\mu = 0.342048$ $\sigma = 1.20350$	$\mu = 0.015447$ $\sigma = 0.303572$
t300	$\mu = -0.352423$ $\sigma = 1.17318$	$\mu = -0.002104$ $\sigma = 0.320406$
t30	$\mu = 0.0913468$ $\sigma = 0.888380$	$\mu = 0.001569$ $\sigma = 0.176339$

TABLE 2
REGISTRATION OF A SET OF RANGE VIEWS OF A TOY SOLDIER

View name	Mean and standard deviation of the distance in mm between points and tangent planes	
	Before refinement	After refinement
s0	$\mu = -0.032932$ $\sigma = 0.401119$	$\mu = -0.002850$ $\sigma = 0.354298$
s60	$\mu = -0.005481$ $\sigma = 0.281938$	$\mu = -0.004468$ $\sigma = 0.220575$
s90	$\mu = -0.025074$ $\sigma = 0.247873$	$\mu = -0.008166$ $\sigma = 0.183995$
s120	$\mu = 0.030175$ $\sigma = 0.269683$	$\mu = -0.005828$ $\sigma = 0.200821$
s240	$\mu = 0.010824$ $\sigma = 0.261182$	$\mu = -0.009330$ $\sigma = 0.224080$
s270	$\mu = 0.030844$ $\sigma = 0.249737$	$\mu = -0.007530$ $\sigma = 0.212676$
s300	$\mu = 0.000080$ $\sigma = 0.297233$	$\mu = -0.003794$ $\sigma = 0.259099$

In order to illustrate these types of problems, sets of unrestricted views of a synthetic scene representing a power line (long cylinder), a beam (polyhedra), and an isolator (hemisphere and short cylinder) were generated and input to the registration algorithm. As expected, the algorithm usually got caught in local minima when started from arbitrary positions in transformation space. These local minima are usually very close to the actual solutions. Still, the histogram is not symmetrical with respect to the zero axis and significant peaks result from image sampling noise and the unresolved degrees of freedom in the registration of the synthetic views.

A possible approach around these problems is to extend our multi-view registration method to include constraints from matched fitted surfaces. These surfaces would be of simple geometric shape, usually with some form of symmetry, e.g., planes, spheres, or cylinders. The parameters of the fitted surfaces would constrain the solution to a sub-space of the overall transformation space. A general multi-view registration technique based on the above multi-view integration technique extended to include constraints from matched fitted surfaces is presently under development.

a)

b)

Fig. 5. (a) Integrated model of the teapot built from the range views of Fig. 3 using the calibrated registration without refinement. Arrow shows the discontinuity artifacts resulting from the misalignment of the integrated views. (b) Integrated model built after refinement of the registration.

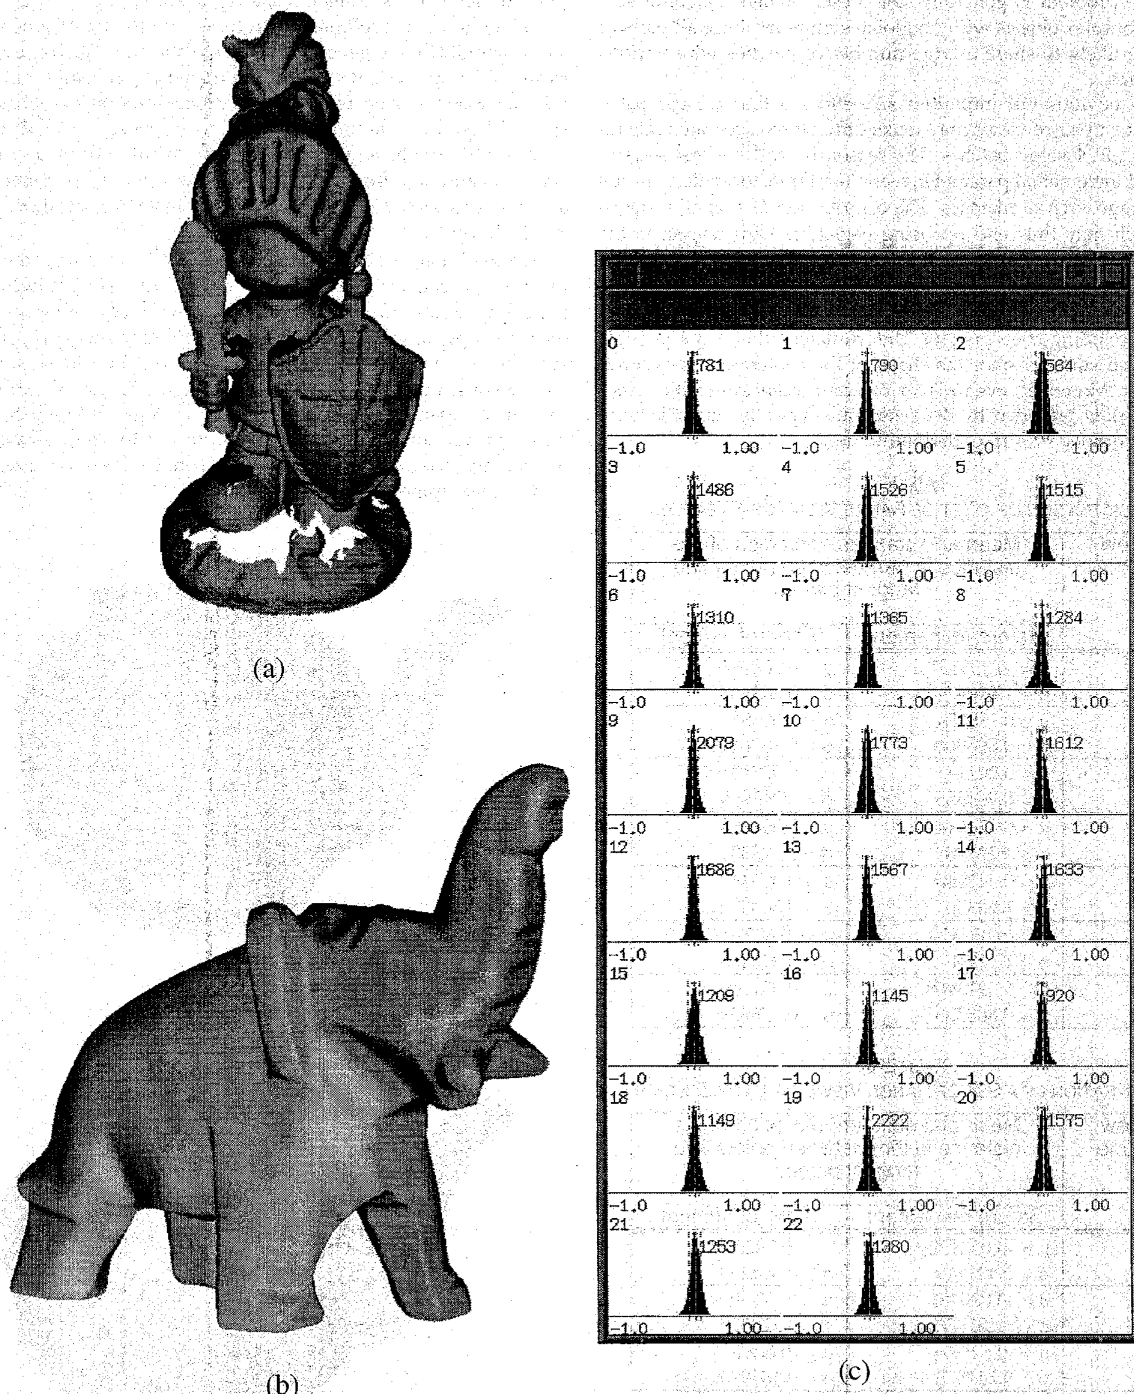

Fig. 6. Integrated models of a toy soldier (a) and an elephant-shaped trinket (b) built from the refined registrations. Final distance histograms (c) for the elephant experiment. Height of the central peak in number of points, mean (continuous) and standard deviation (dotted) vertical lines. Horizontal scale has range $-T_{snt}$ to $+T_{snt}$ (T_{snt}: spatial neighborhood threshold).

The results presented here were obtained using almost all image points (discarding discontinuity points) as control points in the minimization process. In order to improve the performance of our method, we intend to experiment with better control point selection schemes. Preliminary experiments have already shown us that even a simple uniform sub-sampling of the image could improve the performance without affecting much the quality of the results. More extensive experiments are needed in order to evaluate all pertinent parameters such as the number of iterations before convergence. A parallel implementation would also improve the performance of our multi-view registration algorithm. This is also on our agenda of future research.

ACKNOWLEDGMENT

The authors would like to thank Messrs. M. Rioux and J. Domey of the National Research Council of Canada (NRCC) for providing the range images used in this work. The authors are members of the Institute for Robotics and Intelligent Systems (IRIS) and wish to acknowledge the support of the Networks of Centres of Excellence program of the Government of Canada, the Natural Sciences and Engineering Research Council, and the participation of PRECARN Associates Inc.

REFERENCES

[1] J.-A. Beraldin, M. Rioux, F. Blais, G. Godin, and R. Baribeau, "Model-Based Calibration of a Range Camera," *Proc. 11th Intl. Conf. Pattern Recognition*, pp. 163-167, The Hague, The Netherlands, 1992.

[2] R. Bergevin, D. Laurendeau, and D. Poussart, "Registering Range Views of Multipart Objects," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 1-16, 1995.

[3] P.J. Besl and N.D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239-256, Feb. 1992.

[4] G. Blais and M.D. Levine, "Registering Multiview Range Data to Create 3D Computer Objects," TR-CIM-93-16, Center for Intelligent Machines, McGill Univ., 1993.

[5] P. Boulanger, G. Godin, and M. Rioux, "Applications of 3-D Active Vision to Rapid Prototyping and Reverse Engineering," *Proc. Third Int'l Conf. Rapid Prototyping*, pp. 213-223, Dayton, Ohio, June 1992.

[6] Y. Chen and G. Medioni, "Object Modeling by Registration of Multiple Range Views," *Proc. IEEE Int'l Conf. Robotics and Automation*, pp. 2,724-2,729, Sacramento, Calif., 1991.

[7] F.P. Ferrie and M.D. Levine, "Integrating Information From Multiple Views," *Proc. IEEE CS Workshop Computer Vision*, pp. 117-122, Miami, Fla., 1987.

[8] H. Gagnon, M. Soucy, R. Bergevin, and D. Laurendeau, "Registration of Multiple Range Views for Automatic 3-D Model Building," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Seattle, Wash., June 1994.

[9] B. Kamgar-Parsi, J.L. Jones, and A. Rosenfeld, "Registration of Multiple Overlapping Range Images: Scenes Without Distinctive Features," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, San Diego, Calif., pp. 282-290, 1989.

[10] E. Kishon and H. Wolfson, "3-D Curve Matching," *Proc. Workshop on Spatial Reasoning and Multi-Sensor Fusion*, pp. 250-261, St. Charles, Ill., 1987.

[11] M. Potmesil, "Generating Models of Solid Objects by Matching 3D Surface Segments," *Proc. Eighth IJCAI*, pp. 1,089-1,093, Karlsruhe, Germany, 1983.

[12] M. Rioux, "Laser Range Finder Based on Synchronized Scanners," *Applied Optics*, vol. 23, no. 21, pp. 3,837-3,844, Nov. 1984.

[13] J.J. Rodriguez and J.K. Aggarwal, "Navigation Using Image Sequence Analysis and 3-D Terrain Matching," *Proc. IEEE Workshop on Interpretation of 3-D Scenes*, pp. 200-207, Austin, Texas, Nov. 1989.

[14] M. Soucy and D. Laurendeau, "A General Surface Approach to the Integration of a Set of Range Views," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 4, pp. 344-358, Apr. 1995.

[15] M. Soucy and D. Laurendeau, "A Dynamic Integration Algorithm to Model Surfaces From Multiple Range Views," *Machine Vision and Applications*, vol. 8, pp. 53-62, 1995.

[16] M. Soucy, "Modeling the Surface of 3-D Objects Using Multiple Range Views," PhD thesis, Laval Univ., Québec, Canada (in French), 1992.

[17] G. Turk and M. Levoy, "Zippered Polygon Meshes From Range Images," *Proc. SIGGRAPH '94*, pp. 311-318, 1994.

[18] Z. Zhang, "Iterative Point Matching for Registration of Free-Form Curves and Surfaces," *Int'l J. Computer Vision*, vol. 13, no. 2, pp. 119-152, 1994.