Towards Shape Understanding through Non-Parametric Shape Warping

Ulrich Hillenbrand
Institute of Robotics and Mechatronics
German Aerospace Center (DLR)
82234 Wessling, Germany
Ulrich.Hillenbrand@dlr.de

***Abstract*— This paper pursues the idea of understanding shapes of unknown objects through establishing correspondence with points from the surface of known objects. A lot of geometry-related knowledge, such as functionally correct grasps or object constellations, could thus be transferred from known shapes to novel shapes of the same or a similar category. As one critical module in such a system, this paper considers warping of surfaces in 3D across significant and possibly non-smooth shape variations. Results are shown for some objects from the Princeton Shape Benchmark and a range scan.**

I. Introduction

For a robotic agent that is supposed to perform manipulation tasks in largely unconstrained environments, such as the ordinary human living environment, the fundamental problem arises of interpreting its environment in *functional* terms. That is, the agent must determine for various objects it may have never encountered before, *which role*, if any, they play within a given task, and *how to use* them to accomplish that task. This kind of perceptual problem may be denoted *affordance recognition*. For an object manipulating agent, the semantics of a scene is constituted largely by its affordances, i.e., the agent's options to purposefully act on objects in the environment.

For instance, serving a cup of coffee requires the agent to determine where the coffee pot is, as the source of the beverage, and where a cup is, as its destination and as the thing to actually serve. Moreover, to enable a proper pouring action, the agent needs to infer how to grasp the particular coffee pot by its handle and how the constellation of coffee pot and cup should be for pouring the coffee. Finally, for a proper serving of the cup, the agent has to avoid grasps that touch the interior side of the cup. Generally, the agent has to recognize all this for a set of dishes it has never seen before.

In recent years, there has been great progress in recognizing an object's category from images. That is, for images of unknown objects of a learned category, the category label is inferred [1], [2]. Progress has also been made towards finding the outline of the object [3], [4] or segmenting the object's image region [5], [6]. For functionally or semantically correct manipulation, however, a more detailed understanding of an object's surface is required. Indeed, ideally we would like to infer a detailed labeling of all 3D surface points as to their functional role within a given task. This includes a labeling of all functional parts of an object (such as a cup's handle), a labeling of potential grasp regions and non-grasp regions (such as a cup's interior surface), and, more generally, inference of any geometric constraints involving the manipulator or other objects (such as the constellation of coffee pot and cup for pouring).

In a traditional view, this problem has a perceptual and a planning aspect. On the perceptual side, the relevant object categories have to be determined, along with a segmentation of the scene into category-labeled parts. On the planning side, all the actions on the objects have to be synthesized from scratch.

The work presented in this paper is motivated by the idea of pushing the perceptual aspect a bit further to the action side, towards the goal of affordance recognition. The knowledge on object categories and their functionally correct handling within a given task context may be grounded in a few exemplars of objects and their usage. This knowledge can be retrieved by establishing detailed correspondences between shapes from those exemplar objects and similar shapes from the actual scene. In this view, a non-rigid mapping of points from exemplar shapes onto corresponding scene points transfers and adapts geometry-related knowledge to new shapes of unknown objects. Shapes in the actual scene are hence interpreted through a kind of geometric analogy to exemplar shapes.

This article pursues the general idea of shape warping for shape understanding. Some preliminary results on warping surfaces across significant shape variations in 3D are reported. Here the issue of shape category recognition is not considered, as would be needed in a complete system for shape understanding. Also, affordance recognition as proposed here will not generally replace a subsequent task, grasp, or path planning stage. Importantly, however, one can expect that it will give valuable constraints to a planner, thus significantly reducing the prohibitively large search space of an equivalent from-scratch planning problem.

II. Relation to previous work

The work most closely related to the requirements explained above and to the present work has been on non-rigid shape matching. In estimating an aligning transform, detailed correspondence is made between a source shape and target

shape. A number of methods have been used for planar point sets and curves [7], [8], [9], [3], [4], however, we are here interested in correspondences between surfaces in 3D. There has also been work on warping of 3D surfaces, mostly in the context of registering range data sets from different views to a complete model while compensating for low-frequency errors, or modeling of a deformable object such as a beating heart. Some of those methods have been based on non-rigid extensions to the iterative closest point (ICP) algorithm [10], [11], [12], and as such are suited only to local optimization of point correspondences and alignment. While in ICP-type algorithms correspondences and alignment are estimated in an alternating fashion with separate cost functions, an alternative technique optimizes both correspondences and alignment simultaneously in a joint cost function [13]. Generally, methods for nonrigid motion tracking of deformable objects achieve only incremental shape alignment and correspondences [14], [15]. Sometimes markers have to be used to guide convergence of the alignment process.

Common to most existing 3D warping techniques is, therefore, that they can handle just moderate deformations, mainly because they cast the alignment problem as one of global optimization with a just local method. Furthermore, they employ very high-dimensional parametric models to describe the deformation, such as thin-plate splines or locally variant affine transformations, which require careful regularization to sufficiently constrain the solution space. The regularization is achieved by enforcing some smoothness constraints, such as minimal bending energy, which act effectively as a shape prior. While assuming smooth deformations certainly makes sense in the application contexts of the above methods, for modeling shape variability within an object category this assumption is generally not valid.

The algorithm described in this paper differs from previous ones for 3D shape matching in two important respects. i) It realizes a global estimate of correspondences and alignment, handling large deformations between a source and a target shape. ii) It does not use a parametric model for the deformations, nor does it need a prior for expected deformations. In this sense, the estimated deformation is non-parametric and purely data-driven. In particular, even non-smooth deformations are covered by the method.

III. The shape warping algorithm

The three main steps of the shape warping algorithm are

1) deformation-tolerant rigid pose estimation,
2) correspondence estimation,
3) surface-point mapping.

It is worth noting that in this procedure, there is no joint estimation of alignment and correspondences, in contrast to most other schemes. Rather, the initial rigid alignment is estimated without determining correspondences. Once a reasonable rigid alignment is achieved, estimation of correspondences is a comparably simple step. In turn, once correspondences are determined, mapping of surface points is trivial. The most critical step in the proposed procedure is hence the initial rigid alignment.

Each of these steps will now be explained. We will denote the known object as the *source shape* and the unknown object as the *target shape*, as shape warping is supposed to map points, and hence shape-related knowledge, from the former to the latter.

A. Deformation-tolerant rigid pose estimation

The algorithm starts by globally estimating a rigid alignment, or pose, of the source object to the target object. The rigid alignment has to tolerate significant deviations of the target shape from the source shape. A reasonable alignment surely is one where corresponding parts of the two objects come as close as possible to each other. This not only reflects our intuition about the 'right' alignment of two different shapes, it also is crucial for the following step of correspondence estimation. However, we seek such an alignment without a hint as to the correct correspondences.

A correspondence-free alignment that is also robust to large geometric deviations is provided within the framework of parameter-density estimation and maximization, or parameter clustering. This is a robust estimation technique based on location statistics in a parameter space where parameter samples are computed from data samples [16], [17].[1] In the present variant, the sampling is from a surface description based on points with their local surface normal vector, which we shall refer to as *surflets*. In particular, no special geometric features or high-level primitives need to be extracted from the data, making the procedure applicable to dense range data of all kinds of shapes.

The surface normals can be estimated from range data points by local covariance statistics: the eigenvector of the covariance matrix, computed from points within a local surface region, corresponding to the smallest eigenvalue is a useful estimate of surface normal direction. The outward direction of the normal is known from the gaze direction of the sensor.

Let $\mathcal{S} \subset \mathbb{R}^3$ be the given set of points on the source shape and $\mathcal{T} \subset \mathbb{R}^3$ the given set of points on the target shape. Let further be $u(s)$ and $v(t)$ the (normalized) surface normal vectors for points $s \in \mathcal{S}$ and $t \in \mathcal{T}$, respectively.

A pose hypothesis can now be computed from a minimum subset of two source surflets matched against a minimum subset of two target surflets. The sampling proceeds thus as follows.

1) Randomly draw a point pair $s_1, s_2 \in \mathcal{S}$.
2) Randomly draw a point pair $t_1, t_2 \in \mathcal{T}$, such that the surflet pair $\{(t_1, v(t_1)), (t_2, v(t_2))\}$ is geometrically similar to $\{(s_1, u(s_1)), (s_2, u(s_2))\}$.
3) Compute a rigid motion that aligns $\{(s_1, u(s_1)), (s_2, u(s_2))\}$ to $\{(t_1, v(t_1)), (t_2, v(t_2))\}$, up to distortions.

[1] The estimator may be viewed as a continuous version of a generalized, randomized Hough transform. It is fundamentally different from RANSAC-style techniques in that no cost function is evaluated in data space.

4) Compute and store the six parameters of the hypothetical motion.

In step 2 of the sampling procedure, a metric for surflet pairs and an efficient procedure for enforcing similarity when drawing from $\mathcal{T}$ is needed. The intrinsic geometry of a surflet pair can be smoothly described by four parameters, in various ways. The Euclidean metric in such a parameter space will provide a similarity measure for surflet pairs. Similarity of the surflet pair sampled from $\mathcal{T}$ in step 2 to the one sampled from $\mathcal{S}$ in step 1 is efficiently enforced by indexing into a hash table of surflet pairs previously sampled from $\mathcal{T}$. The table is accessed through the four parameters of the drawn $\mathcal{S}$-pair as the key.

Rigidly aligning a surflet pair with another in step 3 requires trading off between positional and directional information. Unlike for pure point sets, there is no unique principled formulation of a cost function. Here we estimated the rotation from the surface normals alone, while the translation has to be estimated from the surface points. More precisely, the rotation between the two surflet pairs $\{(s_1, u(s_1)), (s_2, u(s_2))\}$ and $\{(t_1, v(t_1)), (t_2, v(t_2))\}$ was computed to minimize the squared angles between the normals, i.e.,

$$R = \arg \min_{R' \in SO(3)} \sum_{i=1}^{2} \arccos^2(v(t_i) \cdot R'\, u(s_i)) \ , \qquad (1)$$

and the translation is then the least-squares solution on the points, i.e.,

$$T = \arg \min_{T' \in \mathbb{R}^3} \sum_{i=1}^{2} \|R\, s_i + T' - t_i\|^2 \ . \qquad (2)$$

The parameterization of rigid motions chosen for sampling step 4 may have an influence on the result. In fact, the parameter density from which we sample depends upon this choice. A parameterization that is consistent for clustering is generally recommended, in the sense of [16].

By repeatedly executing the sampling procedure 1 through 4 above, we obtain samples from the parameter density for the rigid alignment problem. This parameter density is similar in spirit to a posterior density, but without assuming a probabilistic observation model; see [16] for details. In fact, for deformation-tolerant alignment, it is the great advantage of a non-parametric technique such as clustering that a probabilistic model of the expected shape variations is not needed.

The parameter samples can be stored in an array or a tree of bins. The sampling stops, when a significant cluster of samples has formed, as judged from the bin counts. Then the location of maximum parameter density is searched by a mean-shift procedure [18], [19]. This location in the 6D parameter space is returned as the pose estimate $(\hat{R}, \hat{T}) \in SO(3) \times \mathbb{R}^3$ of the target shape relative to the source shape. Details of the implementation will be presented elsewhere [20].

B. Correspondence estimation

Correspondence is a relation between points of the source shape and points of the target shape. Usually correspondence is defined as a symmetric relation. Here, however, correspondences are *directed* and symmetry shall be enforced only at the mapping stage; see next section. In particular, not all directed correspondences established will affect the final mapping of surface points. In a sense, the true correspondences will be those realized by the final mapping.

Since rigid alignment of source and target shapes has brought corresponding parts already close to each other, we can exploit again the simple local surface description through surflets to find directed correspondences: correspondences are established locally based on proximity of points and consistence of surface normals.

Let $\hat{u}(s) = \hat{R}\, u(s)$ be the surface normal vector at the source point $s \in \mathcal{S}$ rotated to align with the target surface, where $\hat{R}$ is the rotation estimate from the rigid alignment; cf. previous section. Likewise, let $\hat{v}(t) = \hat{R}^{-1}\, v(t)$ be the surface normal vector at the target point $t \in \mathcal{T}$ rotated to align with the source surface. For a rotated source-surface normal vector $\hat{u}(s)$, let $\mathcal{T}_\delta[\hat{u}(s)] \subseteq \mathcal{T}$ be the set of target points with a normal vector $v(t)$ oriented at most an angle δ away from $\hat{u}(s)$, i.e., $\hat{u}(s) \cdot v(t) > \cos\delta$. We define the *forward correspondence* of point $s \in \mathcal{S}$ as

$$C_\mathrm{f}(s) = \arg \min_{t \in \mathcal{T}_\delta[\hat{u}(s)]} \|\hat{R}\, s + \hat{T} - t\| \ . \qquad (3)$$

Likewise, we define the *backward correspondence* of a point $t \in \mathcal{T}$ as

$$C_\mathrm{b}(t) = \arg \min_{s \in \mathcal{S}_\delta[\hat{v}(t)]} \|\hat{R}\, s + \hat{T} - t\| \ . \qquad (4)$$

The shape warps shown in this paper have used a tolerance parameter $\delta = 35$ degrees for surface normal misalignment.

The constraints on correspondence derived from point distance and inter-normal angle are here enforced in a sequential manner. This avoids trading off Euclidean point differences against angular normal differences between surflets, which is in principle a problem. Nonetheless, instead of using the angle-constrained sets $\mathcal{T}_\delta[\hat{u}(s)]$ and $\mathcal{S}_\delta[\hat{v}(t)]$ and the Euclidean distance in (3) and (4), one could employ a surflet metric that has an Euclidean contribution for the points and an angular contribution for the normals.

C. Surface-point mapping

Mapping of surface points from the known, and hence completely labeled, source shape to the novel target shape endows the latter with a coordinate system in which to measure the relations to functionally relevant parts of the object. Along with that, the functionally relevant parts themselves are also identified. This should be useful for all kinds of grasp and task planning.

The mapping process interpolates the previously found directed correspondences and enforces their symmetry. Interpolation was done here in the most primitive way. For every source point queried, the n closest surface points on the source shape are found, in the ordinary Euclidean metric. The forward-mapped point on the target shape then is the average of the n forward-corresponding target points. Formally, a point

x on the source shape (but not necessarily in $\mathcal{S}$) gets *mapped forward* according to

$$M_{\mathrm{f}}(x) = \frac{1}{n} \sum_{s \in \mathcal{S}_n(x)} C_{\mathrm{f}}(s) \ , \tag{5}$$

where $\mathcal{S}_n(x) \subseteq \mathcal{S}$ is the set of n nearest neighbors of x within $\mathcal{S}$. Likewise, a point y on the target shape (but not necessarily in $\mathcal{T}$) gets *mapped backward* according to

$$M_{\mathrm{b}}(y) = \frac{1}{n} \sum_{t \in \mathcal{T}_n(y)} C_{\mathrm{b}}(t) \ . \tag{6}$$

The shape warps shown in this paper have used a neighborhood of $n = 3$ for interpolation. Equations (5) and (6) represent about the simplest possible interpolation of a mapping. Alternatively, more refined methods of interpolation between correspondences can be tried, such as distance-weighted averaging or higher-order splines.

Symmetry of interpolated correspondences is enforced through a forward/backward consistency check of the mappings. A point x on the source shape is mapped to a point $y = M(x) := M_{\mathrm{f}}(x)$ on the target shape, if and only if its forward and backward mappings agree within a radius of ϵ, that is,

$$\|x - M_{\mathrm{b}}(M_{\mathrm{f}}(x))\| < \epsilon \ . \tag{7}$$

Note that the forward/backward consistency check is similar in spirit to the left/right consistency check in stereo image processing. The shape warps shown in this paper have used a tolerance parameter $\epsilon = 0.04$ object-bounding-box diagonals for forward/backward mapping discrepancy. The result of mapping is the final result of warping.

IV. Results

In this section, the behavior of the shape warping algorithm is illustrated by some processing examples. The test data used have been derived from the Princeton Shape Benchmark (PSB) [21], which provides a public database of triangular-meshed 3D models for research on shape analysis. From this database, all mugs and a liqueur glass have been chosen as example objects for demonstration of the algorithm. Moreover, a scan of my own mug with our laser stripe profiler [22], [23] is included as another test case to also demonstrate the applicability to realistically noisy and partial representations.

For visualization of the shape warps, a labeling of a regular pattern of surface points from a source mug by each of their three Cartesian coordinates is color-coded and mapped through (5), conditioned on (7), to each of the target objects; see figures 1 through 8.

Evidently, the warps are fairly smooth and generally connect corresponding points on the object surfaces. In particular, functional parts like the handle and the mouth of mugs are correctly identified through correspondence with the source (and hence known) mug. If we teach a robot to take the known mug by the handle, the corresponding grasp region can be marked on the novel mug to provide a strong semantic constraint for a grasp planner. If the handle shapes are not too different, even individual contact points may be mapped to the novel handle. Moreover, if we teach a robot to pour a drink through the center of the known mug's mouth, it may be able to do the same with the novel mug – and even with the liqueur glass.

Some observations deserve special notice.

- In the example shown in figure 1, the handle of the target mug is degenerated to a tiny flat extension near the rim of the mouth. As a result, only the upper part of the source handle is mapped to the degenerate target handle.
- In the example shown in figure 3, when warping from the source mug to the target mug, a change in the mugs' surface topology occurs: the handle of the target mug consists of two separate handle parts. The shape warp respects this topological change by mapping the two sides of the source handle to the closer of the two parts of the target handle. Only very few points get interpolated in between the two parts.
- In the example shown in figure 4, the target handle has a rectangular cross section with extended lateral faces, in contrast to the flat source handle with only inner and outer faces. Essentially, the shape warp does not map points from the source handle to the lateral faces of the target handle, which makes sense for a transfer of grasp points: the lateral faces would afford a type of grasp quite different from those on the inner and outer faces. However, some spurious points around the target handle do result from the warp.
- When warping from the source mug to the liqueur glass shown in figure 7, the corresponding parts of the objects are correctly identified. In particular, it can be inferred from the warp that the liqueur glass has no handle, while the mug has no stem.
- When warping from the source mug to the partially scanned mug shown in figure 8, the fraction of the source mug corresponding to the scan data is correctly identified.

Computation times for the presented examples ranged from one to ten minutes, depending on the number of data points to be processed. Almost all this time was spent for computing the forward and backward correspondences (3) and (4) for all several thousand data points of each object, implemented currently under Mathematica. Initial rigid alignment took usually around a second, implemented in C++. Point mapping with consistency checking is on the order of a second, currently running under Mathematica. A pure and optimized C++ implementation can be expected to greatly reduce the current bottleneck of correspondence estimation. Moreover, in a real application, it will usually not be necessary to warp all points from a source to a target shape. Rather, an exemplar shape will be queried for some critical points or a critical region on an unknown object.

V. Discussion

The non-parametric nature of the proposed algorithm for shape warping makes it particularly suited to mapping from a single or just a few given examples of a shape category

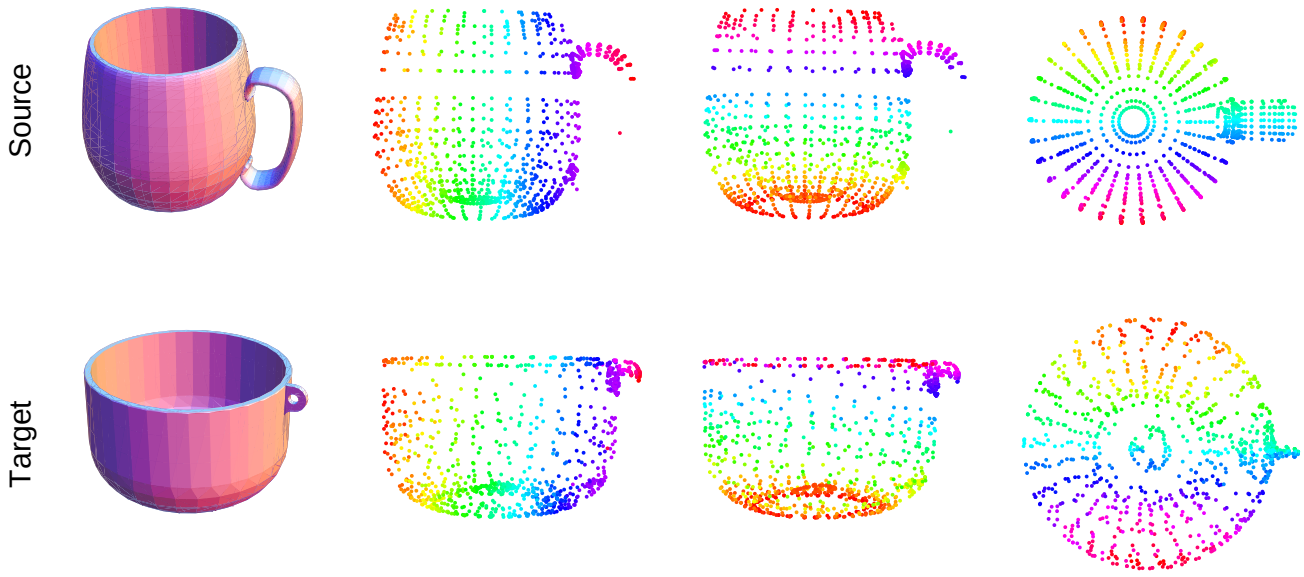

Fig. 1. A regular pattern of points on the source mug is warped to another mug from the PSB (source model number m504, target model number m503). The points are colored to code their three Cartesian coordinates in the reference frame of the source mug, one coordinate in each of the columns of the figure.

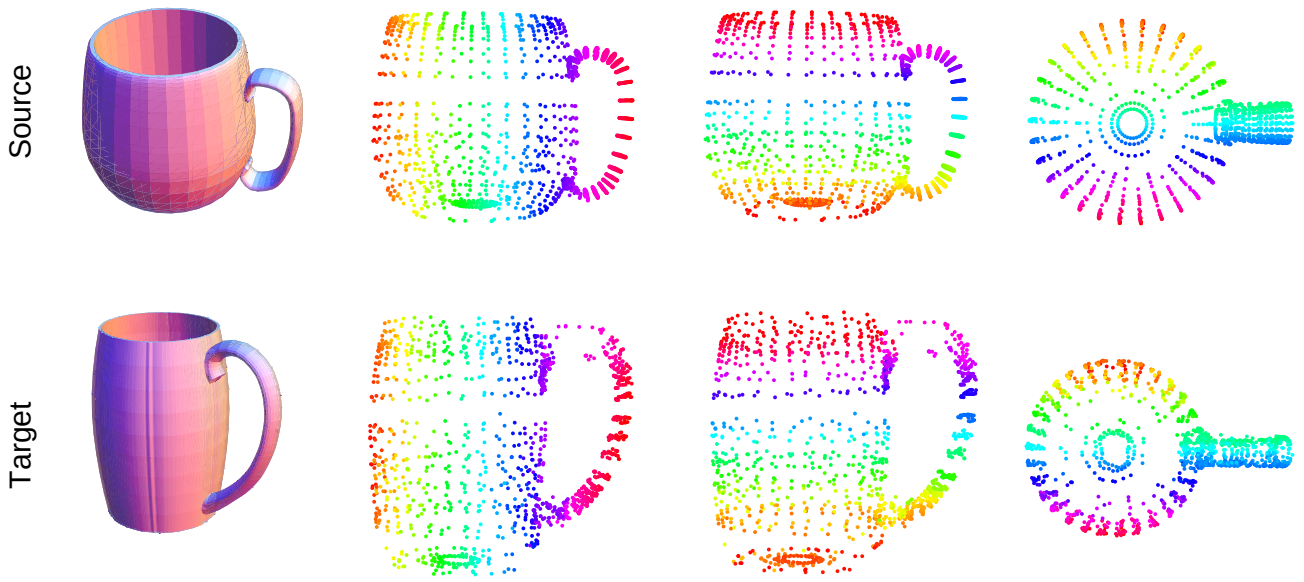

Fig. 2. Same as in figure 1, but for another target mug from the PSB (target model number m505).

to novel instances. This is crucial in cases where the number of examples is very limited. For instance, a user of a service robot may wish to teach the concept of a new object category on the fly, by just showing few examples. Moreover, the algorithm supports an unsupervised refinement of an object category model through acquiring and labeling new examples autonomously. In particular, a statistical shape model, such as an active shape model [24], may be learned from a larger number of examples with estimated point correspondences, similar to the work in [4] for the 2D case.

The processing times, even when bringing them down to a few seconds, may seem long, given that a matching against several exemplar shapes from a database may be desirable. However, the size of such a database may be moderate, if each single exemplar shape can cover a broad range of shape variation, which is precisely the goal of this work.

The results shown in this paper are somewhat preliminary in that they are limited to a relatively small data set that includes just one real case of range sensing. A larger number of realistic cases will have to be studied in the future. Moreover,

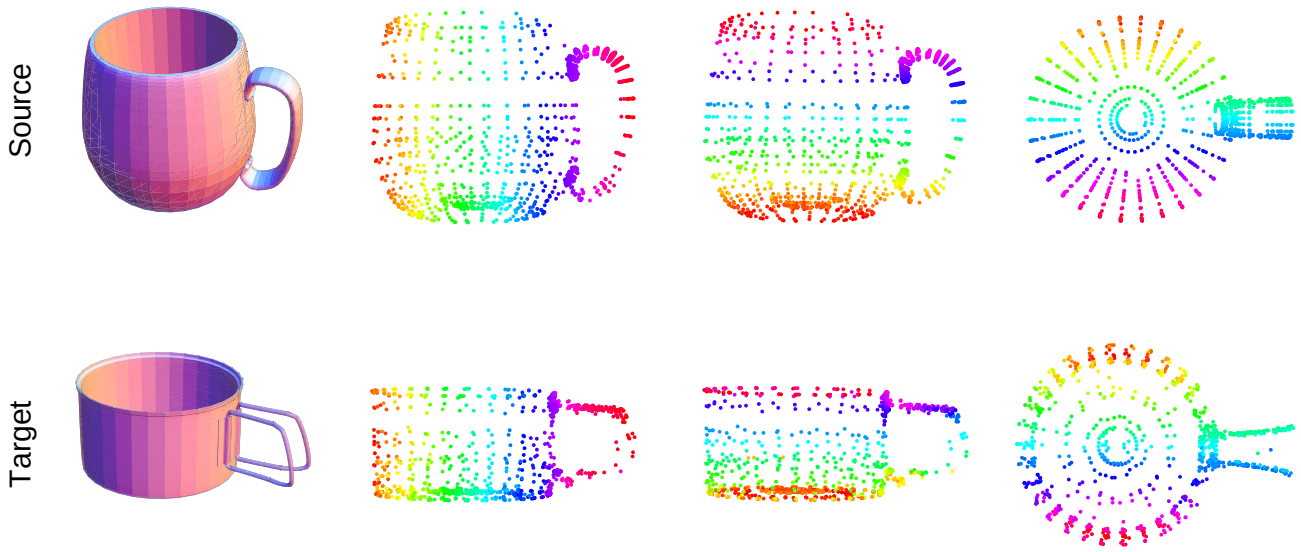

Fig. 3. Same as in figure 1, but for another target mug from the PSB (target model number m506).

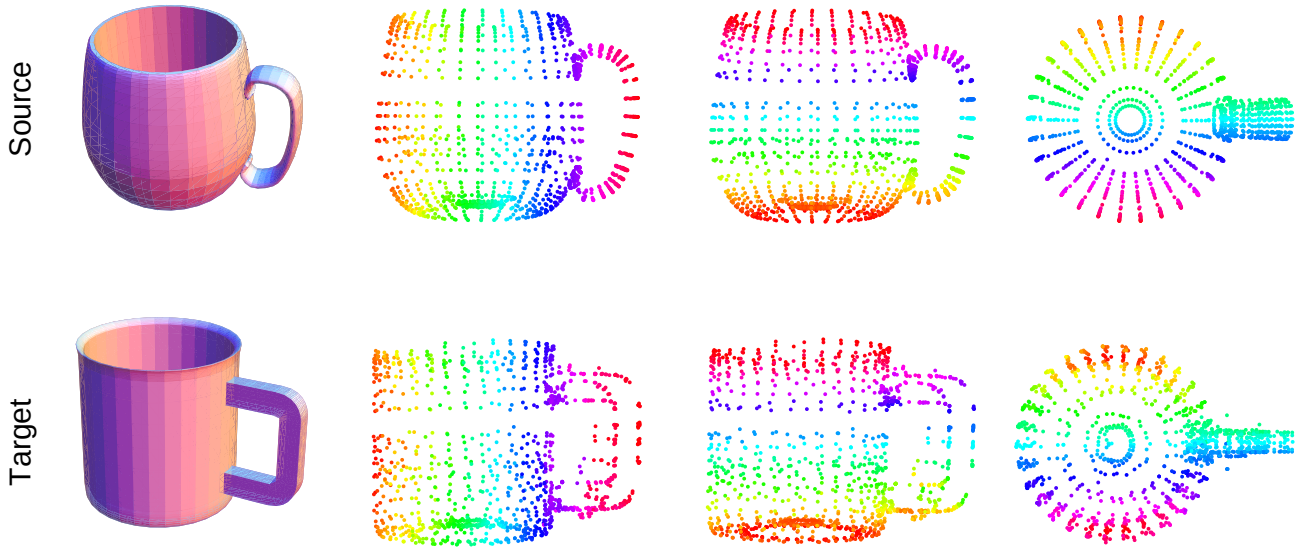

Fig. 4. Same as in figure 1, but for another target mug from the PSB (target model number m507).

an evaluation of the proposed method in the context of robotic manipulation will require mapping of real actions on known objects to novel objects. Such experiments, however, will go beyond the perceptual aspect, involving also planning and control aspects. Nonetheless, the shape warper proposed in this paper is a critical step towards the larger goal of affordance recognition by transferring geometry-related knowledge from known shapes to novel shapes.

REFERENCES

[1] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray, "Visual categorization with bags of keypoints," in *Proc. ECCV Workshop on Statistical Learning in Computer Vision*, 2004.

[2] E. Nowak, F. Jurie, and B. Triggs, "Sampling strategies for bag-of-features image classification," in *Proc. ECCV*, 2006.

[3] A. C. Berg, T. L. Berg, and J. Malik, "Shape matching and object recognition using low distortion correspondences," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2005.

[4] V. Ferrari, F. Jurie, and C. Schmid, "Accurate object detection with deformable shape models learnt from images," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2007.

[5] E. Borenstein, E. Sharon, and S. Ullman, "Combining top-down and bottom-up segmentation," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2004.

[6] D. Hoiem, C. Rother, and J. Winn, "3D layoutCRF for multi-view object class recognition and segmentation," in *Proc. Conf. Computer Vision and*

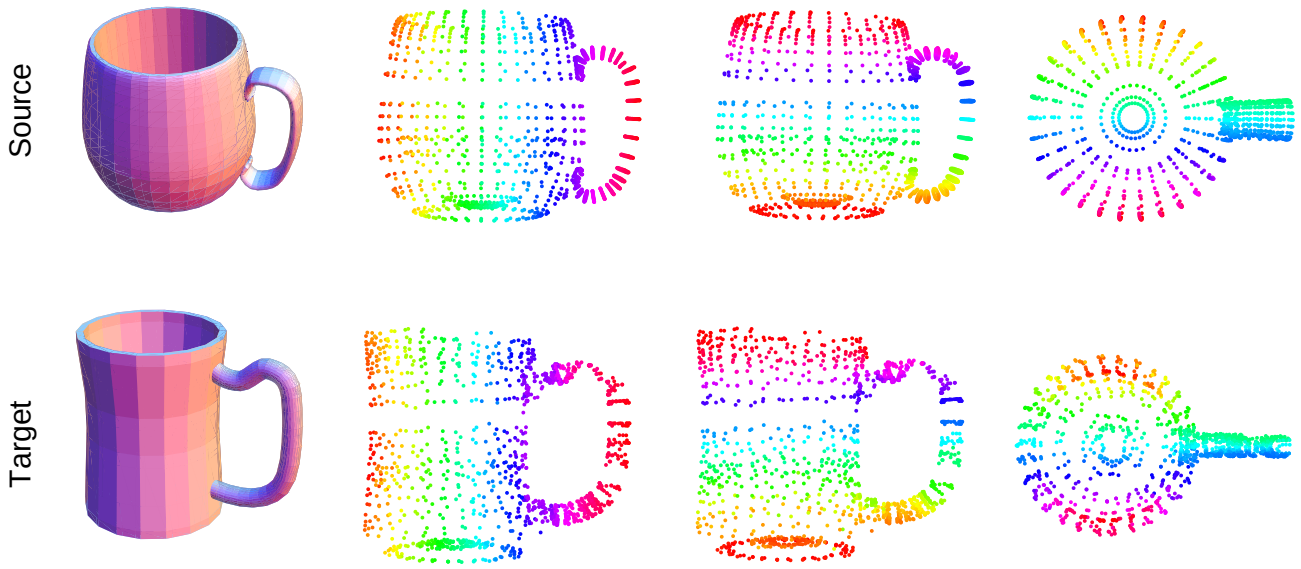

Fig. 5. Same as in figure 1, but for another target mug from the PSB (target model number m508).

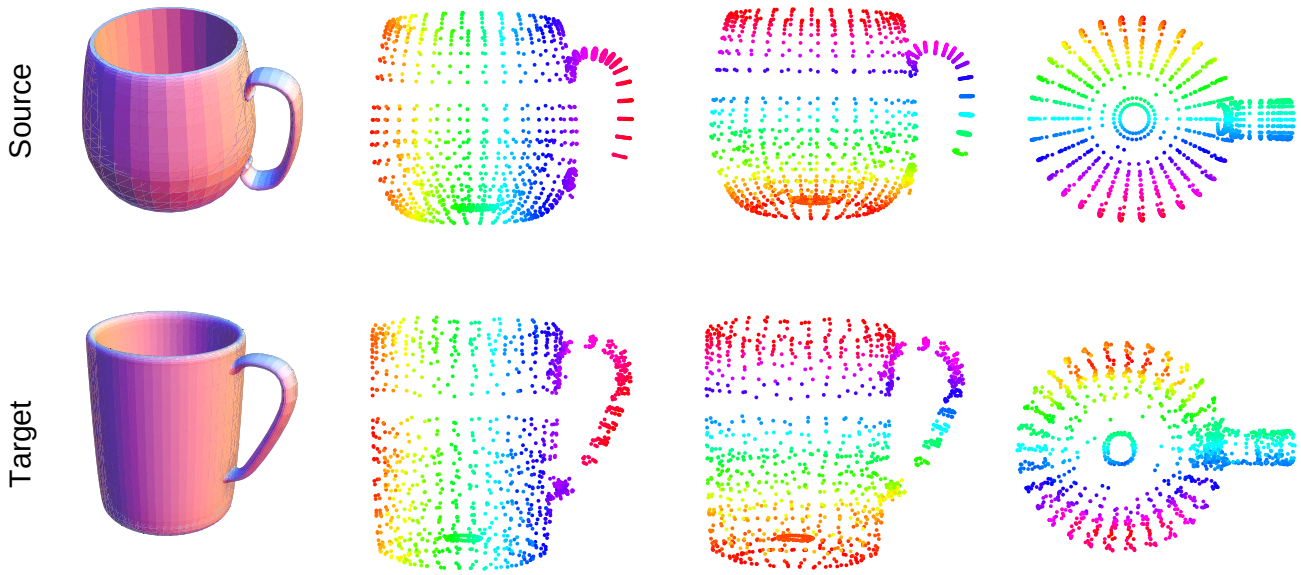

Fig. 6. Same as in figure 1, but for another target mug from the PSB (target model number m509).

Pattern Recognition, 2007.

[7] K.-W. Cheung, D.-Y. Yeung, and R. T. Chin, "On deformable models for visual pattern recognition," *Pat. Recog.*, vol. 35, pp. 1507–1526, 2002.

[8] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 24, pp. 509–522, 2002.

[9] H. Chui and A. Rangarajan, "A new point matching algorithm for non-rigid registration," *Computer Vision and Image Understanding*, vol. 89, pp. 114–141, 2003.

[10] D. Hähnel, S. Thrun, and W. Burgard, "An extension of the ICP algorithm for modeling nonrigid objects with mobile robots," in *Proc. Int. Jt. Conf. Artificial Intelligence*, 2003.

[11] B. J. Brown and S. Rusinkiewicz, "Non-rigid range-scan alignment using thin-plate splines," in *Proc. Int. Symp. 3D Data Processing, Visualization and Transmission*, 2004.

[12] B. Amberg, S. Romdhani, and T. Vetter, "Optimal step nonrigid ICP algorithms for surface registration," in *Proc. Conf. Computer Vision and Pattern Recognition*, 2007.

[13] H. Li, R. W. Sumner, and M. Pauly, "Global correspondence optimization for non-rigid registration of depth scans," in *Proc. Eurograph. Symp. Geometry Processing*, 2008.

[14] Y. Wang, M. Gupta, S. Zhang, S. Wang, X. Gu, D. Samaras, and P. Huang, "High resolution tracking of non-rigid motion of densely sampled 3D data using harmonic maps," *Int. J. Computer Vision*, vol. 76, pp. 283–300, 2008.

[15] M. A. Greminger and B. J. Nelson, "A deformable object tracking algorithm based on the boundary element method that is robust to occlusions and spurious edges," *Int. J. Computer Vision*, vol. 78, pp.

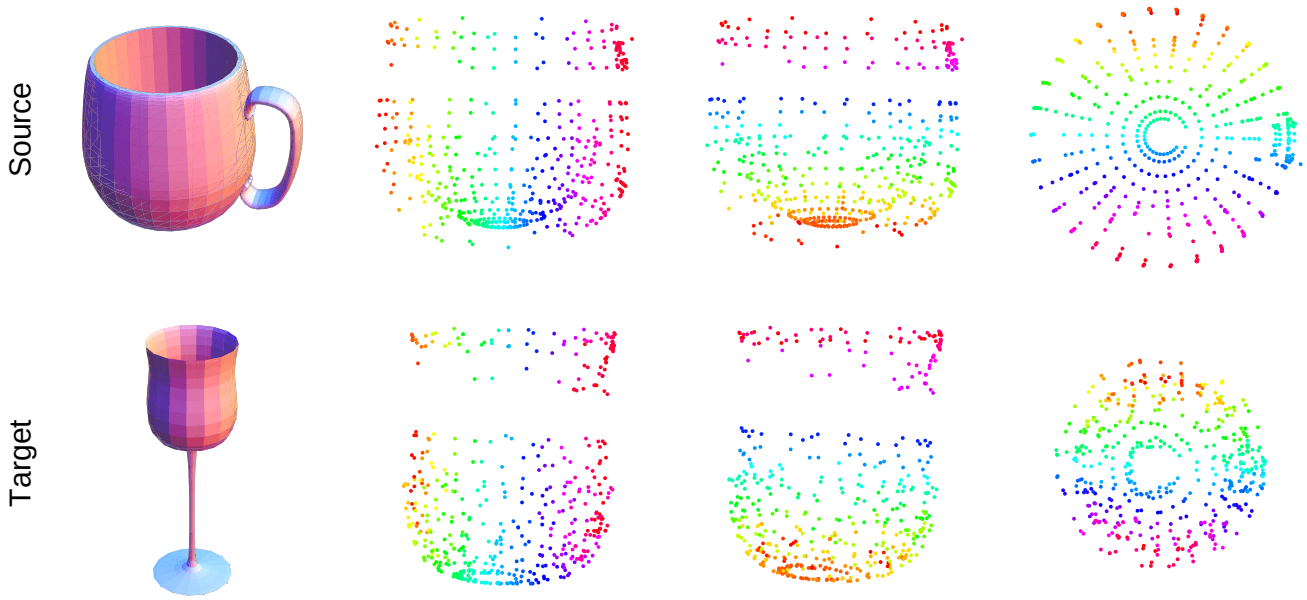

Fig. 7. Same as in figure 1, but for another target object, a liqueur glass from the PSB (target model number m494).

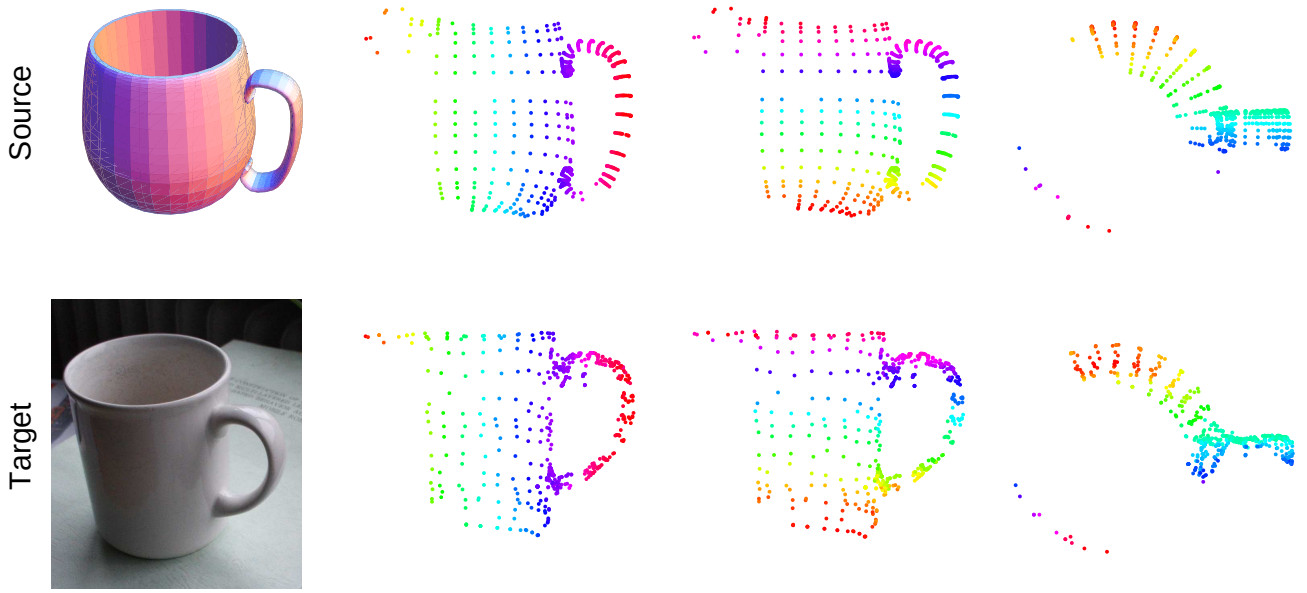

Fig. 8. Same as in figure 1, but for another target mug, acquired through partial range scan by a laser stripe profiler.

29–45, 2008.

[16] U. Hillenbrand, "Consistent parameter clustering: definition and analysis," *Patt. Recogn. Let.*, vol. 28, pp. 1112–1122, 2007.

[17] ——, "Pose clustering from stereo data," in *Proc. VISAPP Int. Workshop on Robotic Perception*, 2008, pp. 23–32.

[18] K. Fukunaga and L. D. Hostetler, "The estimtion of a gradient of a density function, with applications in pattern recognition," *IEEE Trans. Info. Theory*, vol. 21, pp. 32–40, 1975.

[19] D. Comaniciu and P. Meer, "Mean shift: a robust approach toward feature space analysis," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 24, pp. 603–619, 2002.

[20] U. Hillenbrand and A. Nölle-Fuchs, "Four variants of pose clustering from dense range data," 2009, submitted manuscript.

[21] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser, "The Princeton Shape Benchmark," in *Proc. Shape Modeling International*, 2004, http://shape.cs.princeton.edu/benchmark/.

[22] K. H. Strobl, W. Sepp, E. Wahl, T. Bodenmller, M. Suppa, J. F. Seara, and G. Hirzinger, "The DLR multisensory hand-guided device: the laser stripe profiler," in *Proc. IEEE Int. Conf. on Robotics & Automation*, 2004, pp. 1927–1932.

[23] M. Suppa, S. Kielhöfer, J. Langwald, F. Hacker, K. H. Strobl, and G. Hirzinger, "The 3D-Modeller: a multi-purpose vision platform," in *Proc. IEEE Int. Conf. on Robotics & Automation*, 2007.

[24] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models–their training and application," *Computer Vision and Image Understanding*, vol. 61, pp. 38–59, 1995.