Computer Vision and Image Understanding **87,** 27–38 (2002)
doi:10.1006/cviu.2002.0980

Differential Invariants as the Base of Triangulated Surface Registration

Pavel Krsek, Tomáš Pajdla, and Václav Hlaváč[1]

Czech Technical University, Faculty of Electrical Engineering, Center for Applied Cybernetics, 121 35 Prague 2, Karlovo náměstí 13, Czech Republic
E-mail: krsek@cmp.felk.cvut.cz

Received August 31, 2001; accepted September 6, 2002

The paper addresses the problem of 3D model reconstruction from overlapping triangulated range images. A technique for automatic matching of curved freeform surfaces exploiting curvilinear differential structures of the surfaces is presented. We propose a hybrid registration algorithm that combines advantages of working with small amounts of interest points (to attain computational speed), estimates the Euclidean transform matching both surfaces, and uses all available points and the iterative closest reciprocal point algorithm to refine the estimate and finally match surfaces (to attain high precision, good initial estimation avoids local minima). The method works in a bottom-up manner using the hierarchy: points → differential structures (i.e., curvilinear line segments) → surface. The registration is automatic. The only parameter set by the user is the required level of mean curvature. The approach is demonstrated through examples.

Key Words: differential parameters; registration; 3D surface; differential invariants; invariant structures; range image.

1. INTRODUCTION

This contribution tackles the problem of registering two range images of the same 3D scene captured from two different viewpoints. Such a task is of practical importance, e.g., in reverse engineering (related to CAD) in the automotive industry or elsewhere. A 3D CAD model is created from a 3D specimen of an object (e.g., a design made of clay) by capturing a set of overlapping range images that are later transformed into one coordinate system. The 3D model is then created by fusing all range images covering the object's surface.

[1] P. Krsek and V. Hlaváč were supported by the Czech Ministry of Education under Project LN00B096. T. Pajdla was supported by the Czech Ministry of Education under Project MSMT KONTAKT 2001/09 and by the Grant Agency of the Czech Republic under Projects GACR 102/01/0971, GACR 102/00/1679, ME412.

1077-3142/02 $35.00

Individual overlapping range images are related by an unknown Euclidean transformation consisting of a rotation and a translation which sets two overlapping range images into correspondence. Finding the transformation is called *registration of range images*. Sometimes, the Euclidean transformation relating two views or its rough estimate is provided by the range image capturing equipment, e.g., from the turntable holding the object. However, we assume that this information is not available.

This paper suggests a novel method for registering a pair of range images based on significant curvilinear characteristics of the surface. It is assumed that all information needed to find the Euclidean transformation is to be found in the data. Features based on the differential geometry of the surface are used. The registration process is fully automatic. The paper is organized as follows: Section 2 places the solved problem into a wider context and discusses the work of others. Section 3 describes the proposed method. Section 4 demonstrates the behavior of the method on real examples. Section 5 concludes the paper.

2. REGISTRATION ALGORITHMS AND THE STATE-OF-THE-ART

The common and plausible approach to the registration of two range images is to minimize the criterion aggregating distances between pairs of the closest points from overlapping surface patches in the first and the second image. The iterative closest points (ICP) algorithm [1] is a widely used technique. The distance criterion is defined as the sum of distances between two closest points, one from surface S and the second from surface S'. This relationship may not be symmetric. The improved version is called iterative closest reciprocal points ICRP [2] and explores the natural symmetry of the relationship "being the closest." For a pair of points to be considered, the following must hold: For the point $x \in S$, the closest point to it is $x' \in S'$ and vice versa. Points having no counterparts in both directions are ruled out. Such reciprocal relationships are symmetric. The ICP and ICRP algorithms provide precise matches. There are however two major problems with them. The first is that due to the nonconvex character of the optimization task the process usually becomes trapped in a local minima. The second problem is related to the time of computation. The number of points on both surfaces is often in the order of hundreds of thousands and the ICP or ICRP registration is slow if a good initial guess of the Euclidean transformation is not available.

There are several papers including this one that follows the hypothesize and verify paradigm to find good initial transformation for the ICPR algorithm. The hypothesis is the estimation about the Euclidean transformation bringing two surfaces into correspondence. In [3], all points on the overlapped surface are used to generate and test a huge number of hypotheses. In [4], a stochastic optimization was performed. Such approaches become rather impractical if a hundred thousand points have to be registered.

The need to reduce the computational effort requires the Euclidean transformation to be estimated using significantly fewer matching points. A reduced set of surface points that contains the required information is handled similarly to that in stereo vision, motion analysis, or object recognition. These points of interest should be located in a range image and the process has to be invariant to the Euclidean transformation.

Our aim is to perform the initial registration automatically using intrinsic differential structures of the surface. An example of an algorithm exploring differential geometry features (curvatures) was presented in [5]. The approach was called "simplex angle image."

Curvatures of both object surfaces are calculated. The surface is circumscribed by a unit sphere and it is elastically shrunk to the object. The one-to-one correspondence between surface points (and its curvatures) allows registration in 2D only. This method estimates rotation and can be used only for objects topologically equivalent to the sphere.

Another application of differential invariants employs registration to tomographic images [6]. Surface crest lines are estimated by a convolution with Gaussian kernels. Extremal points on crest lines constitute points of interest. They are attributed by four differential parameters. Triples of points with the same attributes are used to estimate the Euclidean transformation.

3. PROPOSED REGISTRATION ALGORITHM

Let us begin with the task formulation. The input is given by two overlapping range images captured by a range finder. We further assume that the triangulation of the measured surface was established. For structured light range finders, this is a natural assumption. The aim of the algorithm is to bring the overlapping surface patches into correspondence. The outputs are the parameters of the Euclidean transformation that align two range images. The approach follows the paradigm of random sample consensus [7] in a bottom-up manner. Sampling is performed in the space of the six parameters of the Euclidean transformation sought, determined by three pairs of corresponding points in a general position.

First, the mean curvature is computed for each point of both range images. Level curves of constant mean curvature constitute a stable feature that is invariant to Euclidean transformation. The curves corresponding to zero mean curvature are often used. We call them *intrinsic local differential structures*. In the ideal case, when a complete closed surface has been measured, which fulfills the requirement of the second fundamental form (without singular points), the intrinsic local differential structures are closed curves. However, in real situations, the intrinsic local differential structures correspond to curvilinear line segments. Some of them are closed but many are open. Second, the aim is to find points of interest on intrinsic local differential structures. For each individual range image, all possible pairings between intrinsic local differential structures are established. In each pairing, the shortest line segment (connector) connecting them is found. The endpoints of connectors define a pair of invariant points—these are points of interest. This pair of points of interest is attributed by the length of the connector and this information is later used when reducing the number of hypotheses.

Two potentially corresponding pairs of connectors from different range images are selected. Information from four interest points is used. The pair of connectors generates the initial estimate of Euclidean transformation, which constitutes the hypothesis. The hypothesis is submitted to the first verification step. The quality of correspondence between all points of intrinsic differential structures is checked. If the estimate does not satisfy a prescribed quality condition, then the hypothesis is rejected. The next best hypothesis is selected and the first verification step is repeated. The hypothesis that passed the first verification is refined by applying the ICRP optimization algorithm using all points on intrinsic differential structures. The quality of the transformation estimate is checked in the second verification step. The correspondence between the whole matched surfaces is tested.

The hypothesis that fulfills the second verification step is refined by the second ICRP optimization step, in which points on whole surfaces are involved. In the final step of the

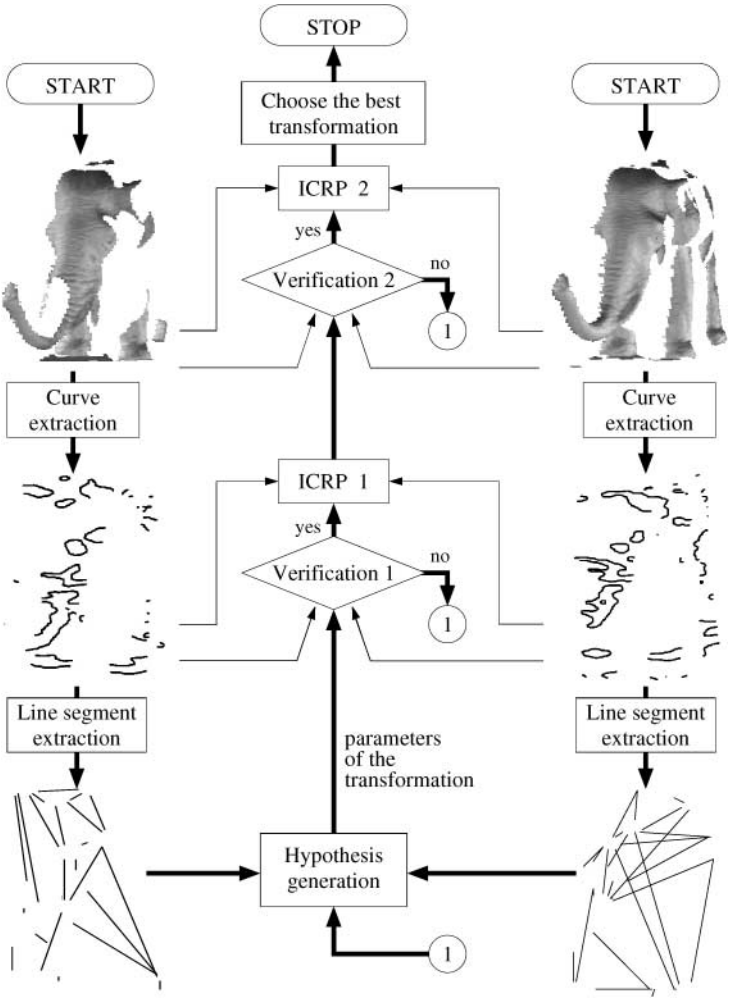

FIG. 1. Flow chart of the proposed registration method.

algorithm, the best Euclidean transformation is selected according to the minimal average distance of reciprocal points. The described method is illustrated as a flow chart in Fig. 1. Thick arrows denote the control flow in the algorithm. Thin arrows represent data involved in a particular step of the method explained in the following four sections.

3.1. Local Differential Parameter Estimation

Many algorithms estimating local differential parameters from range images have been applied and published. The algorithms provide the value of local differential parameters for each point of the input range image. We tested several of them and chose the best algorithm suited to our conditions [8].

We use a full quadric fitting algorithm to estimate locally differential parameters of the surface [8]. The algorithm is based on the local approximation of a surface by a paraboloid in normal position with the axis parallel to the normal direction. The normal direction is estimated from the local approximation by a plane. Transformed points (estimated normal

parallel to axis z) are approximated by the paraboloid in a general position: $z = ax^2 + bxy + cy^2 + dx + ey + f$, where a, b, c, d, e, f are parameters of the paraboloid and x, y, z are coordinates. The estimation of differential parameters is computed from parameters of the paraboloid [9].

Noise in the input data decreases the quality of our estimate. The influence of noise can be alleviated by enlarging the approximation neighborhood while dramatically increasing computation time. Our choice is to use a 2D filter on estimated parameters [10]. We use the median filtration of mean curvature, which preserves surface patches with constant curvature.

3.2. Intrinsic Local Differential Structure and Points of Interest on Them

We chose curves of constant mean curvature as intrinsic local differential structures. Mean curvature $H = (k_1 + k_2)/2$ and Gaussian curvature $K = k_1 k_2$ are first-order differential parameters which are invariant to resampling and Euclidean transformation. Our experiments show [8] that the influence of stochastic errors in measurement on estimation of mean curvature is linear.

Having computed mean curvature at each point of the input range images, the level curves of the constant value of mean curvature, H, have to be found. Available to us are the calculated values of H in a grid induced by sampling. Positions on the surface with mean curvature H are obtained by linearly interpolating between available samples using a planar patch. The result is given in a list of (x, y, z) coordinates for a curvilinear line segment (Fig. 2). Two arrows (one pointing up and the other down) represent the samples of the mean curvature. The dashed line connecting the arrows stands for the linear interpolation that gives the position of a single point $\mathbf{x}$ on the H level curve. The other points on the zero mean curvature level curve are calculated similarly and they are shown as small filled circles. Discovered points lie on the grid given by the triangulation. The way the points are connected and constitute the curvilinear line segment is also induced by the triangulation. There are some singular configurations in which the surface is locally planar or corresponds to a saddle point. Our method ignores such a point.

Choosing the level of mean curvature is important. The best choice depends on the shape of the registered surfaces and, in general, it is difficult to find it automatically. In our experiments it has been chosen a priori. However, a heuristics for automatic adjustment of the level for certain situations can be proposed.

The next step is to find points of interest on the established intrinsic local differential structures. We first thought that we could determine points of interest as the points with

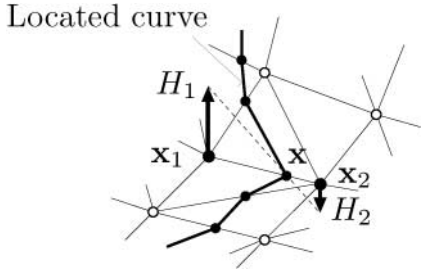

FIG. 2. Extraction of intrinsic invariant structures.

the highest curvature on the curvilinear line segment, similarly as in [2]. However, the experiments show that the extracted points of interest are very sensitive to noise. This is natural because we actually deal with second derivatives along curves which themselves were obtained from second-order derivatives of the surface.

We decided to find more stable structures in the measured range image. We chose the straight line segments corresponding to closest connection between two curvilinear line segments in one range image (connectors). Of course, there are many such pairings. The longer connectors reveal more information.

3.3. Generation of a Hypothesis about the Match

The hypothetical match between two overlapping surfaces given by two range images is sought. The points of interest were already calculated in both input range images. Actually, we have an even richer structure available to us, the connectors. The intention is to use connectors to guess the parameters of the Euclidean transformation. The set from which the hypothesis is going to be selected is given by all possible pairings of connectors from two different range images. To limit this set we adopted a simple heuristics; i.e., we only match connectors of approximately similar length. There is tolerance for registration of connectors. The tolerance is set as twice average distance between connected points on input triangulated surfaces. From this newly created set of hypotheses, the hypotheses are selected according to the length of the connectors—the longer ones first.

3.4. Double Verification and Refinement of the Hypothesis

The hypothesis is submitted for double verification and refinement. The first verification step transforms all intrinsic local invariant structures (curvilinear line segments) from the first range image into the coordinate system of the second one using the hypothesized Euclidean transformation. The degree of match is given by the number of reciprocal points [2] on the structures. Only the hypothesis which has more reciprocal points than the first verification threshold is submitted to the first refinement. Refinement is based on the ICRP algorithm, which explores points of curvilinear line segments. The result of this optimization is a refined transformation that is immediately submitted to the second verification step. The result has to pass a threshold on the minimal number of reciprocal points which are detected on whole range images transformed into one coordinate system by the tested transformation. Transformations which do not pass the threshold are not improved further. The second refinement step is similar to the first one but works with whole range images. The best hypothesis is selected from the list of hypotheses that passed two verifications and refinements. The best hypothesis is the one which has the highest number of reciprocal points after the second refinement or the smallest average distance between reciprocal points if the number of reciprocal points is the same.

Both verifications use a number of reciprocal points thresholds which should be set. Initially both thresholds are set to zero. If the hypothesis passes the verification then the number of its reciprocal points is higher than the threshold. The new threshold is set to 80% of the number of reciprocal points provided by the best hypothesis seen so far. This introduced hysteresis ensures that the potentially better hypothesis is not rejected too early when it can be further successfully improved.

4. EXPERIMENTS AND RESULTS

4.1. Aim of Experiments, Measured Objects, Capturing Device

We will show registrations on two real objects of different complexity. The first object is a curvilinear statuette of an elephant made of glazed ceramic, see Fig. 3a, with semi-opaque surface (dimensions: length 21 cm, width 8 cm, and height 13 cm). The statuette is topologically equivalent to a sphere and its surface is largely nonconvex. Elephant data provided information needed to analyze the properties of the proposed methods. The number of points on a triangulated mesh in one range image is on the order of tens of thousands.

The second object is a skull of an adult human, see Fig. 3b, provided by the Institute of Criminology Prague. Our task is to create a full 3D model of the skull. The criminalists use a software package to add hypothetical soft tissue onto the skull with the aim to generate an approximate 2D intensity picture of a face. The number of points on a triangulated mesh in one range image is on the order of hundreds of thousands.

The range images were captured using the Laser Plane Range Finder constructed in our laboratory [11] (laser diode Lasiris, 5 mW, cylindrical lens, precision 100 micrometers). In the case of the elephant, each range image was captured by translational movement of the camera–laser rig. The elephant was rotated on the turntable to get the second range image, etc. In the case of the skull, the camera–laser rig was fixed and the skull was rotated on the turntable. The result is a cylindrical range image.

4.2. Registration of Two Range Images

Let us demonstrate the proposed registration algorithm on the statuette of the elephant. The extraction of the intrinsic local differential structures is shown in Fig. 4. The points of interest are difficult to display and they are not presented.

The hypotheses generation, verification, and hypothesis refinement are shown in Fig. 5. Figure 5a demonstrates located intrinsic local differential structures in one coordinate system before registration. The curvilinear segments from one range image are shown as black curves and those from the second range image are displayed as gray curves. Figure 5b demonstrates the best initial estimate of the Euclidean transform. Figure 5c exhibits the

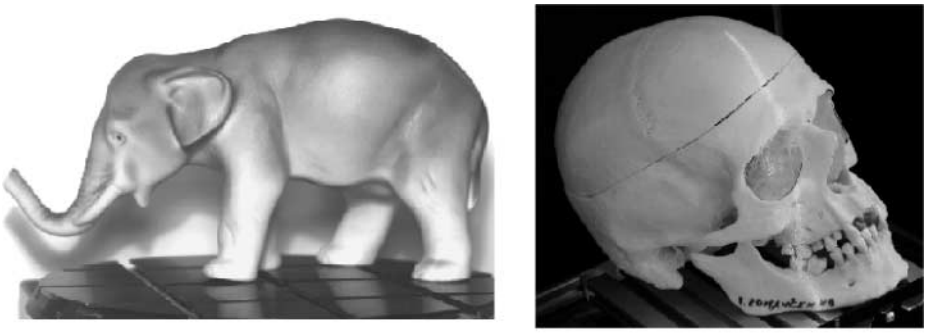

(a) (b)

FIG. 3. Intensity images of objects used in range image registration experiments, (a) the statuette of the elephant from glazed ceramic, (b) the human skull.

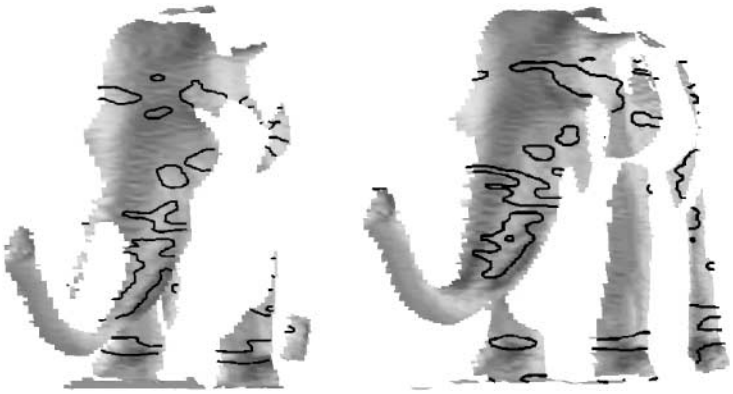

FIG. 4. The range images obtained from two different view directions with superimposed intrinsic local differential structures.

final registration after ICRP refinement on the entire point cloud. The results of registration are shown in Fig. 6.

Table 1 characterizes the performance of the proposed registration method. The first column indicates the step of the algorithm. The labels have the following meaning: IH, initial hypothesis; C, curve registration after the first refinement when points on curvilinear line segments were used; S, surface registration after the second refinement when all surface points were used. The last column gives the number of hypotheses processed in the respective step of the algorithm.

There are four intermediate pairs of columns in Table 1 corresponding to four cases that characterize the way the hypotheses are generated and accepted. Case 1 characterizes the best hypothesis. Case 2 tackles a typical good hypothesis. Case 3 describes a typical rejected hypothesis in the second verification step. Case 4 demonstrates a typical hypothesis rejected in the first verification. The table is vertically divided into two parts. The upper part demonstrates hypotheses generation, verification, and refinement on the level of curvilinear line segments. The lower part shows the performance on the level of whole surfaces. Each of the four cases is represented by numerical values in two columns. The left column (#RP) value gives the number of reciprocal points and characterizes the quality of correspondence. The right column (AD) is the average distance between corresponding reciprocal points.

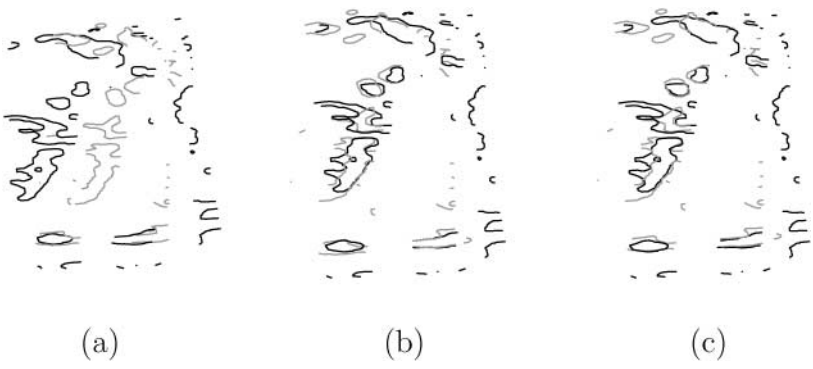

FIG. 5. Intrinsic local differential structures from two views in a common coordinate system. (a) Before registration. (b) The best initial hypothesis. (c) Final registration after refinement on the entire point clouds.

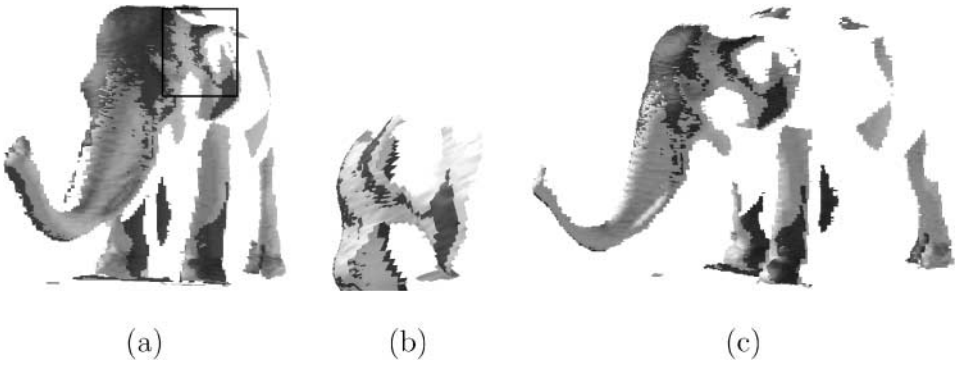

FIG. 6. The entire range image after registration in one coordinate system. Dark surface patches correspond to one range image and the lighter ones correspond to the second range image. (a) View from the front. (b) Detailed view of the elephant's ear. The position in the less detailed image (a) is roughly delimited by a rectangle. (c) View from the back.

We use the average value as it is a relative number that is invariant to the number of involved points. This value provides the information about the distance between the two registered surfaces.

The rightmost column (#H) shows the number of hypotheses tested in that particular step of the algorithm. The number of hypotheses decreases dramatically. The number of reciprocal points should ideally grow in all four cases. This is the case for steps IH and C on the curvilinear segments. In step S, the number of reciprocal points decreases. This phenomenon is caused by the fact that the hypothesis is refined using all surface points and not just points on curvilinear line segments. The bottom part of the table indicates that the number of reciprocal points increases again when all surface points are involved.

Figure 7 helps us to discuss the properties of generated hypotheses. Figure 7a demonstrates the number of reciprocal points, i.e., the quality of match. Notice that the best hypothesis was generated as the 368th. After the hypothesis 900 the quality remains low. This proves that our heuristics selecting hypotheses works. Figure 7b shows the average distance between surface patches attached to individual hypotheses. The first and good hypotheses match rather well.

TABLE 1
Performance Characterization in Selected Steps of Registration Algorithm

	Case 1		Case 2		Case 3		Case 4		
Step	# RP	AD	# RP	AD	# RP	AD	# RP	AD	# H
Detected reciprocal points on curvilinear line segments									
IH	185	3.46	124	7.40	105	12.04	56	54.96	28216
C	205	2.73	201	2.85	140	3.27	—	—	12
S	195	3.25	193	3.26	—	—	—	—	8
Detected reciprocal points on entire surfaces									
C	4086	0.57	3917	0.64	3020	0.73	—	—	12
S	4857	0.34	4855	0.34	—	—	—	—	8

Note. Case 1, The best hypotheses; Case 2, Typical good hypotheses; Case 3, Typical hypotheses rejected by second verification step; Case 4, Typical hypotheses rejected by first verification step.

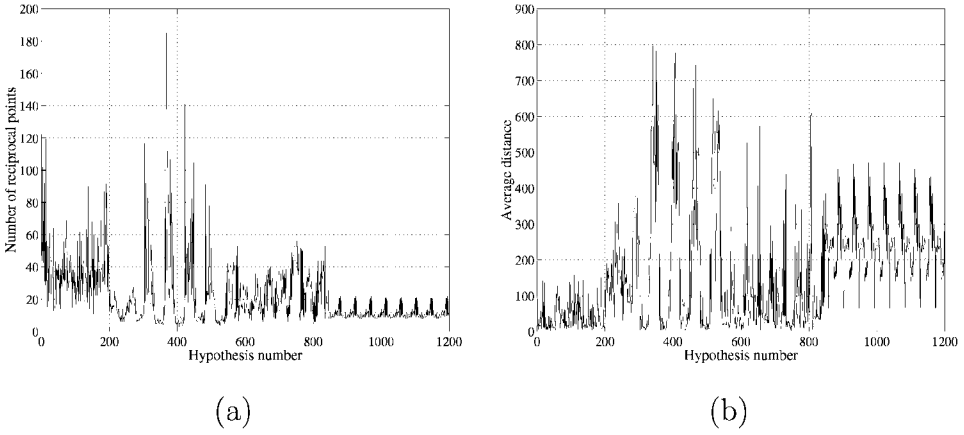

FIG. 7. (a) Number of reciprocal points as a function of hypothesis number. The numbering is given by the order in which hypotheses are generated. (b) Average distance between surfaces as a function of hypothesis number. Only the first 1200 hypotheses are displayed out of 28,000.

Practical reconstruction from range images is demonstrated on the human skull. We have reconstructed five skulls so far. The range images of the skull were obtained as two cylindrical range images. The first range image was measured when the skull was rotated on the turntable around the vertical axis. The second cylindrical range image represents a stripe from the chin to the crown of the skull. The 3D model was composed from these two cylindrical range images (Fig. 8). Registering of the skull is not an easy task. The reason is that curvilinear features are missing in a substantial part of the skull surface, cf. back of the skull. Yet, the proposed algorithm is still able to register such data.

4.3. Registration of Multiple Range Images

The full 3D model is often constructed from more than two range images. The 3D model of the elephant was created from 19 measured range images captured from the statuette. The view circle is covered by 18 views with the step either 15° (10 steps) or 30° (6 steps)

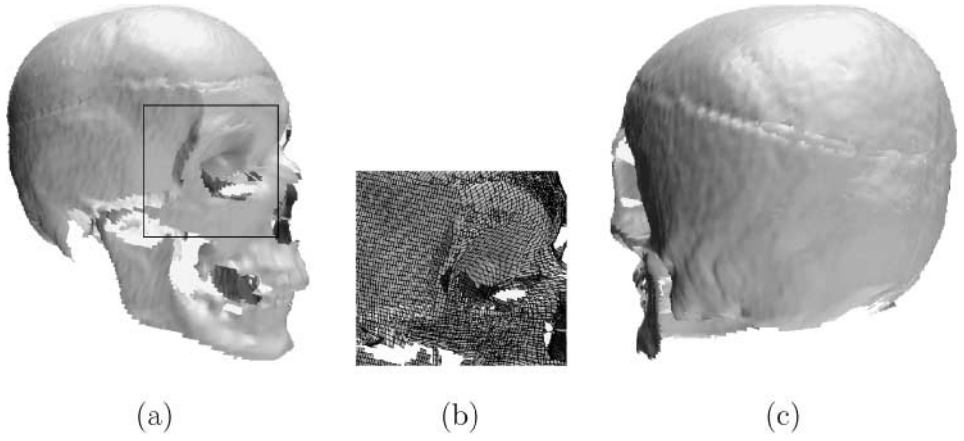

FIG. 8. Skull reconstruction—result of registration. (a) Frontal view. (b) Detail of the reconstructed surface around the right eye socket. The position in the less detailed image. (c) Back view.

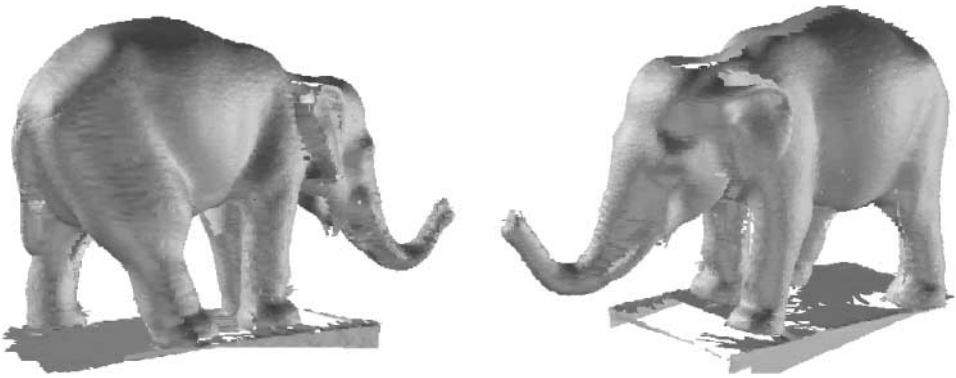

FIG. 9. A full 3D model of the elephant fused from 19 range images.

depending on the complexity of the underlying surface. The 19th view is from the top left and provides missing information about the complicated shape behind the ear of the elephant.

The registration started from the reference range image which is the frontal view of the statuette. The other range images were registered pairwise starting from the reference image to the elephant's tail. The error accumulated. Even so, the obtained results exhibit only tiny artifacts in the tail of the statue. The results are presented in Fig. 9.

Let us discuss the precision of the results. For the 30° steps on the turntable, the method estimated the average rotational angle 30.0° with the standard deviation 0.16°. For the 15° steps on the turntable, the average rotational angle was 15.0° with the standard deviation 0.49°. When all the images were gradually fused into one coordinate frame, the error at the place when the first range image meets the last one was 0.067°. This corresponds approximately to an error in depth equal to 1.2 mm.

5. CONCLUSIONS

A method for automatically registering overlapping range images was proposed and experimentally verified. This allows the creation of a full 3D model of the measured object. The method estimates the Euclidean transformation from curvilinear features on the object from a small amount of points of interest. These points are detected by calculating differential parameters of the surface. A good initial hypothesis avoids the iterative closest reciprocal point algorithm becoming trapped in local minima.

Practical experience with the method was demonstrated on two extensive examples. For more information, see http://cmp.felk.cvut.cz/~krsek/Rireg/.

The advantage of the method is that it copes with very general curved surfaces. However, missing differential structures on the surface lead to failure of the proposed algorithm. This is obvious as surfaces without any features cannot be naturally registered, e.g., sphere, large planar patch. On the other hand, the strategy of hierarchical hypothesis generation and verification can be used for any features, e.g., texture on the surfaces.

REFERENCES

1. P. Besl and N. D. Kay, A method for registration of 3D shapes, *IEEE Pattern Anal. Mach. Intell.* **14**(2).
2. T. Pajdla and L. Van Gool, Matching of 3-D curves using semi-differential invariants, in *5th International Conference on Computer Vision*, pp. 390–395, IEEE Computer Society Press, Cambridge, MA, 1995.

3. C. Chen, Y. Hung, and J. Cheng, A fast and robust approach for registration of partially overlapping range images, in *6th International Conference on Computer Vision, Bombay, India, 1998*, pp. 242–248.
4. A. Stoddart and K. Brunnstrom, Free-form surface matching using mean field theory, in *British Machine Vision Conference* (R. B. Fisher and E. Trucco, Eds.), pp. 33–42, University of Edinburgh, Edinburgh, UK, 1996.
5. K. Higuchi, M. Hebert, and K. Ikeuchi, Building 3-D models from unregistered range images, *Graphical Models Image Process.* **57**, 1995, 315–333.
6. J. Thirion, New feature points based on geometric invariants for 3D image registration, *Internat. J. Comput. Vision* **18**, 1996, 121–137.
7. M. A. Fischler and R. C. Bolles, Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography, *Commun. Assoc. Comput. Mach.* **24**, 1981, 381–395.
8. P. Krsek, G. Lukács, and R. Martin, Algorithms for computing curvatures from range data, in *The Mathematics of Surface VIII* (R. Cripps, Ed.), pp. 1–16, Information Geometers, Winchester, UK, 1998.
9. J. J. Stoker, *Differential Geometry*, Wiley, New York, 1989.
10. P. Krsek, T. Pajdla, and V. Hlaváč, Estimation of differential parameters on triangulated surface, in *Proceedings of the Czech Pattern Recognition workshop'97*, pp. 151–155, Czech Pattern Recognition Society, Prague, Czech Republic, 1997.
11. T. Pajdla, *Laser Plane Range Finder—The Implementation at the CVL*, Technical Report K335-95-98, Czech Technical University, Prague, Oct. 1995.