

Transferring Functional Grasps through Contact Warping and Local Replanning

Ulrich Hillenbrand and Maximo A. Roa

Abstract—We present a method for transferring grasps between objects of the same functional category. This transfer is intended to preserve the functionality of a grasp constructed for one of the objects, thus enabling the analogous action to be performed on a novel object for which no grasp has been specified. Manipulation knowledge is hence generalized from a single example to a class of objects with a significant amount of shape variability. The transfer is achieved through warping the surface geometry of the source object onto the target object, and along with it the contact points of a grasp. The warped contacts are locally replanned, if necessary, to ensure grasp stability, and a suitable grasp pose is computed. We present extensive results of experiments with a database of four-finger grasps, designed to systematically cover variations on grasping the mugs of the Princeton Shape Benchmark.

I. INTRODUCTION

Whenever a robot has to manipulate objects that are not fully specified a priori, or the number of different objects to manipulate is large, a procedure for automatic grasp synthesis is required. The problem of grasp synthesis or planning has been studied extensively. Most methods have regarded the physics of grasping and the kinematics of grippers, planning stable and reachable grasps on objects given their 3D model [1], [2], [3]. Such an approach, however, cannot incorporate constraints imposed by specific manipulation tasks. For instance, a grasp on a mug that makes contact to its interior side may be very stable and also reachable, but is inadequate if the mug is to contain a beverage at some point in the task. An explicit formulation of all such constraints would, in general, be extremely tedious, automatic deduction by reasoning backwards from task goals infeasible in practice.

The alternative is exploitation of the knowledge that is implicit in task-specific example grasps. There have been attempts to transfer grasps constructed on database objects onto new objects with similar shape [5], [6], [7]. In principle, this transfer can preserve the function supported by a database grasp, depending on how the grasp is adapted to the new object. In [6], the approach direction and pre-shape of the database grasps are used, followed by hand closing until contacts are established. It is evident, however, that the strategy will often result in contacts very different from the original ones, depending on the shape difference between source and target objects. The planning system proposed in [8] does not transfer grasps from a database, but only the information on which part to grasp on a segmented object




Fig. 1. Two grasps (top and bottom rows), originally constructed for the mug on the left, and their warps to the mugs in the center and right. The functionality of the original grasps is preserved in the warping process. The numbers refer to the mug models in the Princeton Shape Benchmark [4].

with similar shape. Likewise, the probabilistic model trained in [9] can make inferences about task-specific hand pose relative to an object, but does not transfer a full grasp.

It may be more probable that a specific function of a grasp can be preserved, if contacts are made at locations on the target object corresponding *in detail*, i.e., down to every finger of the hand, to those on the source object. Finding corresponding points across different shapes of the same semantic category is a problem known as shape warping: a nonlinear mapping from points on the source shape to corresponding points on the target shape has to be estimated [10], [11], [12]. Most work on shape warping has been in the domain of tracking deformable or articulated objects, and hence imposes constraints such as some degree of smoothness, stiffness, rigidity, or isometry, which effectively act as a deformation prior. For modeling shape variability within a functional object category, however, such constraints are generally not valid. We therefore propose a non-parametric and data-driven warping technique that covers non-smooth and non-isometric deformations between objects. Contact points can be mapped by the warping from a source object, on which a grasp has been specified, to the target object, on

This work was funded through European Community's Seventh Framework Programme under grant agreement number ICT-248273 (GeRT).

The authors are with the Institute of Robotics and Mechatronics, German Aerospace Center (DLR), 82234 Wessling, Germany
Ulrich.Hillenbrand@dlr.de, Maximo.Roa@dlr.de

which we seek an analogous grasp with the same functionality.

Once that semantic part of grasp transfer is achieved, the physics and kinematics of the new grasp have to be dealt with. Thus, stability of the target grasp must be ensured by allowing small corrections to the warped contacts that are unlikely to change the grasp's functionality. We hence propose a local replanning procedure to find such corrections. Finally, we apply an optimization scheme to compute a grasp pose that reaches the final contacts. Figure 1 shows some examples of grasps warped between different mugs.

This study presents a performance statistics on a large and systematic sample of simulated grasps on object models. Successful experiments on grasping real objects at warped contacts are reported in [13]; there, an imitation learning approach is used to compute the final grasp from the contacts, instead of the present replanning technique.

The article is organized as follows. In section II, we describe our methods for shape warping and grasp replanning. Section III explains the dataset and evaluation procedure we used for our experiments. Results of the experiments are presented and discussed in section IV. Section V concludes this study.

II. METHODS

The proposed method for transferring a functional grasp from a source object to a target object consists of two main steps:

- A. warping the contact points of a grasp from the source surface to the target surface;
- B. checking stability of the warped contacts and, if necessary, refinement through local replanning.

We denote the complete procedure of contact warping and local replanning as *grasp warping*.

A. Warping of Object Surfaces and Grasp Contacts

Let us assume we are given two 3D shapes from the same semantic or functional category through dense sets of range data points. If the shapes are given as mesh models, a dense set of points may be uniformly drawn from the surface polygons. In our approach, the process of shape warping, that is, computing a mapping from the source shape to the target shape, has been broken down into three steps:

- 1) rigid alignment of source and target shapes, such that semantically/functionally corresponding points get close to each other;
- 2) assignment of correspondences between points from the source shape and points on the target shape;
- 3) interpolation of correspondences to a continuous (but possibly non-smooth) mapping.

The alignment step is based upon pose clustering, a robust estimation technique exploiting location statistics in a pose parameter space where parameter samples are computed from data samples [12], [14]. In our case, the data samples are surflet pairs, i.e., pairs of surface points and their local normal vectors, from source and target shapes. The pose

parameter samples are computed through aligning each pair of source surflets to a geometrically similar pair of target surflets. Clusters, i.e., density maxima, are sought in these parameter samples by a mean shift procedure, and the 50% strongest clusters are examined in data space: the source points are aligned using the transform represented by the cluster center, then the Hausdorff distance between the aligned source points and the target points is computed. The alignment with the lowest Hausdorff distance is selected as the final estimate. Since the alignment is sought for two different shapes, the procedure has to tolerate deformation in all its stages.

Once alignment of source and target shapes has brought corresponding parts close to each other, we can again rely on the local surface description by surflets to find correspondences, based on proximity of points and alignment of normal vectors. The correspondence assignment we have used here is an improvement of the method described in [12]. There, correspondences were assigned for each source surflet independently into the set of target surflets. For strong shape variation or unfavorable alignment between source and target, such *point-wise* assignment could result in corresponding parts of the shapes to not be mapped out, or even confusion of similar parts.

In order to cope with larger shape variation, some interaction between assignments of neighboring points has to be introduced. We have now formulated correspondence search as an optimal assignment problem. In this formulation, interaction between assignments of different points is enforced through uniqueness constraints.

Let $\{x_1, \dots, x_N\}$ be points from the source shape, transformed to align with the target shape; let $\{y_1, \dots, y_N\}$ be points from the target shape.¹ Assignment of source point i to target point j is expressed as an assignment matrix,

$$a_{ij} = \begin{cases} 1 & \text{if } i \text{ is assigned to } j, \\ 0 & \text{else.} \end{cases} \quad (1)$$

Let further be

$$d_{ij} = \|x_i - y_j\| \quad (2)$$

the Euclidean point distances between source and target points, and

$$c_{ij} = n_i \cdot m_j \quad (3)$$

the angle cosines between the unit normal vectors n_i and m_j at source point i and target point j , respectively. See Fig. 2 for an illustration of these relations.

The objective is to minimize the sum of distances between correspondences, i.e., mutually assigned points,

$$D(a_{11}, \dots, a_{1N}, a_{21}, \dots, a_{NN}) = \sum_{i=1}^N \sum_{j=1}^N d_{ij} a_{ij}, \quad (4)$$

¹Given sufficiently dense source and target point sets, an equal number N of points can be randomly re-sampled from each, such that source/target correspondences exist between the samples.




Fig. 2. Geometric relations between source and target surflets.

subject to the constraints

$$\sum_{i=1}^N a_{ij} = 1 \quad \forall j \in \{1, \dots, N\}, \quad (5)$$

i.e., to assign every target point to exactly one source point,

$$\sum_{j=1}^N a_{ij} = 1 \quad \forall i \in \{1, \dots, N\}, \quad (6)$$

i.e., to assign every source point to exactly one target point, and

$$c_{ij}a_{ij} \geq 0 \quad \forall i, j \in \{1, \dots, N\}, \quad (7)$$

i.e., to assign only between points with inter-normal angle ≤ 90 degrees. The two equality constraints mediate the desired interaction between assignments of different points. The inequality constraint can exclude points from being assigned. Therefore, to ensure feasibility of the problem (one-to-one assignment), we have to add to the sets imaginary source and target points x_0 and y_0 , respectively, which have no position and no normal direction. They can be accommodated by appending some large entries d_{0j} and d_{i0} to the distance matrix,

$$d_{0j} = d_{i0} > \max_{k,l \in \{1, \dots, n\}} d_{kl}, \quad (8)$$

as well as zero entries to the angle cosine matrix,

$$c_{0j} = c_{i0} = 0. \quad (9)$$

These imaginary points can be assigned to all real points with a penalty, which is sized such that only points without a compatible partner will receive this imaginary assignment; see Fig. 3 for an illustration. We then minimize the cost function

$$\begin{aligned} & C(a_{01}, \dots, a_{0N}, a_{10}, \dots, a_{NN}) \\ &= D(a_{11}, \dots, a_{1N}, a_{21}, \dots, a_{NN}) \\ &+ \sum_{i=1}^N d_{i0}a_{i0} + \sum_{j=1}^N d_{0j}a_{0j}. \end{aligned} \quad (10)$$

As a technical remark, we note that although the assignment variables a_{ij} may take fractional values, there is always an optimal solution with integer values $\{0, 1\}$. This is guaranteed by the total unimodularity of the constraints [15], which in turn may be checked by replacing the matrix c_{ij} by $\text{sign}c_{ij}$, leaving the constraints effectively unchanged.

For solving this constrained optimization problem, we have used the interior-point algorithm, which is guaranteed to find an optimal solution in polynomial time [16]. Nonetheless, a simultaneous treatment of a dense set of points (say,




Fig. 3. Assignments between real source/target points with compatible normal directions, and assignment of real source/target points with incompatible normal directions to imaginary target/source points without normal direction.



Fig. 4. Mug warping example. A dense set of surface points from the source mug (top row) and their mappings to the target mug (bottom row) are colored to code their three Cartesian source coordinates (three columns).

10,000 points) from an object surface is unpractical in terms of computation times. Therefore, for warping the whole source shape to the target shape, we had to resort to some approximation.²

Instead of solving one large assignment problem for all points, we have solved many smaller assignment problems in overlapping regions of the shapes. These *region-wise* assignments then have to be stitched together to yield a coherent set of point correspondences.

Finally, point correspondences may be interpolated to obtain a continuous (but possibly non-smooth) mapping of points from the source domain to the target domain. Options are linear interpolation between a number of source points nearest to a query point, or simply mapping to the average of the targets corresponding to the nearest source points [12]. However, no interpolation has been used in this study on grasp warping. Rather, contact points of a source grasp have simply been “projected” onto their nearest source point that had a target point assigned. These target points have then been taken as the warped contacts.

Figure 4 shows an example of a dense set of surface points warped between two mugs.

²In fact, for warping contact points of a grasp, it may be sufficient to warp the local region of each contact.

B. Stability Checking and Local Replanning of Grasp Contacts

Let points p_i on the object be given relative to the object's center of mass (CM). The force f_i applied on the object at point p_i then generates a torque $\tau_i = p_i \times f_i$ with respect to CM. f_i and τ_i are collected in a wrench vector,

$$\omega_i = \begin{pmatrix} f_i \\ \tau_i/\rho \end{pmatrix}, \quad (11)$$

where ρ is a parameter used to adjust the metric of the wrench space, which depends in general on the object size. To obtain a nicely scaled wrench space, we set ρ to the largest distance from the object's CM to any point of the object. Thus, the maximum potential torque is equal to the maximum applied force (typically unitary). A thorough discussion of other options for this parameter can be found in [17].

We assume Coulomb's friction model for point contacts made between each finger and the object. Hence, to avoid slipping, the force f_i applied at point p_i must lie inside the friction cone $f_i^t \leq \mu f_i^n$, where μ is the friction coefficient and f_i^t and f_i^n are the tangential and normal components of f_i , respectively. In the 3D space this model is nonlinear and, to simplify it, the friction cone is linearized using an m -sided polyhedral convex cone.

The wrench ω_{ij} generated by a unitary force f_i acting at p_i along an edge j of the linearized friction cone is called a primitive wrench. A grasp defined by the set of contact points $C = \{p_1, \dots, p_n\}$ is associated with the set $W = \{\omega_{11}, \dots, \omega_{1m}, \dots, \omega_{n1}, \dots, \omega_{nm}\}$ of primitive contact wrenches.

Once contact points on the target object's surface are obtained through the warping process, we complete the construction of the target grasp in three steps:

- 1) checking of the force closure (FC) condition;
- 2) if not FC, local replanning of grasp contacts;
- 3) computing a hand pose that reaches the grasp contacts.

As a first step, we need to evaluate if the new grasp is FC. The desired FC property means that the grasp is able to resist external disturbances in any direction. A necessary and sufficient condition for a grasp to be FC is that the origin O of the wrench space lies strictly inside the convex hull $CH(W)$ of the primitive contact wrenches [18]. In this work, the condition $O \in CH(W)$ is checked by verifying that O and the centroid of the primitive contact wrenches W (which is always an interior point of $CH(W)$) lie on the same side of the hyperplane H_k containing the facet k of $CH(W)$, for all k [19].

In case a set of warped contact points with its corresponding normal directions does not fulfill the FC condition, a local replanning of the grasp must be invoked. This replanning is performed in the local neighborhood of the warped contact points: a one-at-a-time substitution is performed, trying to move one of the fingers to get an FC grasp while the rest of the fingers are located in the position given by the warping process [17]. As a result of this process, a new FC grasp can be obtained, see Fig. 5 for an example; or the replanning can




Fig. 5. Example of local replanning. Red dots represent the initial warped contact points, and green dots represent the final contact points after local replanning.

fail and it is concluded that the new contact points are not good locations for obtaining an FC grasp.

Finally, we find a feasible hand pose that realizes the grasp. The problem is formulated as an optimization problem subject to constraints derived from the hand kinematics and the location and corresponding normal direction of the contact points [20]. If no solution is found, the grasp contacts are not reachable.

III. EXPERIMENTS

Experiments were conducted with three variants of warping: the point-wise method introduced previously [12], and the new region-wise method introduced here with two different parameterizations: region sizes of 200 and 400 points, to be referred to by 'region-200' and 'region-400', respectively. The number of regions for covering the objects was 400 for region-200 and 200 for region-400, hence having an equal number of $200 \times 400 = 80,000$ points considered in assignment problems for all region-wise warps.

In these experiments, we used mugs with handle as source and target objects, because they have some interesting properties for grasping: they are non-symmetric, non-convex shapes, have surface that is not simply connected, and have reachable interior and exterior parts. They hence afford a variety of distinct grasp types.

A. Grasp Database

We have constructed a database of four-finger precision grasps for the DLR hand II [21] on a mug model from the Princeton Shape Benchmark [4]. The contact points of these grasps were then warped and, if necessary, refined to all the other mugs of the Princeton Shape Benchmark, with the exception of one mug that does not have a graspable handle and thus had to be left out. See Fig. 6 for a view of the mug models involved. Mugs will be referred to by their model number in the Princeton Shape Benchmark, also given in Fig. 6.

For the construction of a database of FC grasps for mugs, five different types of grasp were considered, namely

- grasps with all the fingers located on the handle of the mug, referred to by AH;
- side grasps with one of the fingers touching the handle of the mug, referred to by SH;



Fig. 6. Mugs from the Princeton Shape Benchmark [4] used in our experiments, with their model numbers. Mug 504 was used as the source shape on which the grasp database was constructed, the others as the target mugs onto which the grasps were warped.

- top grasps with one of the fingers touching the handle of the mug, referred to by TH;
- top grasps without touching the mug handle, referred to by TNH;
- top grasps with one finger placed in the interior of the mug, referred to by TIP.

Figure 7 shows one grasp of each type on the source mug. Ten grasps of each type were synthesized on that mug. We have thus constructed a database of 50 four-finger grasps. Warping these from the source mug to the five target mugs yields a total of 250 grasps for our statistics.

B. Evaluation of Grasps

The grasps constructed for the database as well as the grasps warped onto the target objects are analyzed for their stability, quality, reachability, and functionality.

The grasp quality is quantified with the largest perturbation wrench that the grasp can resist independently of the perturbation direction [22]. This grasp quality is equivalent to the radius of the largest hypersphere centered on O and fully contained in $CH(W)$, i.e., the distance from O to the closest facet of $CH(W)$. We used a friction coefficient of $\mu = 0.4$, following the results from [23].

In the absence of a specific robotic task for this study, a functional warped grasp is here defined as one that has its contacts placed at parts on the target mug corresponding to their original locations on the source mug (on handle, on side, inside, outside). The rationale is that such a warped grasp will preserve the function of the original grasp for most tasks. Functionality of the warped grasps has been decided by visual inspection on the target.

We note that functionality is of special relevance for our methodology, as one of the main motivations for grasp warping has been to transfer functional grasps to new objects.

IV. RESULTS

We here present two types of analyses of our results: grasp quality statistics and success ratios across all warped grasps, and frequencies of the different failure cases for each grasp type and target object.

A. Grasp Quality Statistics

Figure 8 shows, for the three warping variants investigated, scatter plots of the quality of warped grasps versus their original quality on the source mug. Colors distinguish cases of grasps obtained through pure contact warping and grasps that have required local replanning for stability. Because local replanning stopped as soon as an FC grasp was found, qualities of replanned grasps are generally lower than those of stable contact-warped grasps. Qualities of warped grasps are mostly lower, but sometimes higher than those of their original version.

The mean quality ratio of warped to original grasps is shown in Fig. 9 top, for the three warping variants. Error bars indicate that there is a significant improvement of the region-wise over the point-wise method, while the two region-wise variants do not differ significantly. The region-wise method achieves a mean grasp quality that is around half that of the original grasp.

It is of interest, how many grasps were stable after pure contact warping, i.e., without local replanning. These ratios are shown in Fig. 9 second from top. Region-wise warping is again superior to point-wise warping, and roughly 70% of their warped grasps are directly FC without replanning.

Another interesting quantity is the ratio of warped grasps, with or without local replanning, that are reachable with the same robot hand as used for the original grasps (the DLR hand II). This is seen in Fig. 9 third from top. Region-wise warping is again superior to point-wise warping, and around 80% of their warped grasps are reachable. We note that local replanning, in the present version, does not take reachability of a grasp into account, hence reachability has to be preserved during contact warping.

Finally, the functionality of the warped grasps in the sense defined in section III-B has to be assessed. The ratios of the warped grasps, with or without local replanning, that are functionally equivalent to their original grasp are shown in Fig. 9 bottom. For all three warping methods, around 85% of their warped grasps preserved their functional character. We note that local replanning does not take functionality of a grasp into account, hence functionality has to be preserved during contact warping.

B. Failure Cases

We have distinguished three categories of failure of a warped grasp. These are instability (not FC), unreachability (by the DLR hand II), and non-functionality (contacts not corresponding to original grasp). Figure 10 shows the number of cases that fall into each failure category, for all grasp types, target objects, and warping variants.

After replanning, only one instance of an unstable grasp remained, with the point-wise method. There is a marked




Fig. 7. One example of each type of grasp in our database.

problem with the grasps that have a contact on the handle when they get warped to mug 506. This does not come as a surprise, as the handle of that mug has no central part to make contact to, but rather consists of two wires with a large gap in between; cf. Fig. 6. Moreover, the much smaller height of mug 506 often caused side contacts on mug 504 to be mapped very close to the rim of mug 506. In fact, that particular mug caused a large part of all failures. It is arguable that, in principle, many grasps on the other mugs are impossible to apply to mug 506. This then marks the limit of the methodology of grasp transfer between objects.

In general, the grasp types that make handle contacts are more difficult to warp correctly. Again, this was to be expected, as the handle is that part of the mugs that varies most strongly in shape between different instances. The region-wise method has very few failures when the handle is not involved, and still treats the large majority of handle contacts properly, with the exception of target mug 506.

V. CONCLUSIONS

The results are very promising for our general strategy of grasp transfer between objects. For the new region-wise warping method, we achieved roughly 80% of at the same time stable and reachable grasps that could hence be executed on the target mugs. Around 85% of warped grasps were functionally equivalent to their original grasp, although not all of them were reachable with the DLR hand II. The high percentage of functional equivalence, however, confirms the validity of our core hypothesis underlying this work: that grasps fulfilling specific task constraints can be transferred through warping and preserve their functionality.

Essentially, the method of grasp transfer is based on warping between similar shapes. Hence, grasps may be transferred between similar shapes even if they belong to different functional categories. What is then lost, of course, is the notion of functional equivalence of source and target grasps.

In this study, we used synthetic 3D data of objects for grasping, which are ideal in terms of accuracy and completeness. For real objects, geometric representations

with comparable quality can be acquired offline with laser scanners and some post-processing. Having acquired such a model for a new object, a robot may, by the present method, automatically generalize its manipulation knowledge to the new geometry [13].

On the other hand, when dealing with unknown and, hence, unmodeled objects, we will have to cope with less perfect range data: scene data will be noisy and incomplete. How the present method will work under such conditions remains to be established.

Contact warping may be improved by taking more descriptive shape features of the contact regions into account. Replanning may be performed to ensure, besides stability, also reachability of the new contacts. These improvements are likely to further increase the success rate of grasp warping, and they will be pursued in the future. Moreover, other functional object classes beside mugs will be explored to study the wider applicability of grasp warping.

REFERENCES

- [1] C. Borst, M. Fischer, and G. Hirzinger, "A fast and robust grasp planner for arbitrary 3D objects," in *ICRA*, 1999, pp. 1890–1896.
- [2] M. Przybylski, T. Asfour, and R. Dillmann, "Unions of balls for shape approximation in robot grasping," in *IROS*, 2010, pp. 1592–1599.
- [3] J. Saut and D. Sidobre, "Efficient models for grasp planning with a multi-fingered hand," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 347–357, 2012.
- [4] 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/>.
- [5] Y. Li, J. L. Fu, and N. S. Pollard, "Data-driven grasp synthesis using shape matching and task-based pruning," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, pp. 732–747, 2007.
- [6] C. Goldfeder, M. Ciocarlei, H. Dang, and P. Allen, "The Columbia grasp database," in *ICRA*, 2009, pp. 1710–1716.
- [7] P. Brook, M. Ciocarlei, and K. Hsiao, "Collaborative grasp planning with multiple object representations," in *ICRA*, 2011, pp. 2851–2858.
- [8] J. Aleotti and S. Caselli, "Part-based robot grasp planning from human demonstration," in *ICRA*, 2011.
- [9] D. Song, C. H. Ek, K. Huebner, and D. Kragic, "Embodiment-specific representation of robot grasping using graphical models and latent-space discretization," in *IROS*, 2011.
- [10] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching," *Proc. Natl. Acad. Sci.*, vol. 103, pp. 1168–1172, 2006.




Fig. 8. Warped versus original grasp quality for our database and the five target mugs. Grasps were obtained through pure contact warping (blue points), or warping with local replanning (red points).

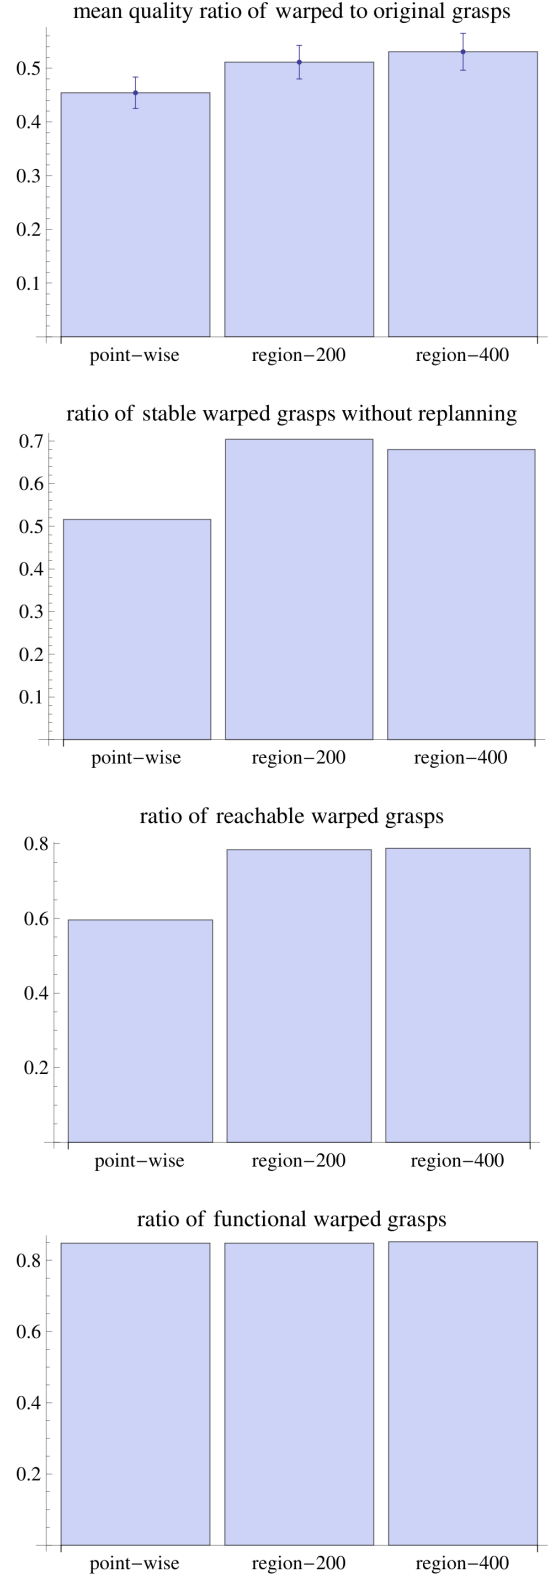


Fig. 9. For the three warping variants (from top to bottom): mean ratio of the warped to original grasp qualities (error bars are standard errors of the sample mean), ratio of warped grasps that were stable without local replanning, ratio of warped grasps that were reachable by the DLR hand II, and ratio of warped grasps that made contacts on the target mug corresponding to the original contacts on the source mug.

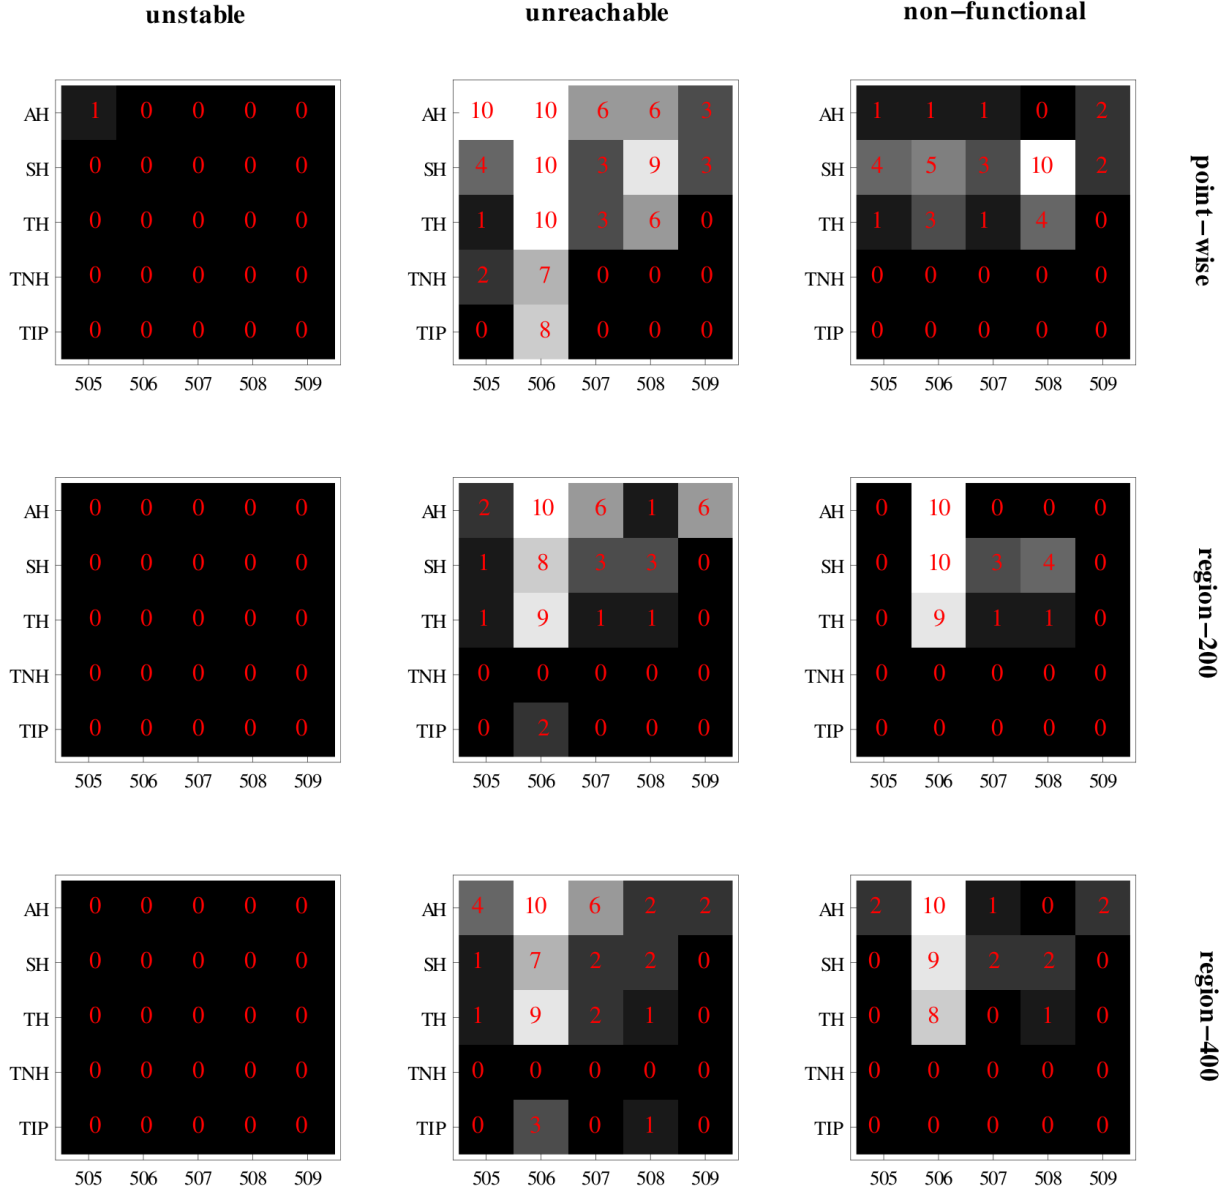


Fig. 10. Frequency of the three failure cases, for all grasp types, target objects, and warping variants. The total number of grasps of each type and for each target object was 10.

- [11] 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.
- [12] U. Hillenbrand, "Non-parametric 3D shape warping," in *Proc. Int. Conf. Pattern Recognition*, 2010.
- [13] H. Ben Amor, O. Kroemer, U. Hillenbrand, G. Neumann, and J. Peters, "Generalization of human grasping for multi-fingered robot hands," in *IROS*, 2012.
- [14] U. Hillenbrand and A. Fuchs, "An experimental study of four variants of pose clustering from dense range data," *Computer Vision and Image Understanding*, vol. 115, pp. 1427–1448, 2011.
- [15] A. Schrijver, *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- [16] R. Vanderbei, *Linear programming: foundations and extensions*. Springer, 2001.
- [17] M. A. Roa and R. Suarez, "Computation of independent contact regions for grasping 3-D objects," *IEEE Trans. Robotics*, vol. 25, pp. 839–850, 2009.
- [18] R. Murray, Z. Li, and S. Sastry, *A mathematical introduction to robotic manipulation*. CRC Press, 1994.
- [19] M. A. Roa and R. Suarez, "Finding locally optimum force-closure grasps," *J. Robotics and Computer Integrated Manufacturing*, vol. 25, pp. 536–544, 2009.
- [20] C. Borst, M. Fischer, and G. Hirzinger, "Calculating hand configurations for precision and pinch grasps," in *IROS*, 2002, pp. 1553–1559.
- [21] J. Butterfass, M. Grebenstein, H. Liu, and G. Hirzinger, "DLR hand II: next generation of a dextrous robot hand," in *ICRA*, 2001, pp. 109–114.
- [22] C. Ferrari and J. Canny, "Planning optimal grasps," in *ICRA*, 1992, pp. 2290–2295.
- [23] K. Matheus and A. M. Dollar, "Benchmarking grasping and manipulation: properties of the objects of daily living," in *IROS*, 2010.