Functional power grasps transferred through warping and replanning

Theodoros Stouraitis, Ulrich Hillenbrand, Máximo A. Roa

***Abstract*— This paper presents a method to transfer functional grasps among objects of the same category through contact warping and local replanning. The method transfers implicit knowledge that enables an action on a class of objects for which no explicit grasp or task information has been given in advance. Contact points on the source object are warped based on global and local shape similarities to the target object. These warped contacts are then used to define a hand posture that reaches close to them, while at the same time provides the desired functionality on the object. The approach is tested on different sets of objects with a success rate of 87.5%, and large benefits are shown when compared to a naive technique that only transfers a suitable hand pose to the novel object.**

I. Introduction

Data-driven methods for grasp synthesis can be classified according to the amount of previous knowledge required from the object to be grasped [1]. For known objects, whose CAD model is given in advance, there is usually an offline-computed database of grasps available for the object, and a suitable grasp is chosen online according to the estimated object pose. For unknown objects there is no previous object model, and the grasp technique looks for specific features on the object that indicate potentially good grasps. In the case of familiar objects, i.e. objects that belong to some predefined category, the novel object is compared to a stored database of known objects (of the same category) to find similar shapes and infer from them how to grasp the new one.

Grasping familiar objects is based on the idea that objects can be grouped together into categories defined by common characteristics, for instance, shape, texture or functionality. When object familiarity is not clearly defined, geometric proximity can be used as an indicator to find similar objects in a predefined grasp database [2]. In this case, preshape and approach directions are retrieved and used as an initial configuration, followed by hand closing until contacts are established. A similar knowledge database for different object types is proposed in [3], including not only shape features (aspect ratio of the bounding box of the object) but also physical features (material and weight). These general database strategies can lead to contacts very different (in terms of relative location or functionality) from the original ones, depending on the shape difference between source and target objects.

Human demonstrations can also be used to identify familiar shapes. An initial approach presented in [4] uses curvature information from the human hand performing different grasping actions to define a hand preshape for grasping objects with similar curvatures. A grasp planner that uses human demonstrations to identify suitable parts on the object to perform good grasps was proposed in [5]. However, only the information on which part to grasp is transferred to a segmented object with similar shape. Likewise, the probabilistic model trained in [6] can make inferences about task-specific hand pose relative to an object, but does not transfer a full grasp. A novel transfer technique was proposed in [7] using a representation based on electrostatic simulation to define the spatial relationship between hand and object, including the envelopment of the hand wrapping around the object. However, information on the hand configuration and finger kinematics is neglected on the process, and no functionality information is included.

Task constraints used in conjunction with the information on object geometry can be used to create semantic constraints (attributes) for grasps [8]. Information on hand kinematics, object geometry and tactile contacts is embedded into a

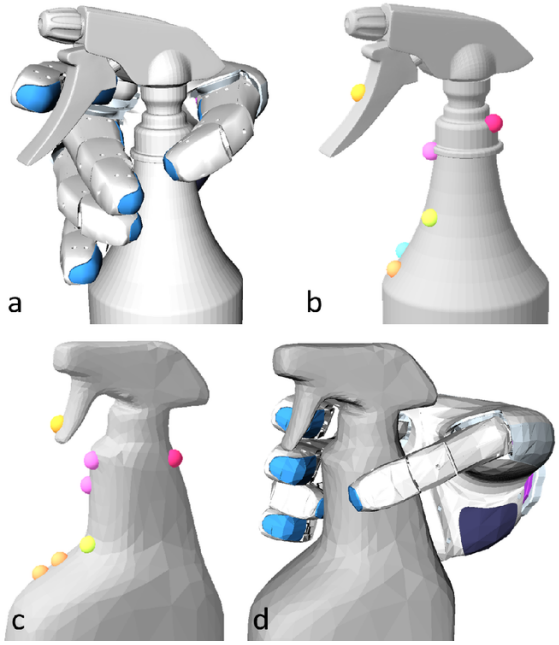

Fig. 1. The proposed method for transferring a grasp specified on a source object (a) takes the set of source contacts (b) between object and hand, warps it to the target object (c), and finds a grasp through replanning (d) with real contacts as close as possible to the warped contacts. Spheres with same color denote corresponding source/target contacts. More than one target contact may be generated by warping of a source contact, thus generating grasp variations that are explored as alternatives by the grasp replanner.

This work was supported by the project SMErobotics, from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 287787.

The authors are with the Institute of Robotics and Mechatronics in the German Aerospace Center (DLR), 82234 Wessling, Germany. Email: {firstname.lastname}@dlr.de

semantic affordance map, used to retrieve ideal approach vectors to grasp a novel object. This ideal vector is used as a seed for an Eigengrasp (synergies-based) planner [9] that provides a good grasp preshape for the hand. After this, as in all previous approaches that provide a preshape for the hand, fingers are simply closed until they produce contact with the object surface. Task and gripper constraints can also be jointly used to deduce possible affordances on novel objects, although the complexity of the grasp transfer is reduced by using only two-finger grippers [10].

An alternative approach to grasp transfer is through emphasizing the relevance of the particular contact points between the object and the robot hand. Specific functionality of a grasp is then assumed to depend on those contact points, and reproducing similar contact locations on new objects with the same parts of the hand yields analogous grasps. Finding corresponding points across different shapes of the same semantic category is a problem known as shape warping [11], [12], [13]. Contact points can be mapped by the warping from a source object, on which a grasp has been specified, to a target object, on which we seek a grasp with the same functionality.

In our previous work along this line [14], [15], a set of specific grasps was defined on a source object representative of a functional category. Its complete surface was warped onto a set of target objects of the same category, and along with it the contact points of the source grasps. Then a local replanning was performed for the warped contacts, in case that the contact set was not force closure. However, the method was used only for precision (fingertip) grasps and the test set of objects was rather limited.

This paper extends our previous work to include the case of power grasps and tests on a wider set of objects. The warping process has been improved to avoid mapping the complete geometry of the source/target objects and to exploit global and local shape similarities to warp the contact locations. Also, the local replanning method is now able to deal with the higher number of contacts usually present in power grasps, and includes reachability and collision information into the replanning loop. Fig. 1 illustrates our process of grasp transfer through images of its main stages from source grasp to transferred grasp.

The article is organized as follows. Section II presents the method used for contact warping. Section III presents the replanning approach for power grasps. Results of experiments with different sets of objects are presented in Section IV. Finally, Section V points out the conclusions of the study and discusses promising research directions.

II. Warping technique

Let us assume we are given two 3D shapes from the same semantic or functional category through dense sets of 3D data points. If the shapes are given as mesh models, a dense set of points may be uniformly sampled from the surface polygons. One shape serves as the source object, on which grasps of interest are specified. The other is the target object, onto which those grasps are to be transferred. The strategy followed in this work is to first warp the grasp contact points, based on global and local similarities of source and target shapes. The contacts warped to the target are then adapted, taking constraints of kinematics, collision and stability into account as described in Section III on grasp replanning. We now describe the procedure for warping the contacts.

For the source contact points, corresponding points on the target shape are sought. This is pursued in two steps.

A. Estimate an alignment of source and target shapes, so functionally corresponding parts get close to each other.

B. Determine points on the target shape that are likely to correspond to grasp contacts given on the source shape.

The first step takes global shape characteristics into account, attempting a pose fit of one shape to the other through minimizing deviations allover. The second step tries to reproduce the contact positions and orientations in the aligned reference frames of the objects, while also maximizing local shape similarity between source and target around each contact. These two consecutive steps will now be explained.

A. Alignment of Source and Target Shapes

The initial alignment is estimated in a way similar to the procedure used in [14], [15]. The global and robust estimator is based upon pose clustering, relying on location statistics in a 6D pose parameter space, where parameter samples are computed from data samples. The data samples used here are pairs of surflets $(\boldsymbol{p}, \boldsymbol{n})$, i.e., surface points $\boldsymbol{p} \in \mathbb{R}^3$ and their local normal vectors $\boldsymbol{n} \in \mathbb{R}^3$, $\|\boldsymbol{n}\| = 1$, from the source and target shapes. As explained thoroughly in [16], the pose parameter samples are computed through aligning pairs of source surflets to geometrically similar pairs of target surflets. Clusters, i.e., density modes, are sought in these parameter samples by a mean shift procedure, and the 50% strongest clusters are examined in data space. To this end, the source points $\mathcal{S} \subset \mathbb{R}^3$ are aligned to the target points $\mathcal{T} \subset \mathbb{R}^3$, using the rigid transform A represented by each of the strongest density modes, and an overlap measure between the aligned source and target points is computed,

$$\Omega(A) = \sum_{\boldsymbol{p} \in \mathcal{S}} \frac{1}{1 + 10\|A(\boldsymbol{p}) - \mathrm{nn}_{\mathcal{T}}(A(\boldsymbol{p}))\|} \ , \qquad (1)$$

where $\mathrm{nn}_{\mathcal{T}}(\boldsymbol{p})$ is the point from $\mathcal{T}$ that is the nearest neighbor of $\boldsymbol{p} \in \mathbb{R}^3$ as measured with the surflet distance

$$\begin{aligned} d((\boldsymbol{p}_1, \boldsymbol{n}_1), (\boldsymbol{p}_2, \boldsymbol{n}_2)) = {} & \frac{1}{4}\|\boldsymbol{p}_1 - \boldsymbol{p}_2\| \\ & + \ \frac{1}{100} \arccos(\boldsymbol{n}_1 \cdot \boldsymbol{n}_2) \\ & + \ \|\boldsymbol{p}_1 - \boldsymbol{p}_2\| \arccos(\boldsymbol{n}_1 \cdot \boldsymbol{n}_2) \ . \end{aligned} \qquad (2)$$

The surflet distance takes into account both the Euclidean distance between surface points $\boldsymbol{p}_1, \boldsymbol{p}_2$ and the angle between their local normal vectors $\boldsymbol{n}_1, \boldsymbol{n}_2$. $\Omega(A)$ quantifies the proximity, after transformation with A, of the points in $\mathcal{S}$ to their most similar surflet in $\mathcal{T}$. The alignment $\hat{A}$ yielding the highest overlap $\Omega(\hat{A})$ is selected as the final estimate. Since alignment is sought for two differing shapes, the procedure has to tolerate shape deviations in all its stages.

B. Warping from Source to Target Contacts

Through the alignment, we have established a common reference frame for the surfaces of source and target objects. We can hence guide the search for corresponding points on the target by proximity to source contacts and orientation similarity of the surface patch around these points. Moreover, shape similarity of the source and target surface patches is also regarded in the search process.

We hence have to compute extrinsic (position, orientation) and intrinsic (shape) distances between localized surface patches. However, these distances are not entirely independent from each other. In fact, the quantifiable distance of shape depends on the assumed rigid alignment of two surface patches. Let $\boldsymbol{c} \in \mathcal{S}$ be a source contact, $\boldsymbol{n_c}$ its local surface normal, and $\mathcal{P_c} \subset \mathcal{S}$ the patch of surface points from a spherical neighborhood around $\boldsymbol{c}$. Likewise, $\boldsymbol{p} \in \mathcal{T}$ is a point from the target, $\boldsymbol{n_p}$ its normal, and $\mathcal{P_p} \subset \mathcal{T}$ the surrounding surface patch. Before comparing the shapes of the patches $\mathcal{P_c}$ and $\mathcal{P_p}$, a sequence of rigid motions of $\mathcal{P_c}$ maps the source contact surflet onto the target surflet, i.e., $(\boldsymbol{p}, \boldsymbol{n_p}) = M_\phi(\boldsymbol{c}, \boldsymbol{n_c})$, with

$$M_\phi = R_{\boldsymbol{p},\boldsymbol{n_p},\phi} \circ R_{\boldsymbol{p},\hat{A}(\boldsymbol{n_c}),\boldsymbol{n_p}} \circ T_{\boldsymbol{p}-\hat{A}(\boldsymbol{c})} \circ \hat{A} \ . \tag{3}$$

Here $\hat{A}$ is the alignment estimated for the source and target shapes (cf. Section II-A), $T_{\boldsymbol{p}-\hat{A}(\boldsymbol{c})}$ is a translation by vector $\boldsymbol{p}-\hat{A}(\boldsymbol{c})$, $R_{\boldsymbol{p},\hat{A}(\boldsymbol{n_c}),\boldsymbol{n_p}}$ is a rotation about point $\boldsymbol{p}$ that aligns $\hat{A}(\boldsymbol{n_c})$ to $\boldsymbol{n_p}$, and $R_{\boldsymbol{p},\boldsymbol{n_p},\phi}$ is a rotation about $\boldsymbol{p}$ around the $\boldsymbol{n_p}$ axis by an arbitrary angle ϕ.[1] The global alignment $\hat{A}$ of source and target shapes does not enter into the comparison of source and target patches. So a measure that captures trade-offs between the extrinsic and intrinsic surface-patch distances is given by

$$\begin{aligned} D_\phi(\mathcal{P_c}, \mathcal{P_p}) = \alpha \, &\|\boldsymbol{p} - \hat{A}(\boldsymbol{c})\| \\ &+ \beta \, \angle(R_{\boldsymbol{p},\boldsymbol{n_p},\phi} \circ R_{\boldsymbol{p},\hat{A}(\boldsymbol{n_c}),\boldsymbol{n_p}}) \\ &+ \gamma \, \Delta(M_\phi(\mathcal{P_c}), \mathcal{P_p}) \ . \end{aligned} \tag{4}$$

The first term gives the positional distance between source and target patches, the second term gives the angle of their relative orientation, and the third term gives their shape distance. The latter is quantified by the minimum displacement of points, after rigid motion of $\mathcal{P_c}$ through M_ϕ, needed to deform one patch of points into the other, i.e.,

$$\Delta(\mathcal{P}_1, \mathcal{P}_2) = \frac{1}{|\mathcal{P}_1|} \sum_{\boldsymbol{p} \in \mathcal{P}_1} \|\boldsymbol{p} - C_2(\boldsymbol{p})\| \ . \tag{5}$$

Here $|\mathcal{P}_1|$ is the number of points in patch $\mathcal{P}_1$, which is not larger than in patch $\mathcal{P}_2$ (given proper assignments to $\mathcal{P_c}$ and $\mathcal{P_p}$); $C_2(\boldsymbol{p}) \in \mathcal{P}_2$ is the corresponding point from $\mathcal{P}_2$ assigned to $\boldsymbol{p} \in \mathcal{P}_1$ through solving the related optimal assignment problem, i.e., minimizing the sum of distances in (5). Possible assignments obey the usual constraints to enforce uniqueness, yielding an injection of the points from $\mathcal{P}_1$ into $\mathcal{P}_2$ [17].

[1]The transforms are taken to act on point and normal vector in the respectively adequate way.

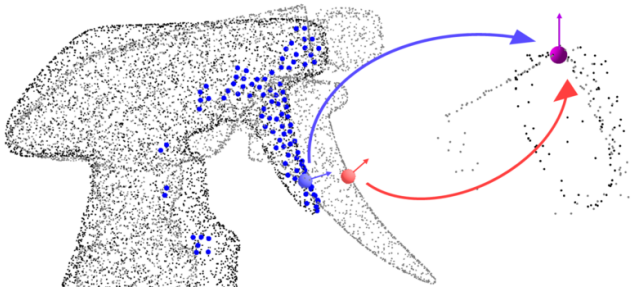

Fig. 2. Left: Source contact (red sphere with arrow) among points of the source object (gray), aligned to points of the target object (black). The warped contact (blue sphere with arrow) is chosen from the 100 target points (small blue dots) closest to the source contact, as measured through (2). The warped contact minimizes the distance (4) over this set. Right: The distance (4) measures the amount of motion between source contact surflet and target surflets (aligned source/target surflets shown as purple sphere and arrow), and the deformation required to conform the surrounding point patches.

Now, our objective is to place each prospective target contact close to a source contact and on a target surface patch with orientation and shape similar to the source contact patch. We thus need to minimize for each source contact $\boldsymbol{c}$ the distance measure $D_\phi(\mathcal{P_c}, \mathcal{P_p})$ with respect to target point $\boldsymbol{p}$ and angle ϕ. The sought target contact is then given by the minimizing point, i.e.,

$$\hat{\boldsymbol{p}} = \arg\min_{\boldsymbol{p} \in \mathcal{T}} \min_{\phi \in (-\pi, \pi]} D_\phi(\mathcal{P_c}, \mathcal{P_p}) \ . \tag{6}$$

All spatial distances were measured in meters and angles in radians. The radial size of the surface patches was set to 0.013, chosen such that each patch contained roughly around 100 surface points at the point density used for this study. The relative weights α, β, γ of the distance terms in (4) were determined through applying some nonlinear deformations to one of our test objects, and searching for corresponding point pairs across the set of thus generated shapes. Parameters were then chosen so as to minimize the maximum occurring positional error between the point $\hat{p}$ and the true correspondence known from the applied deformation. Furthermore, to generate a small number of alternative contacts to be processed by the grasp replanner (cf. Section III), the weight α of the Euclidean distance between source contact and target point was varied within a small range, and the resulting variants of optimal target contacts accumulated. We thus have used $\alpha \in \{1/0.03, 1/0.02, 1/0.01\}$, $\beta = 1$, $\gamma = 1/0.0025$, effectively rescaling the spatial distance terms to suitable length units.

For computational efficiency, the domain of target points over which the minimum of $D_\phi(\mathcal{P_c}, \mathcal{P_p})$ is sought was limited to the 100 points from $\mathcal{T}$ closest to the aligned source contact $\hat{A}(\boldsymbol{c})$, as measured again by the surflet distance (2). Points $\boldsymbol{p}$ further away are not likely to yield a distance $D_\phi(\mathcal{P_c}, \mathcal{P_p})$ below those close points. Likewise, the angles of the checked rotations of the source patch around each target normal were restricted to $\phi \in \{-3\pi/20, -\pi/10, -\pi/20, 0, \pi/20, \pi/10, 3\pi/20\}$, as larger rotations are stronger penalized in D_ϕ. Fig. 2 illustrates the complete contact warping procedure.

III. Grasp replanning technique

Once the warping procedure has defined a set of contact points (and their normals) on the target object, a feasible hand pose that guarantees a collision-free force-closure grasp must be obtained. Its contacts must be close to the warped ones, so original functionality can be preserved on the target object. To increase the chances of retrieving a grasp with these qualities, a replanning step is triggered to explore variations around the hand pose and contacts predicted by the warping procedure, thus obtaining a set of feasible force-closure grasps from which the best functionally fitting grasp can be chosen. The pseudo code for the complete grasp replanning pipeline is presented in Algorithm 1, and the individual phases are described in detail below.

A. Computation of a Feasible Grasp Configuration

Reachability of the Grasp as an Optimization Problem: A grasp configuration is defined here as a pair $G = (T, Q)$ of a hand pose $T \in \mathbb{R}^6$ (position and orientation of the wrist), and corresponding joint values $Q \in \mathbb{R}^k$ for the k degrees of freedom of the hand.

Obtaining the configuration G starting with the contact points can be formulated as an optimization problem [9]. First, an expected set of contact locations is defined on the hand surface. For each finger, reference points are defined in the center of each link, as Fig. 3a shows for the DLR-HIT hand II. Reference points on the palm were not predefined due to the large contact area and the difficulty to guarantee similar contact locations for objects with different geometries. The position $\boldsymbol{r}_i$ and normal direction $\boldsymbol{n}_{\boldsymbol{r}_i}$ for each reference point can be easily computed using forward kinematics.

The cost function for the optimization process is defined as the distance, both angular and Euclidean, between the expected contact points on the hand and the warped contact points on the target object. Let $\boldsymbol{l}_i = \boldsymbol{p}_i - \boldsymbol{r}_i$ be the vector from a reference contact point on the hand, $\boldsymbol{r}_i$, to the corresponding warped contact point on the target object, $\boldsymbol{p}_i$. The cost function to be minimized is then given by

$$f_c = \sum_i \kappa_i \boldsymbol{l}_i^2 + \rho \sum_i \left(1 - \frac{\boldsymbol{n}_{\boldsymbol{p}_i} \cdot \boldsymbol{l}_i}{\|\boldsymbol{l}_i\|}\right) \tag{7}$$

where κ_i are weighting factors that allow ignoring reference points that have no contact on the source grasp or that permit increasing the relative importance of a certain finger (for instance, a finger that should be located on a trigger), and ρ is a scaling factor required to make compatible the range of useful distances with the range of the normalized dot products for the directions ($\rho = 10$ in our experiments).

The optimization brings the reference points on the hand as close as possible to the corresponding target points on the object by minimizing (7). This is a non-linear problem in a high dimensional space; the solution is obtained by using a Sequential Least-Squares Quadratic Programming algorithm (SLSQP) from the NLopt library [18]. The initial guess for the solution is also provided by the warping process. Once the source and target objects are aligned (cf. Section II-A), the hand pose for the source grasp is transformed to the target object, thus providing a transformed hand pose T_{w} that serves as a starting point for the SLSQP algorithm. The initial guess for the joint values, Q_{oh}, correspond to the hand in the fully-opened configuration.

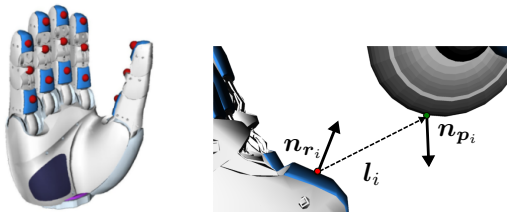

(a) (b)

Fig. 3. (a) Reference contact points $\boldsymbol{r}_i$ are defined for each finger of the DLR-HIT hand II, and depicted here as red spheres; (b) The differences $\boldsymbol{l}_i$ in position between the reference and the warped contact points are minimized, with additional penalty for surface-tangential offset, to find a proper grasp configuration.

Collision Avoidance for the Grasp: The outcome of the optimization method is a grasp configuration $(T_{\mathrm{i}}, Q_{\mathrm{i}})$, which might have self-collisions between the fingers or collisions with the object, as there are no constraints to avoid them in the optimization formulation. First, self-collisions are solved by using a dynamic-based simulation environment built on top of Openrave [19]. Then, the feasibility of the hand motion must be verified, i.e. there must be no collision between the hand and the object when the hand is in the pregrasp pose (considered as the fully-opened configuration). Again, the dynamic simulation is employed; the object is assumed to be static, so we let the hand adapt its configuration around the object. Collisions are removed by letting the fingers/palm move along the perpendicular direction to the force gradient at the local collision area.

Robustness of the Grasp: Once a feasible hand configuration is obtained, fingers are closed around the object and the final contact points are obtained. Those contact points must fulfill the force closure (FC) property, to guarantee a minimum robustness for the grasp. Let the position $\boldsymbol{p}_i$ of the points on the object surface be given with respect to the center of mass (CM). The force $\boldsymbol{f}_i$ applied at a point $\boldsymbol{p}_i$ generates a torque $\boldsymbol{\tau}_i = \boldsymbol{p}_i \times \boldsymbol{f}_i$ around CM. Both elements can be combined in a wrench vector, $\boldsymbol{\omega}_i = (\boldsymbol{f}_i \; \boldsymbol{\tau}_i)^T$.

The friction model is described with Coulomb's law. To simplify computations, the friction cone is linearized using an m-sided polyhedral convex cone. The wrench $\boldsymbol{\omega}_{ij}$ generated by a unitary force $\boldsymbol{f}_i$ along an edge j of the linearized friction cone is called a *primitive wrench*. A grasp defined by the set of contact points $\mathcal{C} = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$ is associated to the set $\mathcal{W} = \{\boldsymbol{\omega}_{11}, \ldots, \boldsymbol{\omega}_{1m}, \ldots, \boldsymbol{\omega}_{n1}, \ldots, \boldsymbol{\omega}_{nm}\}$ of primitive contact wrenches.

A grasp is FC if and only if the origin O of the wrench space lies strictly inside the convex hull of $\mathcal{W}$, represented as $CH(\mathcal{W})$ [20]. The desired FC property means that the grasp is able to resist external disturbances in any direction.

B. Local Replanning

Even after obtaining a feasible grasp configuration that provides a FC grasp on the object, there is no guarantee

Algorithm 1: Grasp replanning pipeline

Given:
- Set of pairs of surflets (warped points and normals) on the target object $\mathcal{Z} = \{(\boldsymbol{p}_i, \boldsymbol{n}_{\boldsymbol{p}_i})\}$
- Transformed hand pose for the target object, T_{w}

Output: Best functional grasp configuration G_{best}

// Computation of a feasible grasp configuration

$G_{\mathrm{i}} = (T_{\mathrm{i}}, Q_{\mathrm{i}}) \leftarrow$ Optimization of the cost function f_c in Eq. (7), with initial guess $(T_{\mathrm{w}}, Q_{\mathrm{oh}})$

$G_{\mathrm{o}} = (T_{\mathrm{o}}, Q_{\mathrm{o}}) \leftarrow$ Removing self-collisions and hand/object collisions from G_{i}

Initialize an empty set of grasp configurations $\mathcal{G} = \{\}$

$\mathcal{C}(G_{\mathrm{o}}) \leftarrow$ Contacts on the object after closing the fingers around it, starting from G_{o}

if *FCtest*$(\mathcal{C}(G_{\mathrm{o}}))$ **then** $\mathcal{G} = \mathcal{G} \cup \{G_{\mathrm{o}}\}$

// Local replanning

// Pose sampling around G_{i}

while *counter* $< h_p$ **do**

  $G_{\mathrm{s}} \leftarrow$ PoseSampling(G_{i})

  $G_{\mathrm{c}} \leftarrow$ Removing self-collisions and hand/object collisions from G_{s}

  $\mathcal{C}(G_{\mathrm{c}}) \leftarrow$ Contacts on the object for G_{c}

  if *FCtest*$(\mathcal{C}(G_{\mathrm{c}}))$ **then** $\mathcal{G} = \mathcal{G} \cup \{G_{\mathrm{c}}\}$

// Sampled realignment of the hand pose

while *counter* $< h_r$ **do**

  $G_{\mathrm{r}} \leftarrow$ Collision-free realignment of grasp configuration G_{i} along a sampled direction

  $\mathcal{C}(G_{\mathrm{r}}) \leftarrow$ Contacts on the object for G_{r}

  if *FCtest*$(\mathcal{C}(G_{\mathrm{r}}))$ **then** $\mathcal{G} = \mathcal{G} \cup \{G_{\mathrm{r}}\}$

// Evaluation of functionality

foreach $G \in \mathcal{G}$ **do** evaluate f_g ; // Eq. (9)

$G_{\mathrm{best}} \leftarrow G$ with minimum f_g

return G_{best}

that the grasp will fulfill the desired functionality, as will be illustrated in Section IV. To increase the chances of obtaining functional FC grasps, a set of candidate grasps is created by producing small variations on the obtained FC grasp configuration. At the end, the best functional grasp from the set of candidate grasps is chosen according to a functionality criterion.

Sampling of the Hand Pose: A number h_p of new hand poses T is generated by taking samples around the pose T_{i}, the one obtained in the first optimization process. The pose for each candidate grasp configuration is obtained by independently taking samples in the 3D space of positions and the 3D space of orientations.

The isotropic normal distribution for the position variation $\boldsymbol{z}$ is

$$f_{3D}(\boldsymbol{z}; \boldsymbol{\sigma}_z) = \frac{\exp -\frac{\|z\|^2}{2\boldsymbol{\sigma}_z^2}}{(\sqrt{2\pi}\boldsymbol{\sigma}_z)^3},\ \boldsymbol{z} \in R^3,\ \boldsymbol{\sigma}_z \in R^3 \qquad (8)$$

A random uniform distribution for unit quaternions $\boldsymbol{q}$ representing 3D orientations can be obtained using one single random value $v \in [0, 1]$, following the approach proposed in [21]. An additional filtering step is implemented here to guarantee that the difference in orientations with respect to the original ones is below a predefined limit.

Each sample taken at this stage is checked for feasibility and robustness, and samples leading to a FC grasp are added to a set $\mathcal{G}$ of valid grasp configurations. The number h_p of samples is empirically chosen as a compromise between variation in grasp configurations and computational time required to evaluate the samples. For the experiments presented in Section IV, $h_p = 100$.

Sampled Realignment of the Hand Pose: When checking for the feasibility of grasps, i.e. removing self-collisions and hand-object collisions (cf. Section III-A), the dynamic engine dictates the trajectory that the hand follows to eliminate the collisions. However, this might lead to grasp configurations that are far from the initial solution of the optimization problem, G_i, which contains valuable information coming from the warping process. To compensate for this effect, an additional sampling was created, applying artificial motions to the hand inside the dynamic engine to force the motion to remove collisions along certain preferred (sampled) directions. These preferred motions are positive and negative rotations around the principal axes, plus combination of rotations along those axes. Again, each sample taken at this stage is checked for feasibility and stability, and samples leading to a FC grasp are added to the set $\mathcal{G}$ of valid grasp configurations. For the experiments in Section IV, the number h_r of samples at this stage was kept low, $h_r = 9$.

C. Selection of the Best Functional Grasp

After completing the replanning phase, there is a pool $\mathcal{G}$ of grasp configurations that lead to a FC grasp on the target object. The final step is then choosing the grasp from this set that best fulfills the grasp functionality demonstrated in the source grasp.

A function to evaluate such grasp functionality is:

$$f_g = \sum_i \kappa_i d((\boldsymbol{p}_i, \boldsymbol{n}_{\boldsymbol{p}_i}), (\boldsymbol{r}_i, \boldsymbol{n}_{\boldsymbol{q}_i})) \qquad (9)$$

where d is the surflet distance defined in (2), taken between the surflet $(\boldsymbol{p}_i, \boldsymbol{n}_{\boldsymbol{p}_i})$ for the warped points and a virtual surflet $(\boldsymbol{r}_i, \boldsymbol{n}_{\boldsymbol{q}_i})$ with the position of the reference points $\boldsymbol{r}_i$ and the normal $\boldsymbol{n}_{\boldsymbol{q}_i}$ at the actual contact points obtained when the fingers are closed around the object, starting with a grasp configuration $G \in \mathcal{G}$. The factors κ_i weigh the relative importance of the contact points, as in (7). Note that the functionality evaluation with f_g compares the contact points for the candidate grasps with the points coming from the warping process, as the warping process keeps most of the information on relative position of the contacts, a key factor to produce a given functionality on an object. The lower the f_g, the better the functionality is for the analyzed grasp.

IV. Experiments

A. Test Data

The proposed method has been tested on four different categories of objects and with a variety of grasps for each

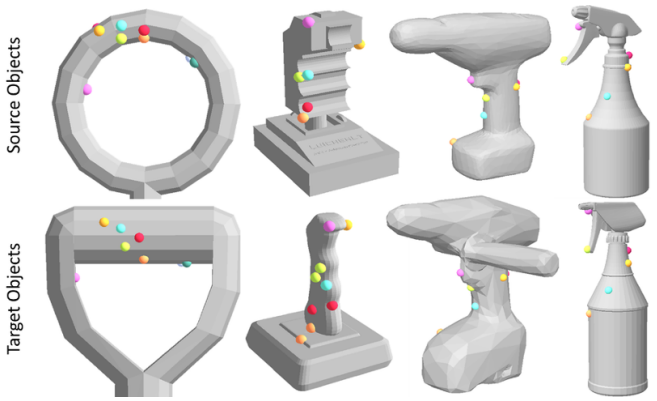

Fig. 4. Contact points for a grasp on each object category from our test set. Top row: contacts of grasps designed on source objects. Bottom row: contacts warped from the above source objects onto target objects.

category. This collection of objects was chosen considering their wide range of shapes and the variation on functional grasp types they afford. In detail, we have run tests with

- 2 source grasps on 1 spray bottle, each transferred to 2 target spray bottles,
- 1 source grasp on each of 2 joysticks, each transferred to 2 target joysticks,
- 4 source grasps on 1 shovel handle, each transferred to 3 target shovel handles,
- 2 source grasps on each of 3 drills, each transferred to 2 target drills.

The object models were taken from the ROS household object database, 3DWarehouse, Digital-X models, and the Princeton Shape Benchmark. In total, the transfer of 32 grasps onto 11 target objects was tested.

B. Success Rates

The analysis of the results is carried out in a stage-wise manner. Thus, we address the outcome of each stage of the complete pipeline individually, along with the finally returned grasps on the target objects.

In our test set of source/target object combinations and grasp variants, there were in total 250 contact points to be warped from a source to a target. For 195 source contacts there was a single warped contact created, for the other 55 source contacts a pair of contact alternatives resulted from the variation of the weight of the positional distance term in (4) (cf. Section II-B).

In Fig. 4 we show for each tested object category one set of grasp contacts on a source object and their warps to a target object. Corresponding source/target contacts are marked with spheres of the same color, pairs of same-colored spheres on the target correspond to the same source contact. Another example of contact warps is shown in Fig. 1.

The grasps generated in the transfer process were analyzed for their reachability, feasibility, robustness, and functionality. From the 32 grasps to be transferred, 30 turned out reachable and stable, i.e., they provided force closure. Out of these, 28 were also functionally equivalent to their source version, as judged by visual inspection. We hence achieved an allover success rate of 87.5% of stable and functional transferred grasps.

TABLE I

NUMBER AND RATIO OF SUCCESSFULLY TRANSFERRED GRASPS PER OBJECT CATEGORY AND PER GENERATING REPLANNING PHASE

	Spray bottles	Joysticks	Shovels	Drills	Total
Feasibility	0/3	1/4	1/11	1/10	3/28
Sampling	2/3	2/4	7/11	7/10	18/28
Realignment	1/3	1/4	3/11	2/10	7/28

TABLE II

MEAN RATIO OF THE FORCE CLOSURE QUALITY VALUES OF TRANSFERRED FUNCTIONAL GRASPS AND THEIR SOURCE GRASP

	Spray bottles	Joysticks	Shovels	Drills	Average
FC ratio	26.43	4.38	2.04	6.88	7.38

The functionally transferred grasps are broken down in Table I by object category and process stage in which they were generated. The ratios are given with respect to the number of functional grasps in each category or to the total of 28 functional grasps. Reading the numbers in the last column of the table, we observe that the success rate of the feasibility stage alone is only 10% due to the low number of just four stable grasps generated at this stage. All four were actually functional but two of those were variants of the same grasp produced on the same target, so only three could have been, and indeed were selected in the end.

The implication, however, is that whenever the warped contacts were reachable and allowed force closure, the resulting grasp was also functional, and even outscored at the selection stage all the other grasps generated later in the replanning pipeline. But the kinematic constraints of the hand and the force closure criterion proved difficult to fulfill simultaneously when taking the warped contacts as the goal contact positions. Hence the activation of the later stages of replanning are necessary for exploring more grasp variations.

Table I shows that the pose sampling stage of replanning is the most successful one, though not completely failsafe. Hence there were 25% of solutions generated only by the realignment stage of replanning. These numbers justify the chain of processes for grasp replanning, set up in order to maximize the number of functionally transferred grasps.

The mean ratios of the force-closure quality values obtained for the transferred functional grasps with respect to their source grasp are shown in Table II for each object category. It is worth noting that the grasp robustness has on average substantially improved through the transfer, a comforting safety feature for real applications.

The total number of grasp variants explored in the replanning stages was determined by the number of contact variants proposed by the warping process as well as by the variations introduced in the sampling stages of replanning. The average total time to replan a grasp was 12.63 min per variant of warped contact set. The number of such contact sets was between one and 64 for each source grasp. As such, the present

TABLE III

SUCCESS RATES OF THE PROPOSED WARP-BASED TRANSFER METHOD IN COMPARISON WITH THE NAIVE APPROACH

	Spray bottles	Joysticks	Shovels	Drills	Total
With warps	3/4	4/4	11/12	10/12	28/32
Naive	2/4	2/4	6/12	7/12	17/32

procedure is clearly for offline grasp transfer only. To make it an online procedure, an intermediate stopping criterion can be introduced to terminate the replanning process as soon as a functional grasp is encountered, or less promising grasp variants can be deleted early in the pipeline. Also, as each variant is processed independently of the others, the complete procedure can efficiently be parallelized.

All the source grasps along with all the grasps transferred to the target objects are presented in Fig. 5. Each source grasp is arranged in a group with all its transferred versions, with the source being leftmost. The grasps can also be viewed in the video attachment to this paper.

C. Failure Cases

We have observed four cases in which the transfer pipeline has failed in the end to generate a simultaneously functional and stable grasp.

The non-stable ones fail to meet the force-closure criterion, the respective target object is displayed without hand in Fig. 5. Thus, the side grasp on the shovel handle failed to be transferred to one shovel in 5(f), the difficulty caused mainly by the very limited space inside the ring of the handle. The other unstable case was the trigger grasp on the drill in 5(k).

The non-functional grasps have contacts on the target object not corresponding to those on the source object. They are marked with a red box overlay in Fig. 5. For the case of the spray bottle in 5(b), the problem occurred at the final grasp selection: although there were a couple of simultaneously functional and force-closure grasps generated, none of these was close enough to the warped contact set. For the case of the drill in 5(i), there were no functional grasps generated.

D. Comparison With a Naive Transfer Approach

We also compared the proposed transfer method with a naive procedure for adapting a source grasp to a target object. In the naive procedure, only the hand pose is transferred through the global alignment transformation estimated initially for the contact warping (cf. Section II-A). A grasp is then attempted with the transformed hand pose, not considering any warped contacts on the target object but simply closing the fingers around the object. In order to isolate in this test the relevance of the warping and kinematic optimization procedures, the rest of the replanning pipeline (collision avoidance, pose sampling, and realignment) was still executed for finding a grasp. The final grasp selection was done by choosing the one minimizing the Euclidean distance of the vector of joint angles to the one of the source grasp. Our original criterion (9) uses the warped contacts as a reference and is simply not applicable without the warps.

As shown in Table III, with the naive procedure the success rate drops from 87.5% down to 53.0%. It is hence evident that warped contacts and their usage as targets for the hand are of critical importance for a useful transfer procedure.

V. CONCLUSIONS

The proposed warp-based grasp transfer method produced good functional grasps with high success rate on diverse objects, and significantly exceeds the performance of a naive grasp planning method. However, there were still a few failed cases, partly due to the final grasp selection step not picking a functional variant from the generated set, although that final grasp set did contain functional grasps. Improvements on the functionality prediction will hence be investigated.

REFERENCES

[1] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis - a survey," *IEEE Trans. Robotics*, vol. 30, no. 2, pp. 289–309, 2014.

[2] C. Goldfeder, M. Ciocarlie, H. Dang, and P. Allen, "The Columbia grasp database," in *Proc. IEEE Int. Conf. Robotics and Automation - ICRA*, 2009, pp. 1710–1716.

[3] N. Curtis and J. Xiao, "Efficient and effective grasping of novel objects through learning and adapting a knowledge base," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems - IROS*, 2008, pp. 2252–2257.

[4] 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, no. 4, pp. 732–747, 2007.

[5] J. Aleotti and S. Caselli, "Part-based robot grasp planning from human demonstration," in *Proc. IEEE Int. Conf. Robotics and Automation - ICRA*, 2011, pp. 4554–4560.

[6] D. Song, C. Ek, K. Huebner, and D. Kragic, "Embodiment-specific representation of robot grasping using graphical models and latent-space discretization," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS*, 2011, pp. 980–986.

[7] P. Sandilands, V. Ivan, T. Komura, and S. Vijaykumar, "Dexterous reaching, grasp transfer and planning using electrostatic representations," in *Proc. IEEE-RAS Int. Conf. on Humanoid Robots*, 2013, pp. 211–218.

[8] H. Dang and P. Allen, "Semantic grasping: Planning robotic grasps functionally suitable for an object manipulation task," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS*, 2012, pp. 1311–1317.

[9] M. Ciocarlie and P. Allen, "Hand posture subspaces for dexterous robotic grasping," *Int. J. Robotics Research*, vol. 28, no. 7, pp. 851–867, 2009.

[10] M. Hjelm, R. Detry, C. Ek, and D. Kragic, "Representations for cross-task, cross-object grasp transfer," in *Proc. IEEE Int. Conf. Robotics and Automation - ICRA*, 2014, pp. 5699–5704.

[11] A. Bronstein, 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.

[12] 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.

[13] U. Hillenbrand, "Non-parametric 3D shape warping," in *Proc. Int. Conf. Pattern Recognition*, 2010.

[14] U. Hillenbrand and M. A. Roa, "Transferring functional grasps through contact warping and local replanning," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS*, 2012, pp. 2963–2970.

[15] H. Ben Amor, O. Kroemer, U. Hillenbrand, G. Neumann, and J. Peters, "Generalization of human grasping for multi-fingered robot hands," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS*, 2012, pp. 2043–2050.

[16] 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.

[17] R. Vanderbei, *Linear programming: foundations and extensions*. Springer, 2001.

[18] S. G. Johnson, "The nlopt nonlinear-optimization package," 2010. [Online]. Available: http://ab-initio.mit.edu/nlopt

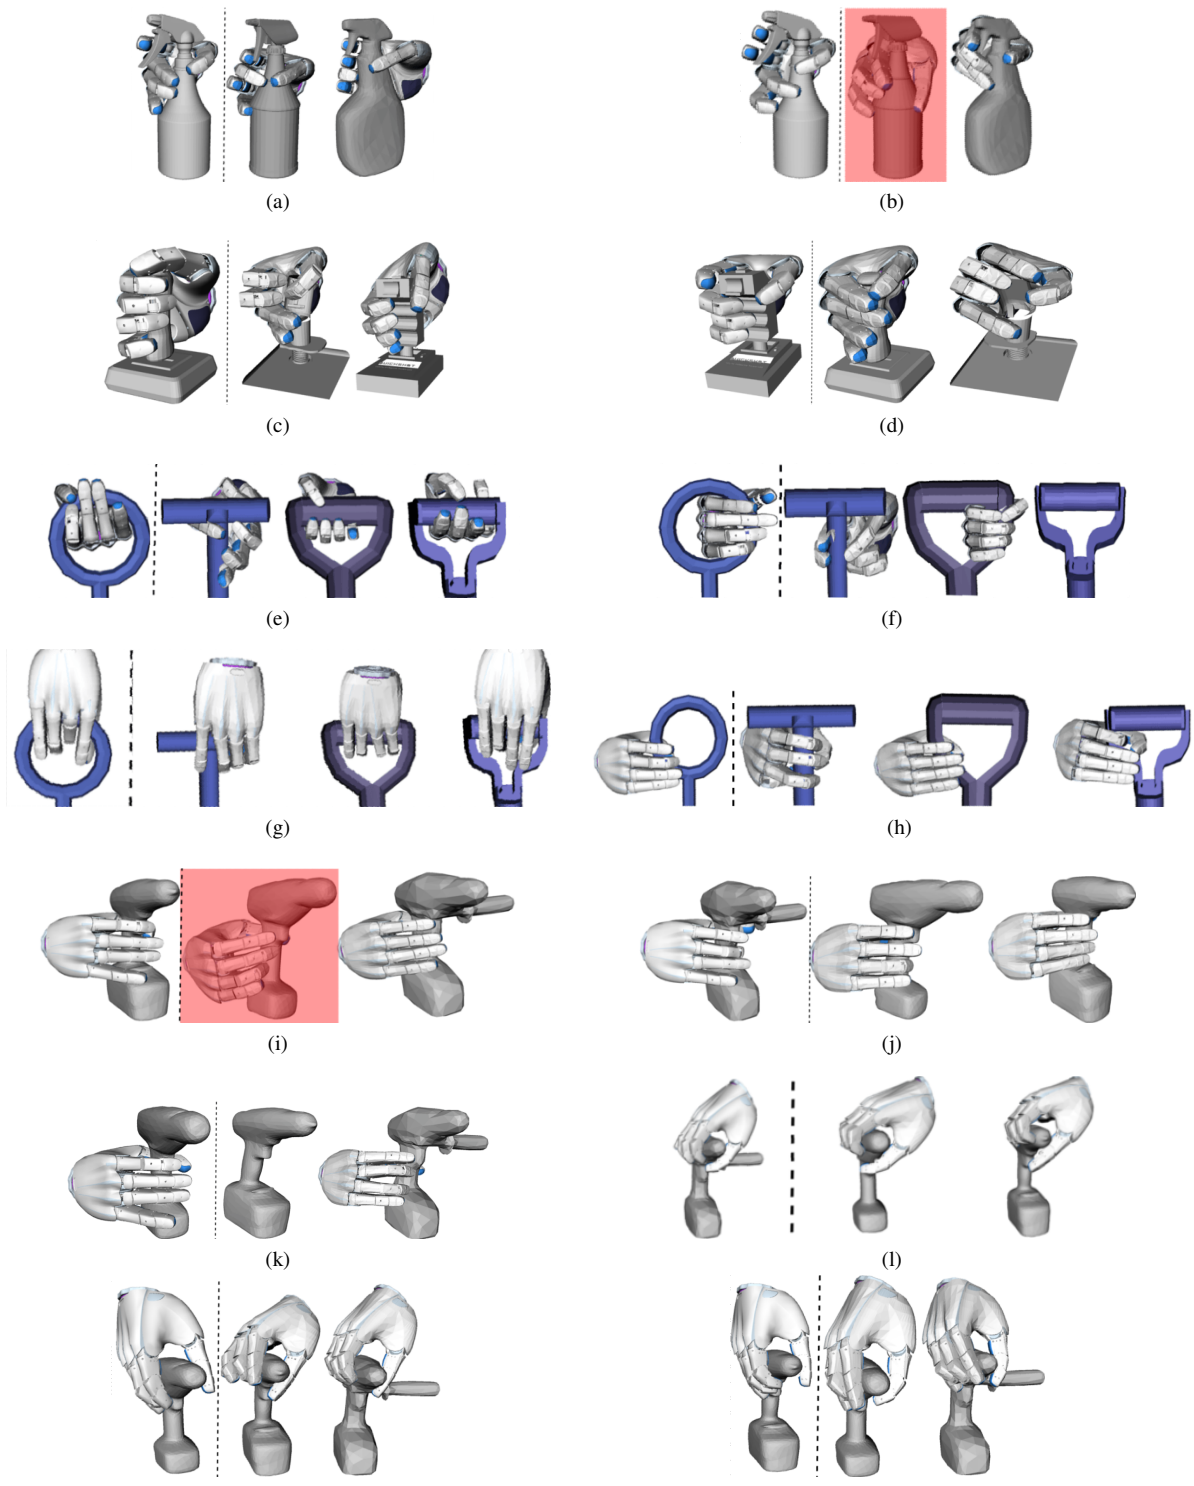

Fig. 5. All source and transferred grasps from our test set. In all the figures, the source grasp is shown on the left side of the dashed line, it is transferred to the target objects on the right side. Trigger-type grasps are shown in (a), (b), (i), (j),(k); handle-type grasps in (e), (g), (f), (h); joystick-like grasps in (c), (d); support-type grasps in (l), (m), (n). In two cases where the hand does not appear in the figure, no stable grasp was found in the transfer; the red overlay marks two cases of non-functional transferred grasps.

[19] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, Carnegie Mellon University, 2010.

[20] R. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. Boca Ratón, Florida: CRC Press, 1994.

[21] J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *Proc. IEEE Int. Conf. Robotics and Automation - ICRA*, 2004.