

TransGrasp: Grasp Pose Estimation of a Category of Objects by Transferring Grasps from Only One Labeled Instance

Hongtao Wen^{*}, Jianhang Yan^{*}, Wanli Peng[†], and Yi Sun

Dalian University of Technology, China
{wenht,yjh97,1136558142}@mail.dlut.edu.cn, lslwf@dlut.edu.cn

Abstract. Grasp pose estimation is an important issue for robots to interact with the real world. However, most of existing methods require exact 3D object models available beforehand or a large amount of grasp annotations for training. To avoid these problems, we propose TransGrasp, a category-level grasp pose estimation method that predicts grasp poses of a category of objects by labeling only one object instance. Specifically, we perform grasp pose transfer across a category of objects based on their shape correspondences and propose a grasp pose refinement module to further fine-tune grasp pose of grippers so as to ensure successful grasps. Experiments demonstrate the effectiveness of our method on achieving high-quality grasps with the transferred grasp poses. Our code is available at <https://github.com/yanjh97/TransGrasp>.

Keywords: Grasp Pose Estimation · Dense Shape Correspondence · Grasp Transfer · Grasp Refinement

1 Introduction

Grasp pose estimation refers to the problem of finding a grasp configuration for the grasping task given an object. It’s one of the most important issues for robots to interact with the real world and has great potential to replace human hands to execute various tasks, such as helping elderly or disabled people carry out everyday tasks. However, predicting a stable grasp of an object is a challenging problem as it requires modeling physical contact between a robot hand and various objects in real scenarios. Currently, most existing works rely on grasp pose annotations on each object or exact 3D object models available beforehand. However, annotating 6D grasp pose is not easy as the rotation, translation and width of the gripper have to be adjusted to ensure a stable grasp. In this paper, we deal with the problems from the perspective of computer vision to develop a category-level grasp pose prediction model. Our approach aims to estimate the grasp poses of all instances by labeling the grasp poses on only one instance of the category and transferring the grasp to other instances, as shown in Fig. 1.

^{*} H. Wen and J. Yan—Equal contributions.

[†] Corresponding author.



Fig. 1. Our method for Grasp Pose Estimation with **only one labeled instance**. We transfer grasps from one labeled source object to target objects, using the learned categorical dense correspondence among a category of objects

As we have known, the study on category-level grasp pose estimation has not been fully explored so far.

Generally, grasp pose estimation can be divided into model-based methods [15,12,33,31,26,10] and data-driven (model-free) methods [18,13,16,17,22,8,2,30]. Both methods have their own limitations. Model-based methods generate grasp poses of an object regarding the criteria such as force closure. It requires precise geometric and complete 3D model of an object which is not always available in real scenarios, while data-driven methods estimate grasp poses by experience that is gathered during grasp execution. It requires dense grasp pose annotations which are extremely time-consuming and is also uncertain whether the experience can be generalized to novel objects. In this paper, we propose a novel grasp pose estimation approach, named TransGrasp, to solve the above problems.

Our TransGrasp is based on the fact that a category of object instances have consistent topologies and grasping affordances, such as mugs with grasping handles for drinking. This motivates us to align grasp contact points within a category of objects by first building their shape correspondences and then transfer labeled grasps on one object to the others, finally these initial grasps on all objects are fine-tuned continuously. In this work, we first learn the categorical shape latent space and a template that captures common structure of the category from the given object instances. The variance of each object instance is represented by the difference of its shape relative to the template, and dense correspondence across instance shapes can be established by deforming their surfaces to the template. Then we design a novel model to transfer grasp points on one object instance across a category of objects based on their shape deformation correspondence. To avoid extreme distortions of grasp points during correspondence deformation, the grasp pose refinement module is introduced to further alter the gripper, allowing the grasp points on the surface to form a stable grasp. Note that the correction on gripper is crucial to guarantee a successful grasp. Finally, grasp points of gripper on any object surface can be aligned across an entire category of objects.

Compared with previous works, the proposed model predicts grasp poses of an entire category of objects with grasp pose annotations of only one object instance, which greatly reduces the amount of 6DoF grasp pose annotations of objects in data-driven methods. Furthermore, category-level grasp pose estimation model learns a categorical shape latent space shared by all instances, which enables shape reconstruction and grasp pose estimation within a category. We

additionally estimate rotation R , translation t , and scale s given object observable point cloud as input to transform grasp poses from Object Coordinates to Camera Coordinates and subsequently conduct simulation and real robot experiments. The experiment results verify that our estimated grasp poses can align well across the instances of a category.

The contributions of this work are as follows:

- We propose a category-level grasp pose estimation model to predict grasp poses of a category of objects by only labeling one object instance, which does not heavily relying on a large amount of grasp pose annotations and 3D model on every object.
- We exploit consistent structure within a category of objects to perform grasp pose transfer across a category of objects. In addition, the grasp pose refinement module further corrects grasp poses to ensure successful grasps.
- We set up simulated and real robot system and conduct sufficient experiments that verify the effectiveness of the proposed method on achieving high-quality grasps with the transferred grasp poses.

2 Related Work

In this section, we briefly review the literatures on 6DoF grasp pose estimation, which can be approximately divided into model-based methods and data-driven (model-free) methods.

Model-based methods: A common pipeline for model-based methods is to generate the grasp pose on the known object models using the criteria such as force closure [15,12], and then transform annotated grasps from Object Coordinates to Camera Coordinates with the estimated object 6D pose. Zeng *et al.* [33] segment objects with fully convolutional network from multiple views and fit pre-scanned 3D models with ICP [1] to obtain object pose. Others apply learning-based and instance-level object pose estimation methods such as [10,26,31] for grasp transformation. However, these model-based methods require exact 3D models beforehand, which cannot be used in general settings because the target objects may be different from the existing 3D models. To estimate the 6D pose of unseen objects, researches for category-level object pose estimation [27,4,25,5] have drawn more and more attention, but they have hardly been applied to the grasp estimation. In this work, we explore the challenging problem for category-level grasp pose estimation by transferring grasps from one labeled instance to the others.

Model-free methods: To estimate the grasp pose on unseen objects, some model-free methods [18,13,16,17,22,8,2,30] that analyze the 3D geometry of the object surface directly using deep neural network are proposed. GPD [18] and PointNetGPD [13] generate a large set of grasp candidates firstly and then build a CNN or PointNet [20] to identify whether each candidate is a grasp or not. Following this pipeline, 6-DOF GraspNet [16] builds a variational auto-encoder (VAE) for efficiently candidates sampling and introduces a grasp evaluator network for grasp refinement. CollisionNet [17] builds a collision evaluation exten-

sion based on [16] to check the cluttered environment for collisions. Since analyzing every candidate costs significantly, some end-to-end methods [22,8,2,30] are proposed. Among them, S⁴G [22] builds a single-shot grasp proposal network based on PointNet++ [21] to predict 6DoF grasp pose and quality score of each point directly. Fang *et al.* [8] build a large scale grasp pose dataset called GraspNet-1Billion and propose an 3D grasp detection network that learns approaching direction, operation parameters and grasp robustness end-to-end. However, all these methods need a sufficient amount of dense grasp pose annotations for all the objects in the scene for training, which is time-consuming and laborious. To avoid this problem, our method estimates grasp poses of a category of objects by labeling only one object instance, which greatly reduces the cost of grasp pose annotations.

Closely related to our work, some methods transfer grasps to other objects by building dense correspondence. DGCM-Net [19] transfers grasps by predicting view-dependent normalized object coordinate (VD-NOC) values between pairs of depth images. CaTGrasp [29] maps the input point cloud to a Non-Uniform Normalized Object Coordinate Space (NUNOCS) where the spatial correspondence is built. DON [9] learns pixel-wise descriptors in a self-supervised manner to transfer grasp points between 2D images. NDF [24] builds 3D dense correspondence using Neural Descriptor Field, through which the grasp poses are transferred directly. Different from these methods, we construct dense surface correspondence between the complete shape of objects to transfer the grasp points on the surface to obtain coarse grasps, and further refine these grasps using a refinement module to ensure reliable and successful grasps for the gripper.

3 Method

Our ultimate goal is to estimate grasp poses of a category of objects by labeling only one object instance, rather than heavily relying on a large amount of grasp pose annotations on every object. Our TransGrasp establishes point-to-point dense correspondence among intra-class shapes (Sec. 3.1) to transfer grasp points on one source object instance to other target object instances by aligning grasp poses on the object surface (Sec. 3.2). Furthermore, to avoid extreme distortions during grasp transfer, we introduce a grasp pose refinement module to further adjust the pose of gripper, encouraging the grasps to satisfy antipodal principle (Sec. 3.3). Our overall architecture is illustrated in Fig. 2. In the remainder of this section, we will discuss each component in detail.

3.1 Learning Categorical Dense Correspondence

Objects can be classified according to shape, function, *etc.* Objects classified by shape usually have consistent topologies and grasping affordances and they are likely to be successfully grasped in a similar way. Based on this prior, we can build the categorical dense correspondence across object instances so that grasp poses could be transferred from one object to others. The categorical correspondence

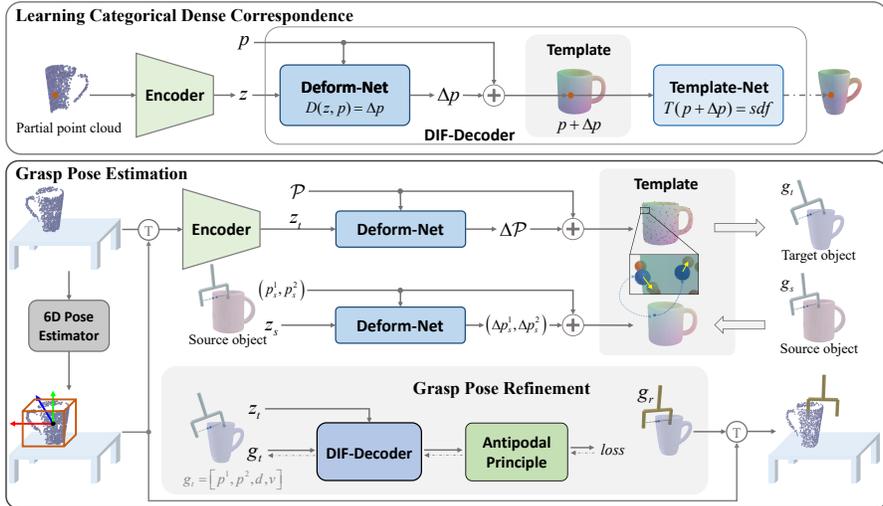


Fig. 2. Overview of the proposed TransGrasp. Top: The network composed of Shape Encoder and DIF-Decoder is trained to learn categorical dense correspondence by deforming object points to the template. Bottom: During inference process, point clouds are first transformed from Camera Coordinates to Object Coordinates and then serve as the input of our trained network, generating the object instance’s deformation to template. Meanwhile, the source model with grasp annotation g_s is fed into Deform-Net to obtain its deformation. The dense correspondence established by their common template subsequently guides the transfer from g_s to g_t by aligning grasp points on object surface. The transferred grasp pose g_t is then refined to g_r by the grasp pose refinement module. Finally, g_r is transformed from Object Coordinates to Camera Coordinates so as to conduct a grasp

across object instances needs to handle the intra-class shape variation and most recent works model the variation explicitly from a pre-learned template [34,25,6]. The variance of each object instance is represented by the difference of its shape relative to the template, and dense correspondence among instance shapes can be established through the template. In this paper, we build an encoder-decoder framework to reconstruct object model given partial point cloud and provide dense shape correspondence between the reconstructed model and the template. The encoder-decoder framework is illustrated at the top of Fig. 2.

Firstly, a latent shape space representing various 3D shapes and a template of a class for shape correspondence are learned by a decoder-only network. Here we adopt Deformed Implicit Field [6] decoder (DIF-Decoder), with a continuous random variable z as input to adapt to the intra-class shape variation. The DIF-Decoder can be divided into two parts, a Template-Net (T) and a Deform-Net (D). The Template-Net can map a 3D point p to a scalar SDF value sdf :

$$T(p) = sdf, p \in \mathbb{R}^3, sdf \in \mathbb{R}. \quad (1)$$

The network weights of T are shared across the whole class, therefore it is enforced to learn common patterns within the class. For each object, the Deform-

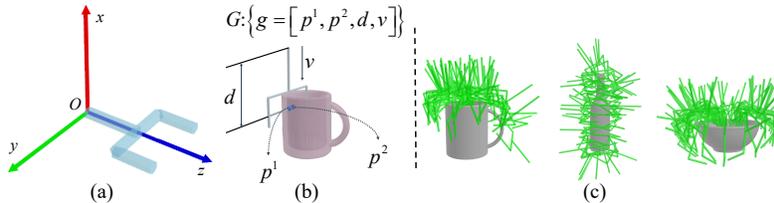


Fig. 3. Grasp pose representation and grasp annotations on source object. (a) two-finger gripper’s mesh model; (b) grasp representation where (p^1, p^2) are the grasp points on the object surface projected from gripper, with d and v denoting approaching depth and approaching vector, respectively; (c) The visualization of 50 grasp annotations on the source object for each category

Net learns a deformation offset to the template $D(z, p) = \Delta p$ for each query point p with the latent code z . And then, the SDF value of point p can be obtained via Template-Net $T(p + \Delta p) = sdf$. This process establishes dense correspondence between an object instance and the template. In other words, the implicit template is “warped” according to the latent codes to model different SDFs. Therefore, dense correspondence among instance shapes can be established through the template. This provides a good reference for grasp transfer across object instances. The DIF-Decoder is trained separately for each category using the CAD models selected from ShapeNetCore [3]. After training, the Template-Net captures the common properties of one certain category and Deform-Net establishes the dense correspondence between instances and the template. Meanwhile, each object instance is mapped to a low-dimensional shape vector z .

However, input data are usually partial point clouds in practice. To complete the shape information, a Shape Encoder learns to map each input to a latent shape code z so that given the shape code a complete shape can be reconstructed using the priors learned by the DIF-Decoder. The Shape Encoder is trained by taking a partial point cloud rendered from its CAD model as input and shape vector z of each object instance as output.

3.2 Grasp Pose Estimation

After learning the categorical dense correspondence among different objects, we then focus on estimating grasp poses for every object instance by grasp pose transfer from only one labeled object. This transfer process is illustrated in the middle of Fig. 2. The source object is the only one object instance with grasp pose annotations, and its grasp poses on the object surface will be transferred to other target object instances by its established correspondence to other ones. We first introduce a novel grasp pose representation that establishes the connection between grasp pose of the gripper and object surface points, and then implement the grasp pose transfer from the source object surface to the target object surface. **Grasp Pose Representation:** To easily implement grasp transfer using dense correspondence between objects, we seek a novel suitable grasp representation for the gripper which follows the two-finger gripper’s mesh model shown in Fig. 3(a).

Unlike previous works [16,17,32] that define the grasp pose as $(R, t, width)$ of the gripper, we introduce a novel grasp representation which connects grasp pose with the object surface points as illustrated in Fig. 3(b). Formally, we denote the two grasp points on the object surface projected from gripper as p^1, p^2 , the vertical distance from gripper’s origin O to the line linking p^1 and p^2 as approaching depth d , and the approaching vector as v . These four parameters form the grasp pose of the gripper where v determines the approaching direction and d the approaching depth along this direction. The width of gripper is determined by the distance between the two grasp points. In this section of Grasp Pose Estimation, only grasp points (p^1, p^2) are transferred from the labeled source object to other target objects keeping the d and v unchanged, while in the next section of Grasp Pose Refinement (Sec. 3.3) all pose parameters (p^1, p^2, d, v) are simultaneously refined to obtain optimal stable grasp poses.

Labeling One Instance: To realize grasp transfer among the instances of a class, it is essential to provide a known instance with grasp pose annotations as the source object. Therefore, for each category, one of the instances needs to be labeled in advance. For convenience, we randomly select one instance for each category from the existing large-scale grasp dataset ACRONYM [7] with pose annotations. Meanwhile, to ensure their robustness and avoid inaccuracy, all the chosen grasps are tested in the simulation environment IsaacGym [14], and only the grasps that can successfully grasp the corresponding source objects in the simulation platform are regarded as the source for grasp transfer. We have about 1k grasp poses for each categorical source object. Figure 3(c) shows 50 grasp annotations on the source object for each category.

Grasp Transfer: With our grasp pose representation and one labeled source object, the grasp poses could then be transferred to other object instances of the same category during the inference procedure of our approach shown in the center of Fig. 2. As the encoder and decoder have been learned previously in Sec. 3.1, the shape dense correspondence between the source and target object can be established through the template immediately. Given grasp points p_s^1, p_s^2 on the source object, their deformation offsets $\Delta p_s^1, \Delta p_s^2$ to the template are predicted by Deform-Net, then its grasp poses on the template will be obtained by this deformation accordingly. To further transfer the grasp points on the template to any target object, the input partial point cloud of the target object is first transformed from Camera Coordinates to Object Coordinates using off-the-shelf 6D object pose estimator such as [25], then encoded to the latent code z_t from which its complete shape is reconstructed using DIF-Decoder. And then query point set \mathcal{P} on the surface of reconstructed model is deformed to the template by Deform-Net. Finally, the grasp poses on the template transferred from source object are naturally aligned to the target object through their shared template. It is noted that although the deformed source and target object shape have similar shapes on the template, the grasp points on the source object cannot be accurately matched to surface points on the target. Here we apply the nearest point matching to transfer the grasp points on the source (blue points) to the target (gold points), as shown in the enlarged detail in Fig. 2.

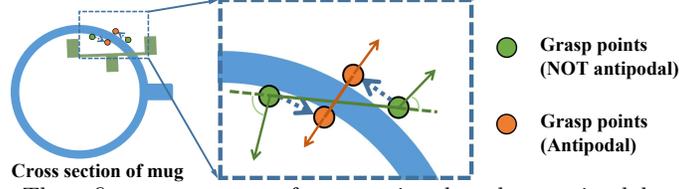


Fig. 4. The refinement process of grasp points based on antipodal principle

3.3 Grasp Pose Refinement

Although point-to-point dense correspondence can realize the grasp transfer from a source object with grasp pose annotations to any target object, when handling some outlier objects with strange shapes, poor deformation could be estimated by Deform-Net, which in turn leads to the failure of grasp transfer. Therefore, we propose a grasp pose refinement module to improve performance after grasp transfer, as shown at the bottom of Fig. 2. To refine grasp pose, we utilize spatial relationship between gripper and object by calculating the distances from the points on gripper to the object surface. Specifically, the transferred grasp points (p^1, p^2) and the sampled points p^g from the gripper, along with the latent code of target object z_t , are input to the DIF-Decoder to get the SDF values of these points. The grasp pose is iteratively refined following the refinement principles for parallel-jaw gripper, until a stable grasp pose g_r is achieved. The gripper poses before and after adjustment are shown as the silver and gold grippers in Fig. 2 respectively.

We first encourage the two grasp points to satisfy the antipodal principle (a simplified version of force closure for parallel-jaw gripper) on the object surface, as shown in Fig. 4. In mathematics, antipodal points of a sphere are those diametrically opposite to each other. For the two-finger grasping problem, the grasp points are antipodal when their normals are in opposite directions and stay in one line with the grasp points. Let n^1 and n^2 be the normal vectors at p^1 and p^2 respectively, then they need to satisfy:

$$\cos \langle p^2 - p^1, n^1 \rangle = 1, \cos \langle p^1 - p^2, n^2 \rangle = 1, \quad (2)$$

where n^1 and n^2 can be easily computed by taking the derivative of sdf with respect to the grasp point: $\partial sdf / \partial p$, using the back-propagation of DIF-Decoder.

Therefore, we define the antipodal loss that encourages the grasp points to satisfy the antipodal principle:

$$L_{anti} = -\cos \langle p^2 - p^1, n^1 \rangle - \cos \langle p^1 - p^2, n^2 \rangle. \quad (3)$$

In addition, we constrain the SDF value of the grasp points (p^1, p^2) to zero, making sure that the grasp points are on the object surface using the touch loss:

$$L_{touch} = \|sdf^1\|_2 + \|sdf^2\|_2, \quad (4)$$

where sdf^1 and sdf^2 are SDF values of the transferred two grasp points (p^1, p^2) .

Moreover, to penalize the interpenetration between the gripper and object, we build a collision avoidance loss to constrain the gripper outside the object, that is, the SDF values of the gripper points are greater than 0 via:

$$L_{collision} = -\frac{1}{N} \sum_i^N \min(0, sdf_i^g), \quad (5)$$

where N is the number of the sampled points p^g of the gripper at its grasp pose, and sdf^g are the SDF values of the these sampled points. Based on such a loss function, the sampled points of the gripper inside the object, of which the SDF values are less than 0, are pushed out of the object surface, so as to avoid the collision between the gripper and the target. Note that this loss function can adjust all the items (p^1, p^2, d, v) of the grasp pose.

Besides, we apply regularization loss to ensure that the grasp pose does not change too much, preventing excessive deviation from stable grasp poses:

$$L_{reg} = \|\Delta p^1\|_2 + \|\Delta p^2\|_2 + \|\Delta d\|_2 + \|\Delta v\|_2, \quad (6)$$

where $\Delta*$ denotes the grasp pose offset during the refinement process.

Among all the losses above, L_{anti} and L_{touch} fine-tune the grasp poses by directly changing the projected grasp points (p^1, p^2) while the effect of $L_{collision}$ and L_{reg} is performed by changing all the items of the grasp pose (p^1, p^2, d, v) . Our final refinement loss is defined as:

$$L_r = \lambda_1 L_{anti} + \lambda_2 L_{collision} + \lambda_3 L_{touch} + \lambda_4 L_{reg}, \quad (7)$$

where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are the trade-off parameters to balance each loss. We set $\lambda_1, \lambda_2, \lambda_3$ and λ_4 as 100, 10, 20 and 200 respectively in practice. During refinement, the weights of the DIF-Decoder and the latent shape code z_t are fixed and L_r is back-propagated through the DIF-Decoder to the grasp pose g_t , so that g_t can be refined to a more stable grasp using Adam algorithm [11].

4 Experiments

In this section, we first introduce the implementation details, along with the dataset and evaluation metrics we adopt. Then we conduct experiments to verify the effectiveness of each part of TransGrasp in Sec. 4.1 and make comparisons with other methods in Sec. 4.2. Finally, real robot experiments are conducted to explore the practical application of our proposed method in Sec. 4.3.

Implementation Details: DIF-Decoder is trained following [34]. For PointNet-like [20] shape encoder, we use the Adam optimizer [11] with an initial learning rate of 0.0001 for training which takes 30 epochs for each category. For the trade-off between the transfer success rate and optimization speed, we set the refine iteration number as 10 in practice.

Dataset: To evaluate our proposed TransGrasp, we select three household categories of objects including 100 mugs, 218 bottles and 107 bowls from ShapeNet-Core [3]. Meanwhile, considering the various shapes of objects within a class, we

introduce data augmentation which enables the learned categorical shape latent space to cover possible object instances of a class. For instance, mugs of different shapes and sizes are generated from the existing mug models by randomly enlarging or shrinking their top or bottom. Please refer to the supplementary for augmentation examples. The training/test distribution of augmented data is 765/135 for mug, 940/150 for bottle and 801/162 for bowl, respectively. We use only one object instance in each class of objects as the source object which are labeled hundreds of possible grasp poses (802 for mug, 1000 for both bottle and bowl, respectively). From this labeled source object, its grasp poses will be transferred to other target objects of the same class. To train Shape Encoder and 6D Object Pose Estimator, We render partial point clouds from 100 random views for each model in training set using Pytorch3D [23].

Evaluation Metrics: In this section, we evaluate the effectiveness of grasp pose estimation for the target objects of each category. As there are no ground truth grasp pose annotations on the target objects for comparison, we utilize IsaacGym [14] to build a parallel-jaw simulation platform where the success rate of grasp transfer is replaced by the success rate of grasp. Our simulation platform is shown in Fig. 6(a). All simulations are done using a Franka Panda manipulator. Different from other grasp generation methods that are interested in the score of one best grasp among all predicted grasps, we focus on every transferred grasp pose and aim to evaluate the effect of our grasp pose transfer. Therefore, for each instance, we count the ratio of successful grasps among all transferred grasps as transfer success rate. And then the transfer success rate for one category is calculated as the average of all its instances’ ones:

$$s_{trans} = \frac{1}{M} \sum_i^M \frac{N_i^{suc}}{N_i^{all}}, \quad (8)$$

where M denotes the number of object instances while N^{suc} and N^{all} denote the number of successful grasps and total grasps respectively. Specifically, the platform consists of free-floating grippers and objects with gravity, and the grippers are placed according to their grasp poses. A grasp is considered as successful one if the object is held for more than 15s.

4.1 Ablation Study

To verify the effectiveness of each part of TransGrasp and avoid other influencing factors, we first conduct the experiments using ground-truth 6D poses and scales of objects. Specifically, the following four experiments are conducted.

Direct Mapping: The grasp points are directly obtained from the source labeled model and then scaled according to the size of target model while both depth d and approaching vector v remain unchanged. **Direct Mapping w/ Refine:** In addition to *Direct Mapping*, it contains our grasp refinement module that fine-tunes grasp poses by changing grasp points (p^1, p^2) , depth d and approaching vector v . **Grasp Transfer:** The grasp points of target model are

Table 1. Performance of grasp transfer for each category

Method	Transfer Success Rate (%)			
	Mug	Bottle	Bowl	Average
Direct Mapping	60.68	79.81	33.27	57.92
Direct Mapping w/ Refine	59.26	79.86	34.83	57.98
Grasp Transfer	82.67	86.93	46.24	71.95
Grasp Transfer w/ Refine	87.77	87.24	61.52	78.84

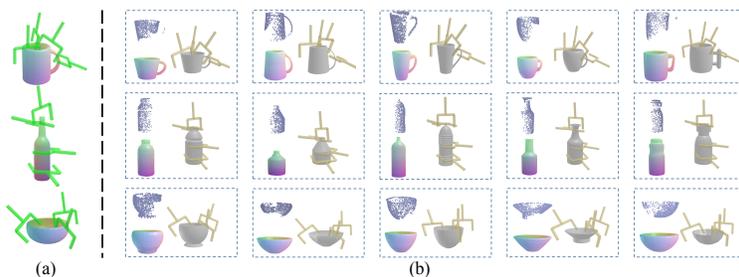


Fig. 5. Qualitative results of three categories. (a) shows source labeled models and some grasp pose annotations. (b) illustrates the transferred grasp poses. For every instance represented inside a dashed box, we show its input partial point cloud (upper left), reconstructed model (bottom left) and transferred grasp poses (right)

transferred from the source labeled model by our method and also scaled according to the size of target model while both depth d and approaching vector v remain unchanged. **Grasp Transfer w/ Refine:** With both grasp transfer and refinement module, it’s the complete version of our proposed method. The quantitative results are shown in Table 1.

Since grasp points after direct mapping may not be on the surface of target object, which provides a poor initialization for refinement, refinement after direct mapping barely has improvement (57.92% to 57.98% in average). Compared to *Direct Mapping*, *Grasp Transfer* greatly increases grasp success rate with 14.03% for the average success rate. This is because through deformation, the transferred grasp points can adapt to object instances with different shapes, which largely avoids the collision of the gripper with the object. It benefits from the categorical shape space learned by TransGrasp which establishes dense correspondences. Besides, *Grasp Transfer w/ Refine* with the proposed grasp pose refinement module further increases the grasp transfer success rates with 6.89% for the average success rate, indicating that our grasp pose refinement module effectively optimizes the transferred grasp poses and improves the stability of grasping. It is worth noting that each method has a relatively low grasp success rate for the bowl. The first reason is that about 28% of the bowl models are single-layer meshes which are not suitable for grasping. The second reason is that the grasp poses of the bowl are always surrounding the rim of bowl, far away from the center of mass of the bowl, which makes the external force insufficient to produce friction that can resist the gravity of the bowl.

Table 2. Quantitative comparison of grasp success rate and inference time

Method	Grasp Success Rate (%)				Inference Time (s) (NVIDIA 1080 GPU)
	Mug	Bottle	Bowl	Average	
GPD [18]	34.81	74.00	44.44	51.08	1.14
6-DOF GraspNet [16]	76.30	78.00	75.93	76.74	1.82
TransGrasp (Ours)	86.67	88.00	72.84	82.50	0.53

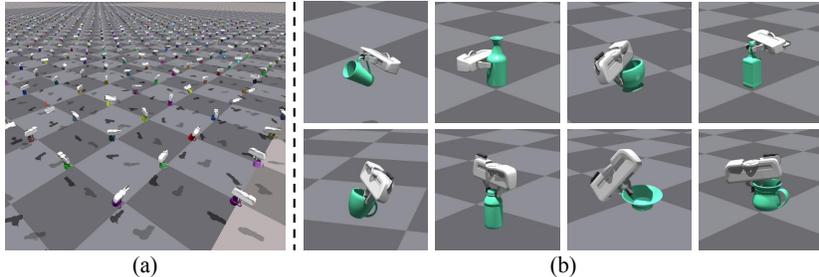
**Fig. 6.** Simulation platform and some grasp results

Figure 5 shows the qualitative results of the transferred grasps after refinement on the test set of each category. We only show a few grasps for clarity. It can be observed that our method can efficiently estimate high-quality grasps even if the input data is incomplete observed point cloud of unseen instance. Figure 6(b) shows the grasp results on the simulation platform [14]. It can be seen that the grippers can grasp objects stably using our transferred grasp poses, even for those objects with slender handle of the mug and thin edge of the bowl.

4.2 Comparisons with Other Methods in Simulation

The above ablation experiments verify the effectiveness of the proposed TransGrasp, we continue to compare our method with others in this section. Among the grasping methods, GPD [18] and 6-DOF GraspNet [16] have similar grippers and object samples as ours, and the two methods are also representative grasping methods at present. Therefore, we make comparison with them. The major difference from the above two methods is that TransGrasp only labels one instance in grasp pose estimation of an entire category of objects, while GPD and 6-DOF GraspNet require grasp pose annotations for every object.

In this experiment, we evaluate grasp success rate on test set of three categories: *mug*, *bottle* and *bowl*. We employ [25] as the 6D Object Pose Estimator to estimate 6D pose of the input point cloud of a target object, as shown in the left of Fig. 2, by which the grasp pose is transformed from Object Coordinates to Camera Coordinates for the next step. We choose the grasp pose with minimal refinement loss as the final one. For GPD and 6-DOF GraspNet, their final grasp poses are chosen according to the grasp scores computed by their own evaluation networks. To be fair, all the methods generate 100 grasp samples for selection.

Table 3. Quantitative comparison of grasp success rate in real robot experiments

Method	Mug	Bottle	Bowl	Average
GPD [18]	7/25(28%)	19/25(76%)	11/25(44%)	37/75(49.3%)
6-DOF GraspNet[16]	16/25(64%)	17/25(68%)	19/25(76%)	52/75(69.3%)
TransGrasp (Ours)	19/25(76%)	21/25(84%)	19/25(76%)	59/75(78.7%)

**Fig. 7.** Real-world experiments setting and some practical results

Table 2 shows the quantitative results of the three methods. Compared with GPD and 6-DOF GraspNet, TransGrasp obtains the highest success rate for the mug and bottle category (88.67% and 88.00%), and achieves comparable performance with 6-DOF GraspNet for the bowl category although our method only labels one instance of an entire category of objects. One dominant factor is that our method reconstructs the complete shapes of objects, rather than directly estimate grasp poses from partial input point clouds like GPD and 6-DOF GraspNet, which introduces more feasible grasps. And benefiting from the reconstructed complete shape, our refinement module can avoid the interpenetration between gripper and unseen parts of object. Furthermore, even our method contains the refinement module, it takes the least inference time with 0.53s among these methods, which satisfies the requirement of grasping household objects.

4.3 Real Robot Experiments

To verify the practical application of our method, we further conduct grasping experiments on a real robot. All experiments are done using a JACO manipulator with an Intel RealSense D435 mounted on its gripper. We choose 5 objects per category with weights from 135.6g to 436.2g for grasping. Figure 7(a) shows the hardware setup and object test set. We make comparison with GPD and 6-DOF GraspNet in our experiments. During the grasping process, we first classify and segment each instance in the scene using off-the-shelf instance segmentation method [28] and then estimate 6D object pose of each instance using [25]. We transfer the grasps with the models trained on our synthetic dataset and test in real scenario directly for our method. For GPD and 6-DOF GraspNet, we use the trained models given by authors to generate grasp poses. After grasps are generated or transferred, we choose the grasp using the grasp selection strategy

consistent with the simulation experiments in Sec. 4.2. If the gripper can lift the object, the grasp is marked as success. We run 5 trials per object and totally 25 trials for each category. As shown in Table 3, our method outperforms GPD and 6-DOF GraspNet on success rate for mug and bottle, and achieves the same performance with 6-DOF GraspNet for bowl. Figure 7(b) shows some of our grasping results in the real-world setting. It can be observed that the grippers can successfully grasp objects with different shapes in real scenes, which indicates that our method can be applied to practical applications. More experiment results can be seen in the supplementary material.

5 Limitations and Future Work

Although the experiments have demonstrated the effectiveness of our method, there is still a limitation to be addressed. As shown in Fig. 8, the grasps are difficult to be transferred across objects with large shape variances. When the grasps on a cylindrical bottle are transferred to a flat bottle or a square bottle, some of the transferred grasps (in red) penetrate the object. We can further improve the transfer success rate by classifying these objects into fine-grained subclass in the future.

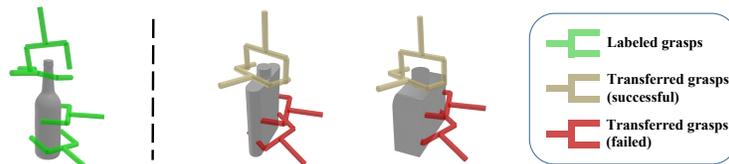


Fig. 8. Grasp Transfer with large shape variances for the bottle category

6 Conclusion

In this work, we present a category-level grasp pose estimation method to predict grasp poses of a category of objects by labeling only one object instance, which does not rely on a large amount of grasp annotations or exact 3D object models available beforehand. We firstly establish dense shape correspondence among an entire category of objects to transfer the grasp poses from the one source object to other instances, and then build a grasp pose refinement module to improve performance after grasp transfer. Sufficient experiments show that our estimated grasp poses can align well across the instances of one category and achieve high-quality grasp performance compared with existing methods.

Acknowledgments. This research was supported by the National Natural Science Foundation of China (No.U1708263 and No.61873046).

References

1. Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: *Sensor fusion IV: control paradigms and data structures*. vol. 1611, pp. 586–606. International Society for Optics and Photonics (1992)
2. Breyer, M., Chung, J.J., Ott, L., Siegwart, R., Nieto, J.: Volumetric grasping network: Real-time 6 dof grasp detection in clutter. *arXiv preprint arXiv:2101.01132* (2021)
3. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: *ShapeNet: An Information-Rich 3D Model Repository*. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago (2015)
4. Chen, D., Li, J., Wang, Z., Xu, K.: Learning canonical shape space for category-level 6d object pose and size estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11973–11982 (2020)
5. Chen, W., Jia, X., Chang, H.J., Duan, J., Shen, L., Leonardis, A.: Fs-net: Fast shape-based network for category-level 6d object pose estimation with decoupled rotation mechanism. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1581–1590 (2021)
6. Deng, Y., Yang, J., Tong, X.: Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10286–10296 (2021)
7. Eppner, C., Mousavian, A., Fox, D.: ACRONYM: A large-scale grasp dataset based on simulation. In: *2021 IEEE Int. Conf. on Robotics and Automation, ICRA* (2020)
8. Fang, H.S., Wang, C., Gou, M., Lu, C.: Graspnet-1billion: A large-scale benchmark for general object grasping. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11444–11453 (2020)
9. Florence, P.R., Manuelli, L., Tedrake, R.: Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. *arXiv preprint arXiv:1806.08756* (2018)
10. He, Y., Sun, W., Huang, H., Liu, J., Fan, H., Sun, J.: Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11632–11641 (2020)
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
12. León, B., Ulbrich, S., Diankov, R., Puche, G., Przybylski, M., Morales, A., Asfour, T., Moiso, S., Bohg, J., Kuffner, J., et al.: Opengrasp: a toolkit for robot grasping simulation. In: *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*. pp. 109–120. Springer (2010)
13. Liang, H., Ma, X., Li, S., Görner, M., Tang, S., Fang, B., Sun, F., Zhang, J.: Pointnetgpd: Detecting grasp configurations from point sets. In: *2019 International Conference on Robotics and Automation (ICRA)*. pp. 3629–3635. IEEE (2019)
14. Makovychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., State, G.: Isaac gym: High performance gpu-based physics simulation for robot learning (2021)
15. Miller, A.T., Allen, P.K.: Graspit! a versatile simulator for robotic grasping. *IEEE Robotics & Automation Magazine* **11**(4), 110–122 (2004)

16. Mousavian, A., Eppner, C., Fox, D.: 6-dof graspnet: Variational grasp generation for object manipulation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2901–2910 (2019)
17. Murali, A., Mousavian, A., Eppner, C., Paxton, C., Fox, D.: 6-dof grasping for target-driven object manipulation in clutter. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 6232–6238. IEEE (2020)
18. ten Pas, A., Gualtieri, M., Saenko, K., Platt, R.: Grasp pose detection in point clouds. *The International Journal of Robotics Research* **36**(13-14), 1455–1473 (2017)
19. Patten, T., Park, K., Vincze, M.: Dgcm-net: dense geometrical correspondence matching network for incremental experience-based robotic grasping. *Frontiers in Robotics and AI* **7** (2020)
20. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
21. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413 (2017)
22. Qin, Y., Chen, R., Zhu, H., Song, M., Xu, J., Su, H.: S4g: Amodal single-view single-shot se (3) grasp detection in cluttered scenes. In: Conference on robot learning. pp. 53–65. PMLR (2020)
23. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv:2007.08501 (2020)
24. Simeonov, A., Du, Y., Tagliasacchi, A., Tenenbaum, J.B., Rodriguez, A., Agrawal, P., Sitzmann, V.: Neural descriptor fields: Se (3)-equivariant object representations for manipulation. arXiv preprint arXiv:2112.05124 (2021)
25. Tian, M., Ang, M.H., Lee, G.H.: Shape prior deformation for categorical 6d object pose and size estimation. In: European Conference on Computer Vision. pp. 530–546. Springer (2020)
26. Wang, C., Xu, D., Zhu, Y., Martín-Martín, R., Lu, C., Fei-Fei, L., Savarese, S.: Densefusion: 6d object pose estimation by iterative dense fusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3343–3352 (2019)
27. Wang, H., Sridhar, S., Huang, J., Valentin, J., Song, S., Guibas, L.J.: Normalized object coordinate space for category-level 6d object pose and size estimation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
28. Wang, X., Kong, T., Shen, C., Jiang, Y., Li, L.: Solo: Segmenting objects by locations. In: European Conference on Computer Vision. pp. 649–665. Springer (2020)
29. Wen, B., Lian, W., Bekris, K., Schaal, S.: Catgrasp: Learning category-level task-relevant grasping in clutter from simulation. arXiv preprint arXiv:2109.09163 (2021)
30. Wu, C., Chen, J., Cao, Q., Zhang, J., Tai, Y., Sun, L., Jia, K.: Grasp proposal networks: An end-to-end solution for visual learning of robotic grasps. arXiv preprint arXiv:2009.12606 (2020)
31. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. arXiv preprint arXiv:1711.00199 (2017)
32. Yang, D., Tosun, T., Eisner, B., Isler, V., Lee, D.: Robotic grasping through combined image-based grasp proposal and 3d reconstruction. In: 2021 IEEE Inter-

- national Conference on Robotics and Automation (ICRA). pp. 6350–6356. IEEE (2021)
33. Zeng, A., Yu, K.T., Song, S., Suo, D., Walker, E., Rodriguez, A., Xiao, J.: Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge. In: 2017 IEEE international conference on robotics and automation (ICRA). pp. 1386–1383. IEEE (2017)
 34. Zheng, Z., Yu, T., Dai, Q., Liu, Y.: Deep implicit templates for 3d shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1429–1439 (2021)