

# Transferring Category-based Functional Grasping Skills by Latent Space Non-Rigid Registration

Diego Rodriguez and Sven Behnke

**Abstract**—Objects within a category are often similar in their shape and usage. When we—as humans—want to grasp something, we transfer our knowledge from past experiences and adapt it to novel objects. In this paper, we propose a new approach for transferring grasping skills that accumulates grasping knowledge into a category-level canonical model. Grasping motions for novel instances of the category are inferred from geometric deformations between the observed instance and the canonical shape. Correspondences between the shapes are established by means of a non-rigid registration method that combines the Coherent Point Drift approach with subspace methods. By incorporating category-level information into the registration, we avoid unlikely shapes and focus on deformations actually observed within the category. Control poses for generating grasping motions are accumulated in the canonical model from grasping definitions of known objects. According to the estimated shape parameters of a novel instance, the control poses are transformed towards it. The category-level model makes our method particularly relevant for on-line grasping, where fully-observed objects are not easily available. This is demonstrated through experiments in which objects with occluded handles are successfully grasped.

**Index Terms**—Dexterous manipulation, Grasping, Multi-fingered hands.

## I. INTRODUCTION

WHILE transferring grasping skills within a category happens frequently and effortless in humans, obtaining that generalization in robots is still an open problem. People can be shown objects that they never saw before, and they often will immediately know how to grasp and operate them. This happens by transferring knowledge from their learned model of the object category, e.g., screw drivers, to novel instances. Although the manipulation of known objects can be planned offline, many open-world applications require the manipulation of unknown instances. Our approach accumulates manipulation knowledge of known instances in category-level models and transfers manipulations skills to novel instances (Fig. 1).

The method presented in this paper focuses on functional grasping, i.e., on motions that allow not only to grasp the object but also to use it. We use the term *grasping* to refer to the




Figure 1: Shape information and grasping knowledge for known object instances are aggregated in a category-level canonical model. Grasping control poses are transferred to novel instances of the category for generating the grasping motion.

process of bringing the object into the hand, and not only to the final configuration of hand and object. We propose a method for generating grasping motions for novel instances by making use of category-level shape information represented by a learned latent shape space. Our method aggregates object shape and grasping knowledge from multiple known instances of a category in a canonical model. The learned latent space of shape variations enables a category-specific shape-aware non-rigid registration procedure that establishes correspondences between a view of a novel object instance and the canonical model. Our method finds a transformation from the canonical model to the view in the latent shape space—linearly interpolating and extrapolating from other transformations found within the category—which best matches the observed 3D points. This estimates the shape parameters of the novel instance and allows for inference of its occluded parts. By the non-rigid transformation and the aggregated manipulation knowledge, control poses for the novel instance are inferred. The grasping motion is finally generated by using those control poses.

In this paper, we extend our previous work [1] by accumulating grasping knowledge in the canonical model in addition to the shape information, which enriches our transferring skill model.

## II. RELATED WORK

### A. Non-Rigid Registration and Shape Spaces

Most of the non-rigid registration methods proposed so far differ mostly by the prior restrictions or regularization on the deformation that the points can undergo. Several restrictions

Manuscript received: November, 21, 2018; Revised January, 26, 2018; Accepted March, 19, 2018.

This paper was recommended for publication by Editor Han Ding upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the European Union's Horizon 2020 Programme under Grant Agreement 644839 (CENTAURO) and the German Research Foundation (DFG) under the grant BE 2556/12 ALROMA in priority programme SPP 1527 Autonomous Learning.

All authors are with the Autonomous Intelligent Systems (AIS) Group, Computer Science Institute VI, University of Bonn, Germany {rodriguez, behnke}@ais.uni-bonn.de

Digital Object Identifier (DOI): see top of this page.

such as conformal maps [2]–[4], isometry [5]–[7], thin-plate splines [8], [9], elasticity [10] and Motion Coherence Theory [11] have been used to encourage or constrain different types of transformations.

For surface reconstruction, many methods use non-rigid registration [12]–[15]. Approaches such as presented by Li *et al.* [12] and Zollhöfer *et al.* [16] sequentially add higher frequency details coming from new depth camera frames to a low-resolution 3D capture through non-rigid registration.

For category-based shape spaces, several methods have been proposed. Hasler *et al.* [17] generate a shape space of human bodies with poses using 3D markers and human scans. Burghard *et al.* [18] developed a shape space of varying geometry based on dense correspondences. Engelmann *et al.* [19] define a shape manifold which models intra-class shape variance; this method is robust with noisy or occluded regions.

### B. Transferring Grasping Skills

Based on segmented objects according to their RGB-D appearance, Vahrenkamp *et al.* [20] transfer grasp poses from a set of template grasps. Ficuciello *et al.* [21] developed an approach to confer grasping capabilities based on a reinforcement learning technique and postural synergies. In [22] and [23], functional grasp poses are warped such that distance between correspondences is minimized, then the warped poses are replanned in order to increase the functionality of the grasp. In [24] a similar contact warping is combined with motor synergies to generalize human grasping. Stueckler *et al.* [25] transfer manipulation skills using a non-rigid registration method based on multi-resolution surfel maps. The non-rigid registration serves as the mechanism to warp available grasping poses.

### C. Discussion

Although current state-of-the-art methods for non-rigid registration yield good results, they have some limitations. Newcombe *et al.* [15] use optical flow constraints and thus this approach does not perform well with large deformations or changes in color and illumination. Moreover, several captures of the object are required. The method by Burghard *et al.* [18] accurately estimates dense correspondences, but does not perform well with incomplete scans or noisy data. To solve these problems, we incorporate category-level information in our approach, such that we are able to register partially-occluded novel instances using a single capture of the object. Methods such as Engelmann *et al.* [19] deal with minor misalignments and occlusions, but do not offer correspondences between points and do not give any kind of transformation. Our method, on the other hand, offers a transformation for each point of the novel instance and even points that do not belong to the object which allows us to transform grasp poses.

Regarding transferring grasping skills, we tackle the problem of requiring a fully observed [22] or a non-occluded [20] object by exploiting the geometrical information residing in our learned categorical model. Unlike [25] we model shape and grasping not for single known instances, but for object categories, which gives us the possibility to learn typical shape

variations and to infer grasping information even when parts of the object are not observed. More importantly, none of previous approaches is able to accumulate and to use knowledge from *several* previous successfully experiences, which is the main focus of this paper.

## III. METHOD

Our approach is composed of a learning phase and an inference phase (Figs. 2 and 3). In the learning phase, a category-specific linear model of the transformations that a category of objects can undergo is built. In this manner, poses in the space of the canonical shape can be transformed into the space of an observed instance. These poses can be added even after the learning phase. The category-specific linear model is learned as follows: First, we select a single instance from the training dataset to be the canonical model of the category. Then, we find the transformations relating this instance to all other instances of the category using Coherent Point Drift (CPD) [11]. Finally, we find a linear latent subspace of these transformations, which becomes our transformation model for the category. For each instance in the training set, an associated grasping descriptor  $\varsigma$  (vector representation of the grasping motion) is also transformed into the canonical space. In this manner, multiple experiences can be aggregated in the canonical model.

In the inference phase, given a novel observed instance, our method searches in the subspace of transformations to find the transformation which best relates the canonical shape to the observed instance. Depending on the resulting latent shape variables and the aggregated grasping knowledge accumulated in the canonical model, a grasping descriptor for the novel instance is inferred.

### A. Categories and Shape Representation

A category is composed by a set of objects which share the same topology and have a similar shape. Each category has a canonical shape  $\mathbf{C}$  that will be deformed to fit the shape of the training and testing sample shapes. To represent a shape, we use point clouds, which can be generated from meshes by ray-casting from several viewpoints on a tessellated sphere and then down-sampling with a voxel grid filter. Each category specifies a canonical pose and reference frame, used for initial alignments.

### B. Coherent Point Drift

Here, we shortly describe the Coherent Point Drift (CPD) [11] and how we use it for our non-rigid registration.

CPD estimates a deformation field mapping between a template point set  $\mathbf{S}^{[t]} = (\mathbf{s}_1^{[t]}, \dots, \mathbf{s}_M^{[t]})^T$  and a reference point set  $\mathbf{S}^{[r]} = (\mathbf{s}_1^{[r]}, \dots, \mathbf{s}_N^{[r]})^T$ . The points in  $\mathbf{S}^{[t]}$  are modeled as centroids of a Gaussian Mixture Model (GMM) from which the points in  $\mathbf{S}^{[r]}$  are drawn. CPD maximizes the likelihood of the GMM while imposing constraints on the motion of the centroids such that points near each other should move coherently and have a similar motion to their neighbors [26]. The likelihood of

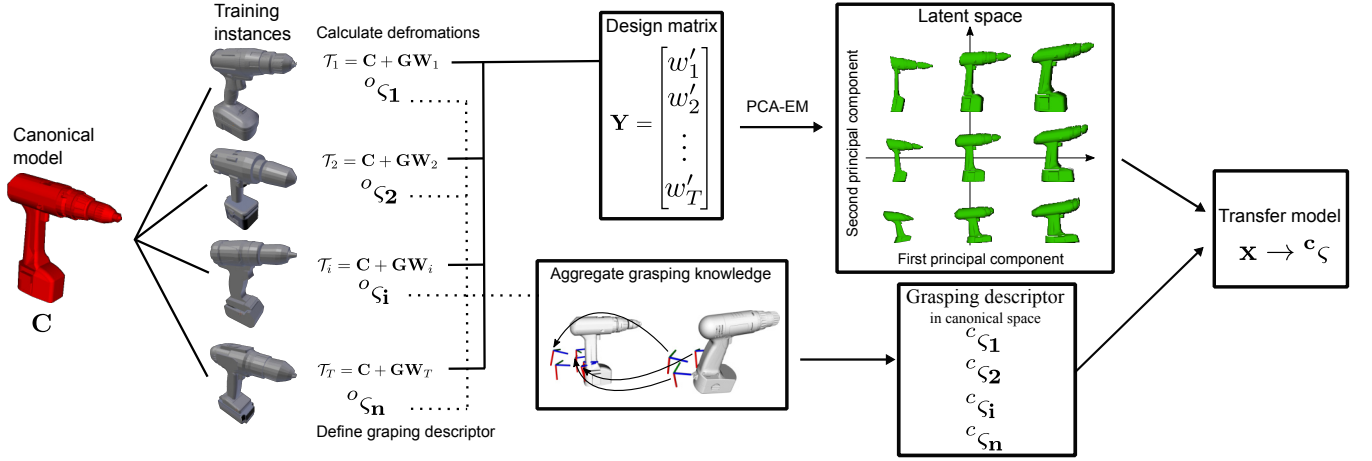


Figure 2: Training phase. The deformations between each instance and the canonical model are calculated using CPD. These deformations are assembled into the design matrix  $\mathbf{Y}$ . Using PCA-EM, the principal components which constitute the latent space are extracted. The grasping descriptor for each training sample is aggregated in the canonical model. The latent variables serve as feature vector while the grasping descriptor is the desired output for the grasping transfer model.

the GMM is not directly maximized, but instead its equivalent negative log-likelihood function is minimized:

$$E(\psi, \sigma^2) = - \sum_{n=1}^N \log \sum_{m=1}^M \exp^{-\frac{1}{2\sigma^2} \|\mathbf{s}_n^{[r]} - \mathcal{T}(\mathbf{s}_m^{[t]}, \psi)\|^2}, \quad (1)$$

where  $\mathcal{T}(\mathbf{s}_m^{[t]}, \psi)$  is a parametrized transformation from the template point set to the reference set, and  $\sigma^2$  is the covariance of the Gaussian density. The transformation  $\mathcal{T}$ , for the non-rigid registration, is defined as the initial position plus a displacement function  $v$ :

$$\mathcal{T}(\mathbf{S}^{[t]}, v) = \mathbf{S}^{[t]} + v(\mathbf{S}^{[t]}). \quad (2)$$

The constraints on the motion of the centroids are realized by regularizing the displacement function  $v$ . Adding this regularization  $\phi(v)$  to the negative log-likelihood Eq. (1), we obtain

$$f(v, \sigma^2) = E(\sigma^2, v) + \frac{\lambda}{2} \phi(v), \quad (3)$$

where  $\lambda$  is a trade-off parameter between the goodness of maximum likelihood fit and regularization. A particular choice of  $\phi(v)$  leads to the following displacement function  $v(\mathbf{Z})$  [11]:

$$v(\mathbf{Z}) = G(\mathbf{S}^{[t]}, \mathbf{Z})\mathbf{W}, \quad (4)$$

for any set of  $D$ -dimensional points  $\mathbf{Z}_{N \times D}$ .  $G(\mathbf{S}^{[t]}, \mathbf{Z})$  is defined as a Gaussian kernel matrix composed element-wise by:

$$g_{ij} = G(\mathbf{s}_i^{[t]}, \mathbf{z}_j) = \exp^{-\frac{1}{2\beta^2} \|\mathbf{s}_i^{[t]} - \mathbf{z}_j\|^2}, \quad (5)$$

$\mathbf{W}_{M \times D}$  is a matrix of kernel weights, and  $\beta$  is a scalar that controls the strength of interaction between points. An additional interpretation of  $\mathbf{W}$  is as a set of  $D$ -dimensional deformation vectors, each associated with one of the  $M$  points of  $\mathbf{S}^{[t]}$ . For convenience in the notation,  $\mathbf{G}_{M \times M}$  will denote  $G(\mathbf{S}^{[t]}, \mathbf{S}^{[t]})$ . Note that  $G(\cdot, \cdot)$  can simply be computed by Eq. (5), but the matrix  $\mathbf{W}$  needs to be estimated.

To minimize Eq. (3), CPD uses an Expectation Maximization (EM) algorithm. In the E-step, the posterior probabilities matrix

$\mathbf{P}$  is estimated using past parameter values. This matrix  $\mathbf{P}$  is composed element-wise by:

$$p_{mn} = \frac{e^{-\frac{1}{2\sigma^2} \|\mathbf{s}_n^{[r]} - (\mathbf{s}_m^{[t]} + G(m, \cdot)\mathbf{W})\|^2}}{\sum_{m=1}^M e^{-\frac{1}{2\sigma^2} \|\mathbf{s}_n^{[r]} - (\mathbf{s}_m^{[t]} + G(k, \cdot)\mathbf{W})\|^2} + \frac{\omega}{1-\omega} \frac{(2\pi\sigma^2)^{\frac{D}{2}}}{N}} \quad (6)$$

where  $\omega$  reflects the assumption on the amount of noise.

In the M-step, the matrix  $\mathbf{W}$  is estimated by:

$$(\mathbf{G} + \lambda\sigma^2 d(\mathbf{P}\mathbf{1})^{-1})\mathbf{W} = d(\mathbf{P}\mathbf{1})^{-1}\mathbf{P}\mathbf{S}^{[r]} - \mathbf{S}^{[t]} \quad (7)$$

where  $\mathbf{1}$  represents a column vector of ones and  $d(\cdot)^{-1}$  is the inverse diagonal matrix. For a more detailed description of the CPD algorithm, please refer to [11].

In our method, we use the canonical shape  $\mathbf{C}$  for the deforming template shape  $\mathbf{S}^{[t]}$  and each training example  $\mathbf{T}_i$  as the reference point set  $\mathbf{S}^{[r]}$ . Therefore, the transformations  $\mathcal{T}_i$  are defined as

$$\mathcal{T}_i(\mathbf{C}, \mathbf{W}_i) = \mathbf{C} + \mathbf{G}\mathbf{W}_i \quad (8)$$

where  $\mathbf{W}_i$  is the  $\mathbf{W}$  matrix computed by taking training example  $\mathbf{T}_i$  as the reference point set  $\mathbf{S}^{[r]}$ .

### C. Latent Space

CPD allows us to define a feature vector representing the deformation field. This vector has the same length for all training examples; additionally, elements in this vector correspond with the same elements in another. This allows us to learn a latent lower-dimensional space.

We observe from Eq. (8) that the deformation field between the canonical and an observed instance is fully determined by  $\mathbf{G}$  and  $\mathbf{W}$ . Moreover, we see that  $\mathbf{G}$  only requires the points of the canonical shape and it remains constant for all training examples. Therefore, the entire uniqueness of the deformation field for each training example is captured by its matrix  $\mathbf{W}$ .

We construct a row vector  $\mathbf{y}_i \in \mathbb{R}^{p=M \cdot D}$  from each matrix  $\mathbf{W}_i$  of each training example  $\mathbf{T}_i$ , that characterizes the

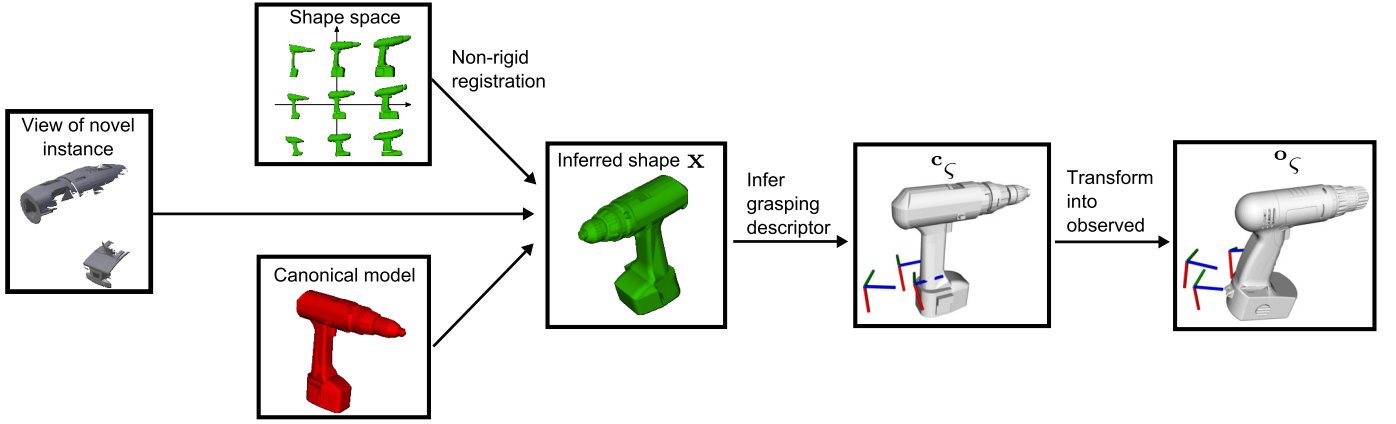


Figure 3: The canonical shape (red) is matched against a partially-occluded target shape (leftmost) by finding its latent shape parameters. The grasping descriptor  $c_\xi$  is inferred from  $\mathbf{x}$ . Finally, the descriptor is transformed to the observed space.

corresponding deformation field. The vectors are normalized to have zero-mean and unit-variance and are then assembled into a design matrix  $\mathbf{Y}$ . Finally, we find a lower-dimensional manifold of deformation fields for the category by applying the Principle Component Analysis Expectation Maximization (PCA-EM) algorithm on the matrix  $\mathbf{Y}$ .

Much like with CPD, we alternate between an E- and M-step. The E-step is given by:

$$\mathbf{X} = \mathbf{Y}\mathbf{L}^T(\mathbf{L}\mathbf{L}^T)^{-1} \quad (9)$$

whereas the M-step is defined by:

$$\mathbf{L} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y}. \quad (10)$$

$\mathbf{L}_{p \times q}$  is the resulting matrix of principle components. So, for a new normalized set of observations  $\mathbf{Y}_o$ , the latent variables can be found by postmultiplying  $\mathbf{Y}_o$  by  $\mathbf{L}$ . In this manner, a deformation field is now described by only  $q$  latent parameters. Similarly, any point  $\mathbf{x}$  in the latent space can be converted into a deformation field transformation by first postmultiplying  $\mathbf{x}$  by  $\mathbf{L}^T$  and by converting the result into a  $\mathbf{W}_{M \times D}$  matrix after the respective denormalization. Thus, moving through the  $q$ -dimensional space linearly interpolates between the deformation fields.

#### D. Grasping Knowledge Aggregation

We aggregate grasping knowledge from different instances into the canonical model in two steps: first, by generating the grasping motion in the observed space and, second, by transforming its grasping descriptor into the canonical space.

A grasping motion is represented as a sequence of parametrized primitives each of them defined by a control pose expressed in the same coordinate system of the shape of the object. The generation of grasping motions can be performed manually for each instance in the training set, which favors accuracy over time and wear off of the system (on real robotic platforms). This imposes however a limit on the number of samples of the training dataset mostly because of time constraints. In order to overcome this limit, we adopt a constrained sample-based motion generation approach.

A sampled motion is created by generating constrained random 6D poses around the control poses of the canonical grasping motion as depicted in Figure 4. Each component of the translation is sampled from a normal distribution. For the rotation, a quaternion is build out of three uniformed points following the approach described in [27]. These orientations are filtered by specific functional constraints of each category, in the case of drills, for example, rotations that occlude or impede the use of the trigger are discarded. If the sampled grasping motion leads to collisions with other objects in the environment including the robotic arm, the motion is discarded as well. Finally, the sampled motion is executed and evaluated. If the object is functionally grasped successfully, the grasping control poses are transformed into the canonical space.

Finding the transformation from the observed space into the canonical space is equal to finding the inverse transformation of Eq. (2) or equivalently to finding the inverse transformation of Eq. (4). However, the inverse function  $v^{-1}$  is not directly available. It can nonetheless be estimated for a point  $\mathbf{o}$  in the space of the observed shape using a set of points  $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_M)^T$  in the canonical space which deform close to  $\mathbf{o}$  by the equation:

$$v^{-1}(\mathbf{o}) = -\frac{\sum_{i=1}^M G(\mathbf{o}, \mathbf{z}_i + v(\mathbf{z}_i))v(\mathbf{z}_i)}{\sum_{i=1}^M G(\mathbf{o}, \mathbf{z}_i + v(\mathbf{z}_i))}. \quad (11)$$




Figure 4: Sampled-based grasping motion generation. 6D constrained random poses are sampled around control poses of the canonical grasping motion.



For transforming the orientation, we apply Eq. (11) to the rotational vector base of each pose and orthonormalize it.

For each instance in our training dataset, we have so far a latent vector  $\mathbf{x}_i$  that represents the shape deformations from the canonical instance to the observed instance and a grasping descriptor  $\varsigma_i$  transformed into the canonical space. We set the latent vector  $\mathbf{x}_i$  as a feature vector and the grasping descriptor  $\varsigma_i$  as the corresponding target output and train a linear regression model. In other words, grasping knowledge is aggregated in the canonical model by serving as a training label of a regression model (Fig. 5). Algorithm 1 summarizes the training phase (Figure 2).

### E. Shape Inference

A shape transformation is specified by the  $q$  parameters of the latent vector  $\mathbf{x}_i$  plus additional seven parameters of a rigid transformation  $\theta_i$ . The rigid transformation is meant to account for minor misalignments between the observed shape and the canonical shape at the global level.

We concurrently optimize for the latent parameters and the rigid transformation using gradient descent. As CPD and ICP, our method requires an initial coarse alignment of the observed shape because of the expected local minima. We want to find an aligned dense deformation field which when applied to the canonical shape  $\mathbf{C}$  minimizes the distance to corresponding points in the observed shape  $\mathbf{O}$ . Specifically, we want to minimize the energy function:

$$E(\mathbf{x}, \theta) = - \sum_{m=1}^M \log \sum_{n=1}^N \exp \frac{1}{2\sigma^2} \|\mathbf{O}_n - \Theta(\mathcal{T}_m(\mathbf{C}_m, \mathbf{W}_m(\mathbf{x})), \theta)\|^2 \quad (12)$$

where the function  $\Theta$  applies the rigid transformation given parameters  $\theta$ .

When a minimum is found, we can transform any point or set of points into the observed space by applying the deformation

field using Eq. (4) and Eq. (2) and then applying the rigid transformation  $\Theta$ . Moreover, CPD provides a dense deformation field, allowing us to find deformation vectors for novel points, even those added after the field is created.

### F. Transferring Grasping Skills

The transfer of grasping skills for novel instances is performed as follows. A latent vector  $\mathbf{x}$  describing the shape deformation of the object from the canonical instance is calculated as explained in Section III-E. This vector constitutes a test sample of the linear regression, whose inference is a grasping descriptor  ${}^c\varsigma$ . Then,  ${}^c\varsigma$  is transformed into the observed space. This transformation is performed in two steps. First, the control poses of the grasping motion are warped using Eq. (2) replacing  $\mathbf{S}^{[t]}$  by the translational part and the rotational vector base of the control poses. Because the warping process can violate the orthogonality of the orientation, we orthonormalize the warped orientation. Second, we apply the rigid transformation  $\Theta$  defined by the parameters  $\theta$ .

The resulting transformed control poses  ${}^o\varsigma$  are expressed in the frame of the object. Thus, for executing the motion each of the poses has to be adapted relative to the pose of the observed object by premultiplying the control poses by the pose of the object w.r.t. the base of the manipulator. Algorithm 2 summarizes the inference of grasping skills.

## IV. SETUP AND EVALUATION

In this section, we evaluate only the grasping skill transfer because the latent space non-rigid registration method was already evaluated in [1]. We tested our method on two categories: *Drill* and *Spray Bottle*, containing 13 and 17 instances respectively. We obtained the object models from two online CAD databases: GrabCad<sup>1</sup> and 3DWarehouse<sup>2</sup>. The CAD models were converted into meshes in order to generate the input point clouds for our method. They were obtained by




Figure 5: Grasping knowledge aggregation. Grasping descriptors of observed instances are transformed and aggregated in the canonical model by Eq. (11)

#### Algorithm 1 Training phase

**Input:** A set of training shapes in their canonical pose with corresponding grasping descriptors  ${}^o\varsigma$ .

- 1: Select a canonical shape  $\mathbf{C}$  via heuristic or pick the one with the lower reconstruction energy.
- 2: Estimate the deformation fields between the canonical shape and the other training examples using CPD.
- 3: Concatenate the resulting set of  $\mathbf{W}$  matrices from the deformation fields into a design matrix  $\mathbf{Y}$ .
- 4: Perform PCA-EM on the design matrix  $\mathbf{Y}$  to compute the latent space of deformation fields  $\mathbf{x}$ .
- 5: Transform the grasping descriptors  ${}^o\varsigma$  into the canonical space  ${}^c\varsigma$ .
- 6: Train the Linear Regressor  $\mathcal{R} : \mathbf{x} \rightarrow {}^c\varsigma$ .

**Output:** A canonical shape  $\mathbf{C}$ , a latent space of deformation fields  $\mathbf{L}$  and a trained model for inferring grasping descriptors  $\mathcal{R}$ .

<sup>1</sup><https://grabcad.com/library>

<sup>2</sup><https://3dwarehouse.sketchup.com/>

**Algorithm 2** Grasping Skills Inference

**Input:** Transformation model ( $\mathbf{C}$ ,  $\mathbf{L}$ ), trained regressor  $\mathcal{R}$  and observed shape  $\mathbf{O}$

- 1: Use gradient descent to estimate the parameters of the underlying transformation ( $\mathbf{x}$  and  $\boldsymbol{\theta}$ ) until the termination criteria is met. To calculate the value of the energy function, in each iteration:
  - Using the current values of  $\mathbf{x}$  and  $\boldsymbol{\theta}$ :
  - 1) Create vector  $\hat{\mathbf{Y}}$  and convert it into matrix  $\mathbf{W}$ .
  - 2) Use Eq. (4) and Eq. (2) to deform  $\mathbf{C}$ .
  - 3) Apply the rigid transformation  $\Theta$  to the deformed  $\mathbf{C}$ .
- 2: Use the resulting  $\mathbf{x}$  to infer a grasping descriptor  $^c\zeta$  inferred by  $\mathcal{R}$ .
- 3: Transform the grasping descriptor into the observed space.

**Output:** Grasping descriptor in observed space  $^c\zeta$ .

ray-casting from several viewpoints on a tessellated sphere and down-sampling with a voxel grid filter.

We use the five-fingered Schunk hand with a total of 9 fully actuated Degrees of Freedom (DoF) and 11 mimic joints. The experiments were carried out in the Gazebo simulation environment. The collision model of the finger links were modeled by capsules using an automatic ROS optimal capsule generator based on the RoboOptim library [28] as shown in Fig. 6. The inertia tensors of the graspable objects were approximated using Meshlab. For building the shape latent spaces, we parametrized CPD with  $\beta=1$ ,  $\lambda=3$  and  $\sigma^2=0.01$ . The number of latent variables was set to capture at least 95% of the variance of each class. The grasping motions for each object in the training set were sampled as described in Section III-D with a maximum distance of 0.04 m and a maximum angular deviation of 0.2.

For each category, we select the canonical model manually. We use cross validation leaving two samples out. We trained six *drill* and seven *spray bottle* grasping transfer models. Because our method is able to infer category-alike geometries, we also evaluated our method with partially-observed point clouds. For this, we generate a single view of the test objects of each cross validation model. In total, we evaluated the method on 12 fully observed and 12 partially observed *drills* and 14 fully observed and 14 partially observed *spray bottles*. For each instance, one



Figure 6: Visual and collision model of the robotic hand. At rightmost both models are displayed simultaneously to show the goodness of the capsule approximation.

TABLE I  
RATIO OF SUCCESSFULLY TRANSFERRED GRASPS.

	Grasp	Drill Func. Grasp	Spray Bottle Grasp	Func. Grasp
Fully observed	7/12	4/12	8/14	3/14
Partially observed	6/12	3/12	9/14	6/14

simulation trial was performed because the execution of the generated motion is fully deterministic in simulation.

From the 52 instances to be grasped 30 were successfully grasped; that yields a success rate of 57.7%. Note, however, that a successfully grasped instance in our approach considers the entire motion, not only the last grasp configuration.

Regarding functional grasps, i.e., the index finger is able to trigger the tools, 16 instances were successfully grasped which results in a 32% success rate. The results are presented in Table I. Compared to the results presented in [22], although the success rate of our method is lower, our method is able to handle partially-occluded objects and an inference takes in average 7 s compared to the 12.6 min which is only suitable for offline applications. Figure 7 shows for each category two different—a fully observed and a partially-observed—samples that were successfully grasped.

Our method was also tested in real-robot experiments. We created only one latent transformation model for the *drill* category using all the 12 available meshes plus the canonical model. The observed object was inferred from one single view captured by the Kinect v2 sensor [29]. The tests were carried out on two different platforms: a UR10 arm and the CENTAURO robot. The hand was controlled by a PID position-current cascade controller, such that the joint position controller defines the desired joint currents. The saturation values of the current controller together with the PID values of the position controller were set to provide a certain level of compliance which contributed mainly at the last stage of the grasping motion. Using the UR10 robotic arm, our method was able to grasp two different drills twice without any failure. Similarly, with the CENTAURO robot, our approach grasped one instance of a drill twice without any failure (Fig. 8).

A video illustrating our approach is available online<sup>3</sup>.

#### A. Discussion

Real experiments with two different robotic arms demonstrate that our method does not depend on the kinematics of the arm holding the hand. We assume however that the kinematics of the arm is able to reach 6D poses in its workspace. Our method is also agnostic to the robotic hand; a canonical grasping motion that is suitable to the hand is the only requirement for applicability.

Most of the grasping motions that failed exhibited a high deviation with respect to the canonical control poses which indicates a large variance in the learned transfer model. This suggests a need for more sample-efficient inference methods and the need for more training data.

<sup>3</sup>[http://www.ais.uni-bonn.de/videos/RA-L\\_2018\\_Rodriguez](http://www.ais.uni-bonn.de/videos/RA-L_2018_Rodriguez)

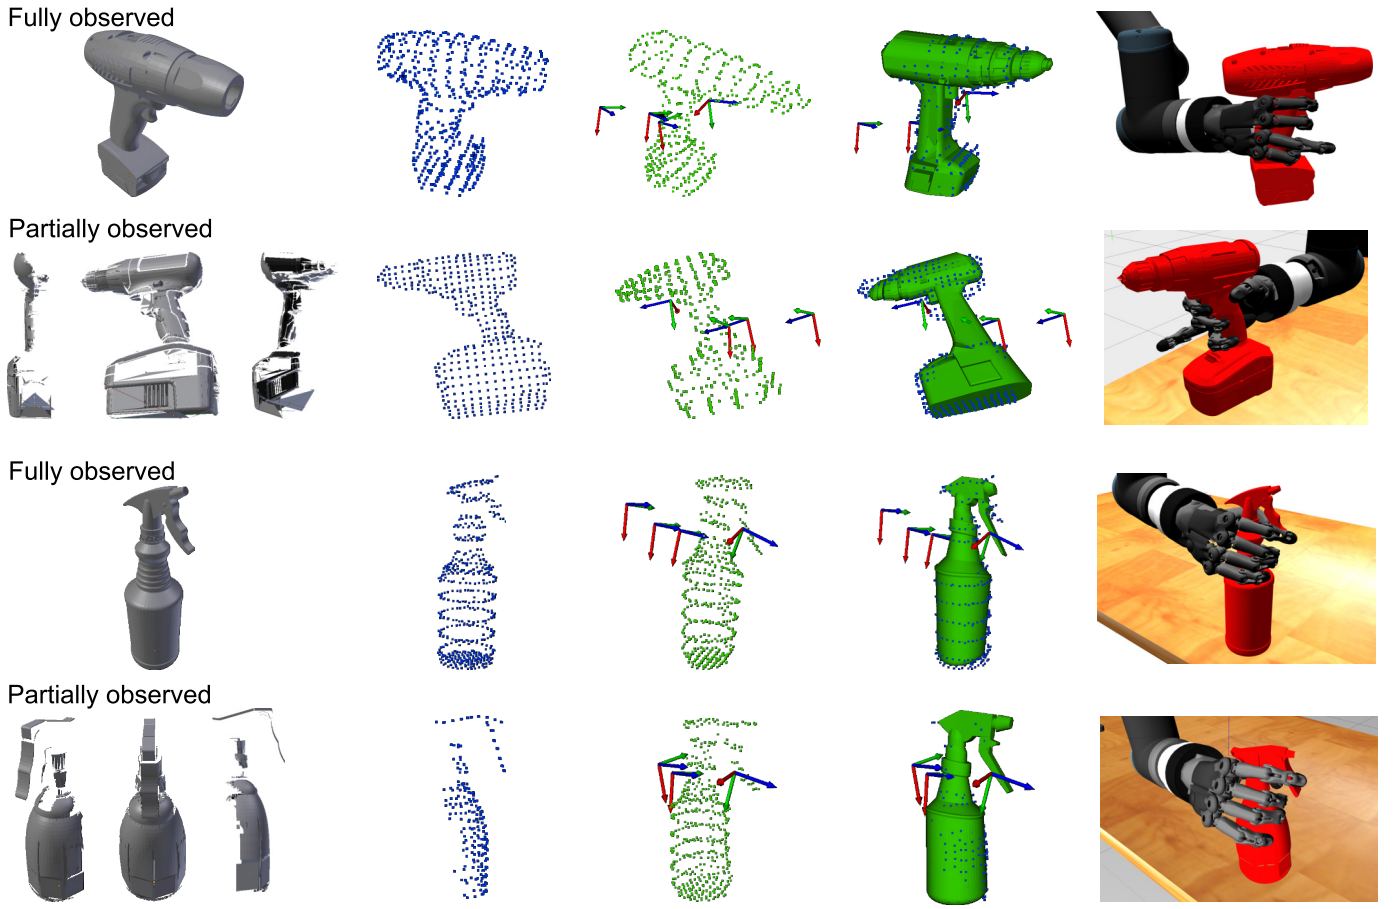


Figure 7: At the leftmost the meshes are shown. For illustration purposes we show additional perspectives of the same single view of the partially observed objects. The respective point clouds are shown in blue. The inferred instances (green point clouds) together with the transformed control points that define the motion are also displayed. In order to observe how good the inference matches the observed points, the mesh of the canonical models is transformed and displayed (green meshes) together with the observed data (blue points). Finally, the resulting grasped object in Gazebo is also depicted at the rightmost.




Figure 8: Experiments performed with the Centauro robot grasping autonomously a novel instance of a drill.



## V. CONCLUSION

In this paper, we proposed a new approach of transferring grasping skills between objects within a category that is based on the knowledge aggregation of different training samples into a canonical model. Thanks to the learned latent shape space, our method is capable of completing missing or occluded object surfaces from partial views. Our method was able to transfer grasping skills with real robotic platforms from experiences collected only in simulation. This demonstrates the feasibility regarding the available sensory data (single-view point clouds) and runtime of our approach.

For future work, we want to consider more complex categories that impose higher variations in the joint configuration of the hand. So, more dimensionality reduction will be expected. As we realized the reduced number of training samples limits the presented approach, we start looking into automatic generation of plausible meshes from the canonical model. We also want to explore variants of the CPD algorithm in order to speed our current implementation. Finally, we would like also to exploit additional sensory modalities such as joint currents and force-torque sensors.

## REFERENCES

- [1] D. Rodriguez, C. Cogswell, S. Koo, and S. Behnke, "Transferring grasping skills to novel instances by latent space non-rigid registration", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [2] B. Lévy, S. Petitjean, N. Ray, and J. Maillot, "Least squares conformal maps for automatic texture atlas generation", in *ACM Transactions on Graphics (TOG)*, vol. 21, 2002, pp. 362–371.
- [3] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios, "Dense non-rigid surface registration using high-order graph matching", in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pp. 382–389.
- [4] V. G. Kim, Y. Lipman, and T. Funkhouser, "Blended intrinsic maps", in *ACM Transactions on Graphics (TOG)*, vol. 30, 2011, p. 79.
- [5] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "Efficient computation of isometry-invariant distances between surfaces", *SIAM Journal on Scientific Computing*, vol. 28, no. 5, pp. 1812–1836, 2006.
- [6] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel, "Isometric registration of ambiguous and partial data", in *Computer Vision and Pattern Recognition (CVPR), 2009 IEEE Conference on*, pp. 1185–1192.
- [7] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas, "One point isometric matching with the heat kernel", in *Computer Graphics Forum*, Wiley Online Library, vol. 29, 2010, pp. 1555–1564.
- [8] B. Allen, B. Curless, and Z. Popović, "The space of human body shapes: reconstruction and parameterization from range scans", in *ACM Transactions on Graphics (TOG)*, vol. 22, 2003, pp. 587–594.
- [9] B. J. Brown and S. Rusinkiewicz, "Global non-rigid alignment of 3-d scans", in *ACM Transactions on Graphics (TOG)*, vol. 26, 2007, p. 21.
- [10] D. Hahnel, S. Thrun, and W. Burgard, "An extension of the ICP algorithm for modeling nonrigid objects with mobile robots", in *18th International Joint Conference on Artificial Intelligence (IJCAI)*, 2003, pp. 915–920.
- [11] A. Myronenko and X. Song, "Point set registration: Coherent point drift", *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 32, no. 12, pp. 2262–2275, 2010.
- [12] H. Li, B. Adams, L. J. Guibas, and M. Pauly, "Robust single-view geometry and motion reconstruction", in *ACM Transactions on Graphics (TOG)*, vol. 28, 2009, p. 175.
- [13] J. Süßmuth, M. Winter, and G. Greiner, "Reconstructing animated meshes from time-varying point clouds", in *Computer Graphics Forum*, Wiley Online Library, vol. 27, 2008, pp. 1469–1476.
- [14] M. Wand, B. Adams, M. Ovsjanikov, A. Berner, M. Bokeloh, P. Jenke, L. Guibas, H.-P. Seidel, and A. Schilling, "Efficient reconstruction of nonrigid shape and motion from real-time 3D scanner data", *ACM Transactions on Graphics (TOG)*, vol. 28, no. 2, p. 15, 2009.
- [15] R. A. Newcombe, D. Fox, and S. M. Seitz, "Dyadicfusion: Reconstruction and tracking of non-rigid scenes in real-time", in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 343–352.
- [16] M. Zollhöfer, M. Nießner, S. Izadi, C. Rehmann, C. Zach, M. Fisher, C. Wu, A. Fitzgibbon, C. Loop, C. Theobalt, et al., "Real-time non-rigid reconstruction using an RGB-D camera", *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, p. 156, 2014.
- [17] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel, "A statistical model of human pose and body shape", in *Computer Graphics Forum*, Wiley Online Library, vol. 28, 2009, pp. 337–346.
- [18] O. Burghard, A. Berner, M. Wand, N. Mitra, H.-P. Seidel, and R. Klein, "Compact part-based shape spaces for dense correspondences", *ArXiv preprint arXiv:1311.7535*, 2013.
- [19] F. Engelmann, J. Stückler, and B. Leibe, "Joint object pose estimation and shape reconstruction in urban street scenes using 3D shape priors", in *German Conference on Pattern Recognition (GCPR)*, 2016, pp. 219–230.
- [20] N. Vahrenkamp, L. Westkamp, N. Yamanobe, E. E. Aksoy, and T. Asfour, "Part-based grasp planning for familiar objects", in *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2016, pp. 919–925.
- [21] F. Ficuciello, D. Zaccara, and B. Siciliano, "Synergy-based policy improvement with path integrals for anthropomorphic hands", in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1940–1945.
- [22] T. Stouraitis, U. Hillenbrand, and M. A. Roa, "Functional power grasps transferred through warping and replanning", in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4933–4940.
- [23] U. Hillenbrand and M. A. Roa, "Transferring functional grasps through contact warping and local replanning", in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 2963–2970.
- [24] H. B. Amor, O. Kroemer, U. Hillenbrand, G. Neumann, and J. Peters, "Generalization of human grasping for multi-fingered robot hands", in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 2043–2050.
- [25] J. Stueckler, R. Steffens, D. Holz, and S. Behnke, "Real-time 3D perception and efficient grasp planning for everyday manipulation tasks.", in *European Conference on Mobile Robots (ECMR)*, 2011, pp. 177–182.
- [26] A. L. Yuille and N. M. Grzywacz, "The motion coherence theory", in *Computer Vision, 2nd International Conference on (ICCV)*, IEEE, 1988, pp. 344–353.
- [27] K. Shoemake, "Uniform random rotations", in *Graphics Gems*, Morgan Kaufmann, 1992, pp. 124–132.
- [28] A. E. Khoury, F. Lamiraux, and M. Taix, "Optimal motion planning for humanoid robots", in *2013 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3136–3141.
- [29] P. Fankhauser, M. Bloesch, D. Rodriguez, R. Kaestner, M. Hutter, and R. Siegwart, "Kinect v2 for mobile robot navigation: Evaluation and modeling", in *International Conference on Advanced Robotics (ICAR)*, 2015, pp. 388–394.