

kPAM: KeyPoint Affordances for Category-Level Robotic Manipulation

Lucas Manuelli*, Wei Gao*, Peter Florence, Russ Tedrake

CSAIL, Massachusetts Institute of Technology,
{manuelli, weigao, peteflo, russt}@mit.edu

**These authors contributed equally to this work.*

Abstract. We would like robots to achieve purposeful manipulation by placing any instance from a category of objects into a desired set of goal states. Existing manipulation pipelines typically specify the desired configuration as a target 6-DOF pose and rely on explicitly estimating the pose of the manipulated objects. However, representing an object with a parameterized transformation defined on a fixed template cannot capture large intra-category shape variation, and specifying a target pose at a category level can be physically infeasible or fail to accomplish the task – e.g. knowing the pose and size of a coffee mug relative to some canonical mug is not sufficient to successfully hang it on a rack by its handle. Hence we propose a novel formulation of category-level manipulation that uses semantic 3D keypoints as the object representation. This keypoint representation enables a simple and interpretable specification of the manipulation target as geometric costs and constraints on the keypoints, which flexibly generalizes existing pose-based manipulation methods. Using this formulation, we factor the manipulation policy into instance segmentation, 3D keypoint detection, optimization-based robot action planning and local dense-geometry-based action execution. This factorization allows us to leverage advances in these sub-problems and combine them into a general and effective perception-to-action manipulation pipeline. Our pipeline is robust to large intra-category shape variation and topology changes as the keypoint representation ignores task-irrelevant geometric details. Extensive hardware experiments demonstrate our method can reliably accomplish tasks with never-before seen objects in a category, such as placing shoes and mugs with significant shape variation into category level target configurations. The video, supplementary material and source code are available on our project page <https://sites.google.com/view/kpam>.

1 Introduction

This paper focuses on pose-aware robotic pick and place at a category level. Contrary to single-instance pick and place, the manipulation policy should generalize to potentially unknown instances in the category with different shape, size, appearance, and topology. These tasks can be easily described using natural language, for example “put the mugs upright on the shelf,” “hang the mugs on the rack by their handle” or “place the shoes onto the shoe rack.” However, converting these intuitive descriptions into concrete robot actions remains a significant challenge. Accomplishing these types of tasks is of significant importance to both industrial applications and interactive assistant robots.

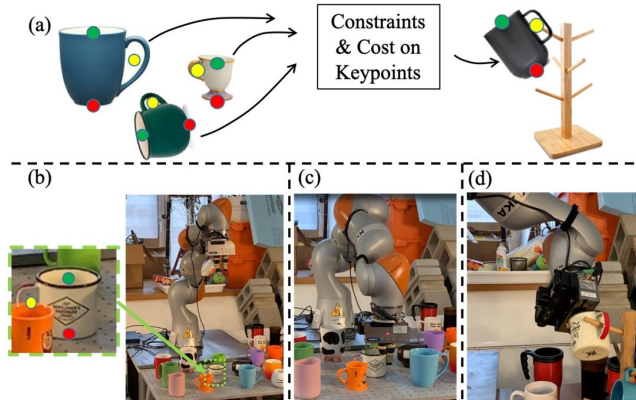


Fig. 1: kPAM is a framework for defining and accomplishing category level manipulation tasks. The key distinction of kPAM is the use of semantic 3D keypoints as the object representation (a), which enables flexible specification of manipulation targets as geometric costs/constraints on keypoints. Using this framework we can handle wide intra-class shape variation (a) and reliably accomplish category-level manipulation tasks such as perceiving (b), grasping (c), and (d) placing any mug on a rack by its handle. A video demo for this task is available on our [project page](#).

While a large body of work addresses robotic picking for arbitrary objects [7, 9, 31], existing methods have not demonstrated pick *and* place with an interpretable and generalizable approach. One way to achieve generalization at the object category level, and perhaps the most straightforward approach is to attempt to extend existing instance-level pick and place pipelines with category-level pose estimators [21, 29]. However, as detailed in Sec. 4, representing an object with a parameterized pose defined on a fixed geometric template, as these works do, may not adequately capture large intra-class shape or topology variations, and can lead to physically infeasible target pose for certain instances in the category. Other recent work has developed dense correspondence visual models, including at a category level, as a general representation for robot manipulation [4], but did not formulate how to specify and solve the task of manipulating objects into specific configurations. As a different route to address category-level pick and place, without an explicit object representation, [6] trains end-to-end policies in simulation to generalize across the object category. It is unclear, however, how to measure the reward function for this type of approach in a fully general way without an object representation that can adequately capture the human’s intention for the task.

Contributions. Our main contribution is a novel formulation of the category-level pick and place task which uses semantic 3D keypoints as the object representation. This keypoint representation enables a simple and interpretable specification of the manipulation target as geometric costs and constraints on the keypoints, which flexibly generalizes existing pose-based manipulation targets. Using this formulation, we contribute a manipulation pipeline that factors the problem into 1) instance segmentation, 2) 3D keypoint detection, 3) optimization-based robot action planning 4) geometric grasping and action execution. This factorization allows us to leverage well-established solutions for these submodules and combine them into a general and effective manipu-

lation pipeline. The keypoint representation ignores task-irrelevant geometric details of the object, making our method robust to large intra-category shape and topology variations. We experimentally demonstrate the use of this keypoint representation with our manipulation pipeline on several category-level pick and place tasks implemented on real hardware. We show that our approach generalizes to novel objects in the category, and that this generalization is accurate enough to accomplish tasks requiring centimeter level precision.

Paper Organization. In Sec. 2 we review related work. Sec. 3 describes our formulation of category-level manipulation tasks. The formulation is introduced in Sec. 3.1 using a concrete example, while Sec. 3.2 describes the general formulation. Sec. 4 compares our formulation with pose-based pick and place pipelines, highlighting the flexibility and generality of our method. Sec. 5 describes the results of hardware experiments on 3 different category-level manipulation tasks, specifically showing generalization to novel object instances. Sec. 6 discusses limitations and future work and Sec. 7 concludes.

2 Related Work

2.1 Object Representations and Perception for Manipulation

There exist a number of object representations, and methods for perceiving these representations, that have been demonstrated to be useful for robot manipulation. For a pick and place task involving a known object the standard solution starts by estimating the object’s 6DOF pose. This allows the robot to then move the object from its estimated pose to the specified target pose. Pose estimation is an extensively studied topic in computer vision and robotics, and existing methods can be generally classified into geometry-based algorithms [5, 19] and learning-based approaches [21, 27, 29]. There exist several datasets [29, 30] annotated with aligned geometric templates, and pose estimators [21, 29] trained on these datasets can produce a category-level pose estimation. Consequently, a straightforward approach to category-level pose-aware manipulation is to combine single object pick and place pipelines with these perception systems. However, pose estimation can be ambiguous under large intra-category shape variations, and moving the object to the specified target pose for the geometric template can lead to incorrect or physically infeasible states for different instances within a category of objects. For example knowing the pose and size of a coffee mug relative to some canonical mug is not sufficient to successfully hang it on a rack by its handle. A more technical discussion is presented in Sec. 4.

Other work has developed and used representations that may be more generalizable than object-specific pose estimation. Recent work has demonstrated dense visual descriptors [22] as a fully self-supervised object representation for manipulation that can generalize at the category level [4]. In comparison with our present work based on 3D keypoints: (i) it is unclear how to extend dense visual descriptors to represent the full object configuration due to self-occlusions which would require N layers of occluded descriptors, (ii) the sparse keypoint representation may in practice be more effective at establishing task-relevant correspondence across significant topology variation, and (iii)

correspondence alone may not fully define a class-general configuration-change manipulation task, but the addition of human-specified geometric costs and constraints on 3D keypoints may. Keypoints have also been used in prior works as components of manipulation pipelines. Several prior works demonstrate the manipulation of deformable objects, and keypoint detection plays a role in their respective perception pipelines. The detected keypoints are typically used as grasp points [13, 24] or building blocks for other shape parameterizations, e.g. the polygons in [17, 18, 28] on which the manipulation policy is defined. These approaches tackled various challenging manipulation tasks such as bed making and towel folding. In contrast, we propose a novel category-level manipulation target specification using costs and constraints defined on 3D keypoints. Additionally the state-machine approach in [13, 24] and manipulation primitives of [18, 28] are specific to cloth and hence our manipulation task is out of scope for these approaches.

2.2 Grasping Algorithms

In recent years there have been significant advances in grasping algorithms that allow robots to reliably pick up a wide range of objects, including potentially unseen objects. Among various approaches for grasping, model-based methods [12, 32] typically rely on a pre-built grasp database of common 3D object models labeled with sets of feasible grasps. During execution, these methods associate the sensor input with an object entry in the database for grasp planning. In contrast, model-free methods [7, 11, 31] directly evaluate the grasp quality from raw sensor inputs. Many of these approaches achieved promising robustness and generality in the Amazon Picking Challenge [23, 31, 32]. Several works also incorporate object semantic information using instance masks [23], or non-rigid registrations [20] to accomplish tasks such as picking up a specific object or transferring a grasp pose to novel instances.

In this work we focus on category-level manipulation tasks which require placing novel instances of a category into desired goal states. Although the ability to reliably grasp an object is an important part of our manipulation pipeline, it doesn't help with the problem of deciding what to do with the object after it has been grasped. Thus the tasks that we consider are out of scope for the aforementioned grasping works.

2.3 End-to-End Reinforcement Learning

There have been impressive contributions [1, 6] in end-to-end reinforcement learning with applications to robotic manipulation. In particular, [6] has demonstrated robotic pick and place across different instances and is the most related to our work. These end-to-end methods encode a manipulation task into a reward function and train the policy using trial-and-error.

However, in order to accomplish the category level pose-aware manipulation task, these end-to-end methods lack a general, flexible, and interpretable way to specify the desired configuration, which is required for the reward function. In [6], the target configuration is implemented specific to the demonstrated task and object category. Extending it to other desired configurations, object categories and tasks is not obvious. In this way, using end-to-end reinforcement learning allows the policy to be learned from experience




Fig. 2: An overview of our manipulation formulation using the “put mugs upright on the table” task as an example: (a) we train a category level keypoint detector that produces two keypoints: $p_{\text{bottom_center}}$ and $p_{\text{top_center}}$. The axis of the mug $v_{\text{mug_axis}}$ is a unit vector from $p_{\text{bottom_center}}$ to $p_{\text{top_center}}$. (b) Given an observed mug, its two keypoints on bottom center and top center are detected. The rigid transform $T_{\text{action}} \in SE(3)$, which represents the robotic pick-and-place action, is solved to move the bottom center of the mug to the target location p_{target} and align the mug axis with the target direction $v_{\text{target_axis}}$.

without worrying about the details of shape variation, but only transfers the burden of shape variation to the choice and implementation of the reward function. We believe that our proposed object representation of 3D keypoints could be used as a solution to this problem.

3 Manipulation Formulation

In this section, we describe our formulation of the category level manipulation problem. Sec. 3.1 describes the approach using a concrete example while Sec. 3.2 presents the general formulation.

3.1 Concrete Motivating Example

Consider the task of “put the mug upright on the table”. We want to come up with a manipulation policy that will accomplish this task for mugs with different size, shape, texture and topology.

To accomplish this task, we pick 2 semantic keypoints on the mugs: the bottom center $p_{\text{bottom_center}}$ and the top center $p_{\text{top_center}}$, as shown in Fig. 2(a). Additionally, we assume we have a keypoint detector, discussed in Section 3.2, that takes as input raw observations (typically RGBD images or point clouds) and outputs the 3D locations of the specified keypoints. Note that there is no restriction that the keypoints be on the object surface, as evidenced by keypoint $p_{\text{top_center}}$ in Fig. 2(a). The 3D keypoints are usually expressed in the camera frame, but they can be transformed to an arbitrary frame using the known camera extrinsics. In the following text, we use $\mathbf{p} = \{p_i\}_{i=1}^N \in \mathbb{R}^{3 \times N}$ to denote the detected keypoint positions in world frame, where p_i is the i^{th} detected keypoint, and N is the total number of keypoints. In this example $N = 2$.

For robotic pick-and-place of mostly rigid objects, we represent the robot action as a rigid transform T_{action} on the manipulated object. Thus, the keypoints associated with

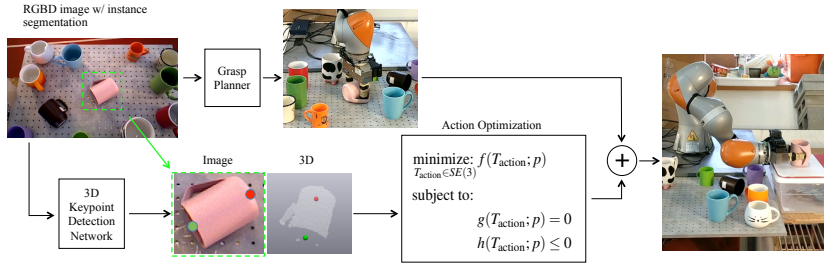


Fig. 3: An overview of the category level pick and place pipeline using our manipulation formulation. Given a RGBD image with instance segmentation, the semantic 3D keypoints of the object in question are detected. We then feed these 3D keypoints into an optimization based planning algorithm to compute the robot pick and place actions, which is represented by a rigid transformation T_{action} . Finally, we use an object-agnostic grasp planner to pick up the object and apply the computed robot action.

the manipulated object will be transformed as $T_{\text{action}}\mathbf{p} \in \mathbb{R}^{3 \times N}$ using the robot action. In practice, this action T_{action} is implemented by first grasping the object using the algorithm detailed in Sec. 3.2 and then planning and executing a trajectory which ends with the object in the desired target location. This trajectory may require approaching the target from a specific direction, for example in the “mug upright on the table” task the mug must approach the table from above.

Given the above analysis, the manipulation task we want to accomplish can be formulated as finding a rigid transformation T_{action} such that

1. The transformed mug bottom center keypoint should be placed at some target location:

$$\|T_{\text{action}}p_{\text{bottom_center}} - p_{\text{target}}\| = 0 \quad (1)$$

2. The transformed direction from the mug bottom center to the top center should be aligned with the upright direction. This is encoded by adding a cost to the objective function

$$\|1 - \langle v_{\text{target_axis}}, \text{rot}(T_{\text{action}})v_{\text{mug_axis}} \rangle\|^2 \quad (2)$$

where $\text{rot}(T)$ is the rotational component of the rigid transformation T , the target orientation $v_{\text{target_axis}} = [0, 0, 1]^T$, and

$$v_{\text{mug_axis}} = \frac{p_{\text{top_center}} - p_{\text{bottom_center}}}{\|p_{\text{top_center}} - p_{\text{bottom_center}}\|} \quad (3)$$

An illustration is presented in Fig. 2 (b). The above problem is an inverse kinematics problem with T_{action} as the decision variable, a constraint given by Eq. (1) and cost given by Eq. (2). This inverse kinematics problem can be reliably solved using off-the-shelf optimization solvers such as [26]. We then pick up the object using robotic grasping algorithms [7, 11] and execute a robot trajectory which applies the manipulation action T_{action} to the grasped object.

3.2 General Formulation

Given an arbitrary category-level manipulation task we propose to solve it in the following manner. First the modeler must specify a set of semantic 3D keypoints $\mathbf{p} = \{p_i\}_{i=1}^N$ for the category, together with a set of geometric costs and constraints on these keypoints which fully specify the task, e.g. Eq. (1)-(2). It is up to the modeler to choose the keypoints, costs and constraints that encode the task. This step can be seen as analogous to choosing costs and constraints in a trajectory optimization or planning method, or specifying a reward function in a reinforcement learning approach. The only restriction on the choice of the 3D keypoints is that they must be well defined for all instances of the category that might be encountered at test time. In particular the 3D keypoints need not lie on the object surface, as demonstrated by $p_{\text{top_center}}$ in Fig. 6. The keypoints can also lie in regions of the object that may be occluded at test time, as exemplified by the $p_{\text{bottom_center}}$ keypoint in Fig. 6 which is subject to self-occlusion when the mug is viewed from the side.

Once we have chosen keypoints together with geometric costs and constraints as the problem specification there exist natural formulations for each remaining piece of the manipulation pipeline. This enables us to factor the manipulation policy into 4 sub-problems: 1) object instance segmentation 2) category level 3D keypoint detection, 3) a kinematic optimization problem to determine the manipulation action T_{action} and 4) grasping the object and executing the desired manipulation action T_{action} . An illustration of our complete manipulation pipeline is shown in Fig. 3. In the following sections, we describe each component of our manipulation pipeline in detail.

Instance Segmentation and Keypoint Detection As discussed in Section 3.1 the kPAM pipeline requires being able to detect category-level 3D keypoints from RGBD images of specific object instances. Here we present a specific approach we used to the keypoint detection problem, but note that any technique that can detect these 3D keypoints could be used instead.

We use the state-of-the-art integral network [25] for 3D keypoint detection. For each keypoint, the network produces a probability heatmap and a depth prediction map as the raw outputs. The 2-D image coordinates and depth value are extracted using the integral operation [25]. The 3-D keypoints are recovered using the calibrated camera intrinsic parameters. These keypoints are then transformed into world frame using the camera extrinsics.

We collect the training data for keypoint detection using a pipeline similar to LabelFusion [14]. Given a scene containing the object of interest we first perform a 3D reconstruction. Then we manually label the keypoints on the 3D reconstruction. We note that this does not require pre-built object meshes. Keypoint locations in image space can be recovered by projecting the 3D keypoint annotations into the camera image using the known camera calibration. Training dataset statistics are provided in Fig. 7(c). In total labeling our 117 training scenes took less than four hours of manual annotation time and resulted in over 100,000 labeled images. Even with this relatively small amount of human labeling time we were able to achieve centimeter accurate keypoint detections, enabling us to accomplish challenging tasks requiring high precision, see Section 5. More details on the keypoint detection network are contained in the supplementary material.

The keypoint detection network [25] requires object instance segmentation as the input, and we integrate Mask R-CNN [8] into our manipulation pipeline to accomplish this step. The training data mentioned above for the keypoint detector [25] can also be used to train the instance segmentation network [8]. Please refer to the supplemental material for more detail.

kPAM Optimization The optimization used to find the desired robot action T_{action}^* can in general be written as

$$\begin{aligned} & \underset{T_{\text{action}} \in SE(3)}{\text{minimize:}} f(T_{\text{action}}; \mathbf{p}) \\ & \text{subject to:} \end{aligned} \tag{4}$$

$$\begin{aligned} g(T_{\text{action}}; \mathbf{p}) &= 0 \\ h(T_{\text{action}}; \mathbf{p}) &\leq 0 \end{aligned}$$

where f is a scalar cost function, g and h are the equality and inequality constraints, respectively. The robot action T_{action} is the decision variable of the optimization problem, and the detected keypoint locations enter the optimization parametrically.

In addition to the constraints used in Sec. 3.1, a wide variety of costs and constraints can be used in the optimization (4). This allows the user to flexibly specify a large variety of manipulation tasks. Below we present several costs/constraints used in our experiments:

1. L2 distance cost between the transformed keypoint with its nominal target location:

$$\|T_{\text{action}} p_i - p_{\text{target},i}\|^2 \tag{5}$$

This is a relaxation of the target position constraint presented in Sec. 3.1

2. Half space constraint on the keypoint:

$$\langle n_{\text{plane}}, T_{\text{action}} p_i \rangle \leq b_{\text{plane}} \tag{6}$$

where $n_{\text{plane}} \in R^3$ and $b_{\text{plane}} \in R$ defines the separating plane of the half space. Using the mug in Sec. 3.1 as an example, this constraint can be used to ensure all the keypoints are above the table to avoid penetration.

3. The point-to-plane distance cost of the keypoint

$$\| \langle n_{\text{plane}}, T_{\text{action}} p_i \rangle - b_{\text{plane}} \|^2 \tag{7}$$

where $n_{\text{plane}} \in R^3$ and $b_{\text{plane}} \in R$ defines the plane that the keypoint p_i should be in contact with. By using this cost with keypoints that should be placed on the contact surface, for instance the $p_{\text{bottom_center}}$ of the mug in Sec. 3.1, the optimization (4) can prevent the object from floating in the air.

4. The robot action T_{action} should be within the robot’s workspace and avoid collisions.

Robot Grasping Robotic grasping algorithms, such as [7, 11], can be used to apply the abstracted robot action $T_{\text{action}} \in SE(3)$ produced by the kPAM optimization (4) to the manipulated object. If the object is rigid and the grasp is tight (no relative motion between the gripper and object), applying a rigid transformation to the robot gripper will




Fig. 4: A pose representation cannot capture large intra-category variations. Here we show different alignment results from a shoe template (blue) to a boot observation (red). (a) and (b) are produced by [5] with variation on the random seed, and the estimated transformation consists of a rigid pose and a global scale. In (c), the estimated transformation is a fully non-rigid deformation field in [19]. In these examples, the shoe template and transformations can not capture the geometry of the boot observation. Additionally, there may exist multiple suboptimal alignments which make the pose estimator ambiguous. The subsequent robotic pick and place action from these estimations are different, despite these alignments being reasonable geometrically.

apply the same transformation to the manipulated object. These grasping algorithms [7, 12] are object-agnostic and can robustly generalize to novel instances within a given category.

For the purposes of this work we developed a grasp planner which uses the detected keypoints, together with local dense geometric information from a pointcloud, to find high quality grasps. This local geometric information is incorporated with an algorithm similar to the baseline method of [31]. In general the keypoints used to specify the manipulation task aren't sufficient to determine a good grasp on the object. Thus incorporating local dense geometric information from a depth image or pointcloud can be advantageous. This geometric information is readily available from the RGBD image used for keypoint detection, and doesn't require object meshes. Our grasp planner leverages the detected keypoints to reduce the search space of grasps, allowing us to focus our search on, for example, the heel of a shoe or the rim of a mug. Once we know which aspect of the local geometry to focus on, a high quality grasp can be found by any variety of geometric or learning-based grasping algorithms [7, 11].

We stress that keypoints are a sparse representation of the object sufficient for describing the manipulation task. However grasping, which depends on the detailed local geometry, can benefit from denser RGBD and pointcloud data. This doesn't detract from keypoints as an object representation for manipulation, but rather shows the benefits of different representations for different components of the manipulation pipeline.

4 Comparison and Discussions

In this section we compare our approach, as outlined in Sec. 3, to existing robotic pick and place methods that use pose as the object representation.

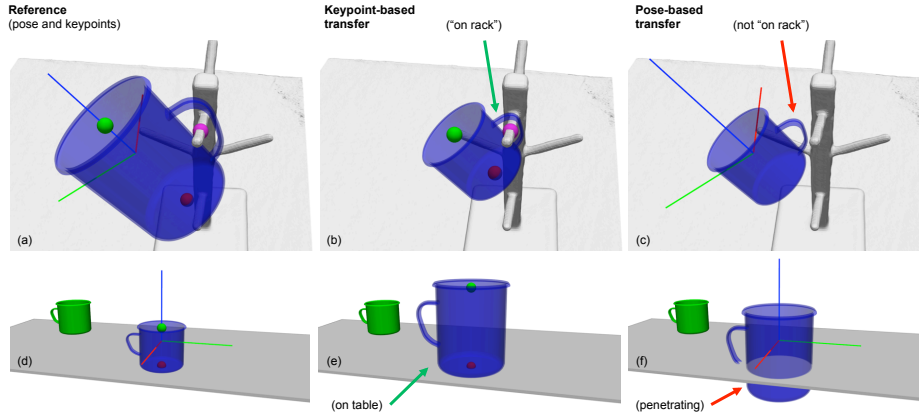


Fig. 5: A comparison of the keypoint based manipulation with pose based manipulation for two different tasks involving mugs. The first row considers the mug on rack task, where a mug must be hung on a rack by its handle. (a) Shows a reference mug in the goal state, (b) and (c) show a scaled down mug instance that could be encountered at test time. (b) uses keypoint based optimization with a constraint on the handle keypoint to find the target state for the mug. The optimized goal state successfully achieves the task of hanging the mug on the rack. In contrast (c) shows the scaled mug instance at the pose defined by (a), which leads to the handle of the mug completely missing the rack, a failure of the task. The second row shows the task of putting a mug on a table. Again (a) shows a reference mug in a goal state, (b) - (c) show a scaled up mug that could be encountered at test time. (b) uses keypoint based optimization with costs/constraints on the bottom and top keypoints to place the mug in a valid goal state. (c) directly uses the pose from (a) on the new mug instance which leads to an invalid goal state where the mug is penetrating the table.

4.1 Keypoint Representation vs Pose Representation

At the foundation of existing pose-estimation methods is the assumption that the geometry of the object can be represented as a parameterized transformation defined on a fixed template. Commonly used parameterized pose families include rigid, affine, articulated or general deformable. For a given observation (typically an RGBD image or pointcloud), these pose estimators produce a parameterized transformation that aligns the geometric template to the observation.

However, the pose representation is not able to capture large intra-category shape variation. An illustration is presented in Fig. 4, where we try to align a shoe template (blue) to a boot observation (red). The alignments in Fig. 4 (a) and (b) are produced by [5] where the estimated transformation consists of a rigid pose and a global scale. Fig. 4 (c) is produced by [19] and the estimated transformation is a fully non-rigid deformation field. In these examples, the shoe template and transformations cannot capture the geometry of the boot observation. Additionally, there may exist multiple suboptimal alignments which make the pose estimator ambiguous, as shown in Fig. 4. Feeding these ambiguous estimations into a pose-based manipulation pipeline will produce different pick and place actions and final configurations of the manipulated object.

In contrast, we use semantic 3D keypoints as a sparse but task-specific object representation for the manipulation task. Many existing works demonstrate accurate 3D keypoint detection that generalizes to novel instances within the category. We leverage these contributions to build a robust and flexible manipulation pipeline.

Conceptually, a pose representation can also be transformed into keypoint representation given keypoint annotations on the template. However, in practice the transformed keypoints can be inaccurate as the template and the pose cannot fully capture the geometry of new instances. Using the shoe keypoint annotation in Fig. 6 as an example, transforming the keypoints p_5 and p_6 to a boot using the shoe to boot alignment in Fig 4 would result in erroneous keypoint detections. A general non-rigid kinematic model (and the associated estimator) that can handle large variations of shape and topology, such as in the example of Fig. 4, remains an open problem. Our method avoids this problem by sidestepping the geometric alignment phase and directly detecting the 3D keypoint locations.

4.2 Keypoint Target vs Pose Target

For existing pose-based pick and place pipelines, the manipulation task is defined as a target pose of the objects. For a given scene where the pose of each object has been estimated, these pipelines grasp the object in question and use the robot to move the objects from their current pose to the target pose.

The proposed method can be regarded as a generalization of the pose-based pick and place algorithms. If we detect 3 or more keypoints and assign their target positions as the manipulation goal, then this is equivalent to pose-based manipulation. In addition, our method can specify more flexible manipulation problems with explicit geometric constraints, such as the bottom of the cup must be on the table and its orientation must be aligned with the upright direction, see Sec. 3.1. The proposed method also naturally generalizes to other objects within the given category, as the keypoint representation ignores many task-irrelevant geometric details.

On the contrary a pose target is object-specific and defining a target pose at the category level can lead to manipulation actions that are physically infeasible. Consider the mug on table task from Section 3.1. Fig. 5(d) shows the target pose for the reference mug model. Directly applying this pose to the scaled mug instance in Fig. 5(f) leads to physically infeasible state where the mug is penetrating the table. In contrast, using the optimization formulation of Section 3 results in the mug resting stably on the table, shown in Fig. 5(e).

In addition to leading to states which are physically infeasible, pose-based targets at a category level can also lead to poses which are physically feasible but fail to accomplish the manipulation task. Figures 5(a) - (c) show the mug on rack task. In this task the goal is to hang a mug on a rack by its handle. Fig. 5(a) shows the reference model in the goal state. Fig. 5(c) shows the result of applying the pose based target to the scaled down mug instance. As can be seen even though the pose unambiguously matches the target pose exactly, this state doesn't accomplish the manipulation task since the mug handle completely misses the rack. Fig. 5(b) shows the result of our kPAM approach. Simply by adding a constraint that handle center keypoint should be on the rack, a valid goal state is returned by the kPAM optimization.

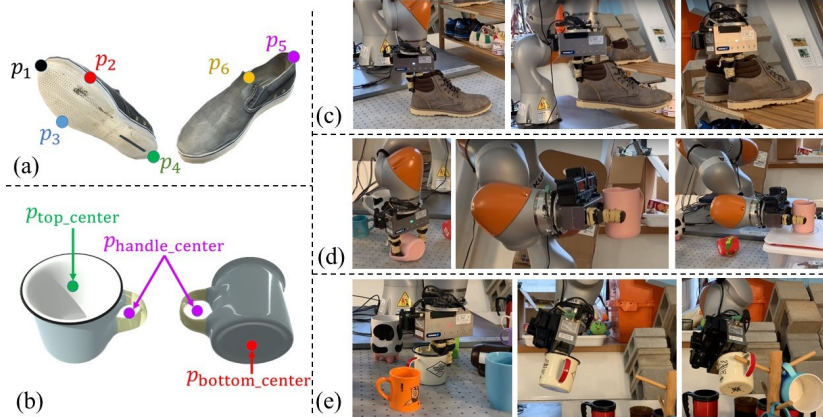


Fig. 6: An overview of our experiments. (a) and (b) are the semantic keypoints we used for the manipulation of shoes and mugs. We use three manipulation tasks to evaluate our pipeline: (c) put shoes on a shelf; (d) put mugs on a mug shelf; (e) hang mugs on a rack by the mug handles. The video of these experiments are available on our [project page](#).

5 Results

In this section, we demonstrate a variety of pose-aware pick and place tasks using our keypoint-based manipulation pipeline. The particular novelty of these demonstrations is that our method is able to handle large intra-category variations without any instance-wise tuning or specification. We utilize a 7-DOF robot arm (Kuka IIWA LBR) mounted with a Schunk WSG 50 parallel jaw gripper. An RGBD sensor (Primesense Carmine 1.09) is also mounted on the end effector. The video demo on our [project page](#) best demonstrates our solution to these tasks. More details about the experimental setup are included in the supplemental material.

5.1 Put shoes on a shoe rack

Task Description Our first manipulation task is to put shoes on a shoe rack, as shown in Fig. 6(c). We use shoes with different appearance and geometry to evaluate the generality and robustness of our manipulation policy. The six keypoints used in this manipulation task are illustrated in Fig 6(a), and the costs and constraints in the optimization (4) are

1. The L2 distance cost (5) between keypoints p_1, p_2, p_3 and p_4 to their nominal target locations.
2. The sole of the shoe should be in contact with the rack surface. In particular, the point-to-plane cost (7) is used to penalize the deviation of keypoints p_2, p_3 and p_4 from the supporting surface.
3. All the keypoints should be above the supporting surface to avoid penetration. A half-space constraint (6) is used to enforce this condition.

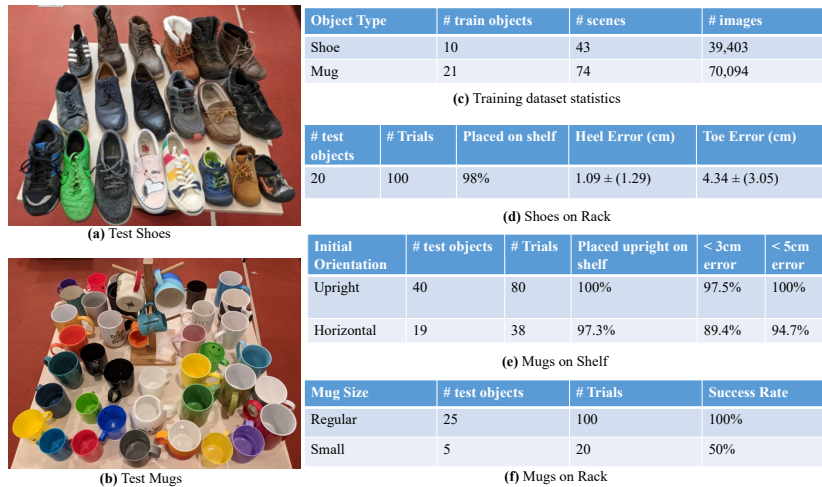


Fig. 7: Quantitative results from the 3 hardware experiments. (a) and (b) show some of the test objects for the experiments. (c) statistics of the training data (d) We report the average heel and toe errors (along the horizontal direction) from their desired locations as well as the standard deviation. (e) The reported errors for the mug on shelf task are the distance from the bottom center keypoint to the target location of that keypoint in the optimization program. (f) reports success rates for the mug on rack task for different sized mugs. Mugs with handles having either height or width less than 2cm are classified as “small” (more details in supplementary material). A trial was deemed successful if the mug ended up hanging on the rack by the mug handle. Videos of the experiments are available on our [project page](#).

For our experiments we place the shoe rack in a known position, but this constraint could be easily relaxed by adding a pose-estimation module for the shoe-rack.

Experimental Results The shoe keypoint detection network was trained on a labeled dataset of 10 shoes, detailed in Figure 7(c). Experiments were conducted with a held out test set of 20 shoes with large variations in shape, size and visual appearance (more details in the video and supplemental material). For each shoe we ran 5 trials of the manipulation task. Each trial consisted of a single shoe being placed on the table in front of the robot. Using the kPAM pipeline the robot would pick up the shoe and place it on a shoe rack. The shoe rack was marked so that the horizontal deviation of the shoe’s toe and heel bottom keypoints (p_1 and p_4 respectively in Fig. 6) from their nominal target locations could be determined. Quantitative results are given in Fig. 7(d). Out of 100 trials only twice did the pipeline fail to place the shoe on the rack. Both failures were due to inaccurate keypoint detections. One led to a failed grasp and another to an incorrect T_{action} . For trials which ended up with the shoe on the rack average errors for the heel and toe keypoint locations are given in Fig. 7(d). During the course of our experiments we noticed that the majority of these errors come from the fact that when the robot grasps the shoe by the heel the closing of the gripper often results in the object shifting from the position it was in when the RGBD image used for keypoint detection was captured. This accounts for the majority of the errors observed in the final heel and

toe keypoint locations. The keypoint detections and resulting T_{action} would have almost always results in heel and toe errors of less than 1 cm if we were able to exactly apply T_{action} to the object. Since our experimental setup relies on a wrist mounted camera we are not able to re-perceive the object after grasping it. We believe that these errors could be further reduced by adding an external camera that would allow us to re-run our keypoint detection after grasping the object to account for any object movement during the grasp. Overall kPAM approach was very successful at the shoes on rack task with a greater than 97% success rate.

5.2 Put mugs upright on a shelf

Task Description We also perform a real-world demonstration of the “put mugs upright on a shelf” task described in Sec. 3.1, as shown in Fig. 6(d). The keypoints used in this task are illustrated in Fig. 6(b). The costs and constraints for this task include the target position constraint (1) and the axis alignment constraint (2). This task is very similar to the mugs task in [6]. For this set of experiments we didn’t place any costs or constraints on the yaw degree of freedom of the mug, but if a specific yaw orientation was desired this could be incorporated by adding an L2 cost (5) between the $p_{\text{handle_center}}$ keypoint with its target location.

Experimental Results The mug keypoint detection network was trained on a dataset of 21 standard sized mugs, detailed in Fig. 7(c). Experiments for the mug on shelf task were conducted using a held out test set of 40 mugs with large variations in shape, size and visual appearance (more details in the video and supplemental material). All mugs could be grasped when in the upright orientation, but due to the limited stroke of our gripper (7.5cm when fully open) only 19 of these mugs could be grasped when lying horizontally. For mugs in that could be grasped horizontally we ran two trials with the mug starting from a horizontal orientation, and two trials with the mug in a vertical orientation. For the remaining mugs we ran two trials for each mug with the mug starting in an upright orientation. Quantitative performance was evaluated by recording whether the mug ended up upright on the shelf, and the distance of the mug’s bottom center keypoint to the target location. Results are shown in Fig. 7(e). Overall our system was very reliable, managing to place the mug on the shelf within 5cm of the target location in all but 2 trials. In one of these failures the mug was placed upside down. In this case the mug was laying horizontally on the table and the RGB image used in keypoint detection (see Fig. 14 in the Appendix) was taken from a side-on profile where the handle is occluded and it is very difficult to distinguish the top from the bottom of the mug. This led our keypoint detector to mix up the top and bottom of the mug, causing it to be placed upside down. The keypoint detection error is understandable in this case since it is very difficult to distinguish the top from the bottom of this mug in the single RGBD image. In addition this particular instance was a small kids sized mug, whereas all the training data for mugs contained only regular sized mugs. See Section 6 for more discussion on this failure.

Overall the accuracy in the mug on shelf task was very high, with 97% of upright trials, and 88% of horizontal trials resulting in bottom keypoint final location errors of less than 3cm. Qualitatively the majority of this error arose from the object moving slightly during the grasping process with the rest attributed to the keypoint detection.

5.3 Hang the mugs on the rack by their handles

Task Description To demonstrate the accuracy and robustness of our method we tasked the robot with autonomously hanging mugs on a rack by their handle. An illustration of this task is provided in Fig. 6 (e). The relatively small mug handles (2-3 centimeters) challenge the accuracy of our manipulation pipeline. The costs and constraints in this task are

1. The target location constraint (1) between $p_{\text{handle_center}}$ to its target location on the rack axis.
2. The keypoint L2 distance cost (5) from $p_{\text{top_center}}$ and $p_{\text{bottom_center}}$ to their nominal target locations.

In order to avoid collisions between the mug and an intermediate goal for the mug was specified. Using the notation of (4) this intermediate goal T_{approach} was gotten by shifting T_{action} away from the rack by 10cm along the direction of the target peg. We then executed the final placement by moving the end effector in a straight line connecting T_{approach} to T_{action} . During these experiments the mug-rack was placed in a fixed known position, but this constraints be easily relaxed by adding a pose-estimation module for the mug-rack.

Experimental Results For the mug on rack experiments we used the same keypoint detection network as for the mug on shelf experiments. Experiments were conducted using a held out test set of 30 mugs with large variation in shape, texture and topology. Of these 5 were very small mugs whose handles had a minimum dimension (either height or width) of less than 2cm (see the supplementary material for more details). We note that the training data did not contain any such “small” mugs. Each trial consisted of placing a single mug on the table in front of the robot. Then the kPAM pipeline was run and a trial was recorded as successful if the mug ended up hanging on the rack by its handle. Five trials were run for each mug and quantitative results are reported in Fig. 7 (e). For regular sized mugs we were able to hang them on the rack with a 100% success rate. The small mugs were much more challenging but we still achieved a 50% success rate. The small mugs have very tiny handles, which stresses the accuracy of the entire system. In particular the total error of the keypoint detection, grasping and execution needed to successfully complete the task for the small mugs was on the order of 1-1.5 cm. Two main factors contributed to failures in the mug on rack task. The first, similar to the case of shoe on rack task, is that during grasping the closing of the gripper often moves the object from the location at which it was perceived. Even a small disturbance (i.e. $< 1\text{cm}$) can lead to a failure in the mug on rack task since the required tolerances are very small. The second contributing factor to failures is inaccurate keypoint detections. Again an inaccurate detection of even 0.5-1cm can be sufficient for the mug handle to miss the rack entirely. As discussed previously, the movement of the object during grasping could be alleviated by the addition of an external camera that would allow us to re-perceive the object after grasping.

6 Limitations and Future Work

Our current data collection pipeline in Sec. 3.2 requires human annotation, although the use of 3D reconstruction somewhat alleviates this manual labor. An interesting direction

for future work is to train our keypoint detector using synthetic data, as demonstrated in [27, 29].

Representing the robot action with a rigid transformation T_{action} is valid for robotic pick-and-place. However, this abstraction does not work for deformable objects or more dexterous manipulation actions on rigid objects, such as the in-hand manipulation in [1]. Combining these learning-based or model-based approaches with the keypoint representation to build a manipulation policy that generalizes to categories of objects would be a promising direction for future work. In addition to the usage in our pipeline, we believe that the keypoint representation can potentially contribute to various learning-based manipulation approaches as 1) a reward function to flexibly specify the manipulation target or 2) an alternative input to the policy/value neural network, which is more robust to shape variation and large deformation than the widely-used pose representation.

7 Conclusion

In this paper we contribute a novel formulation of category-level manipulation which uses semantic 3D keypoints as the object representation. Using keypoints to represent the object enables us to simply and interpretably specify the manipulation target as geometric costs and constraints on the keypoints, which flexibly generalizes existing pose-based manipulation methods. This formulation naturally allows us to factor the manipulation policy into the 3D keypoint detection, optimization-based robot action planning and grasping based action execution. By factoring the problem we are able to leverage advances in these sub-problems and combine them into a general and effective perception-to-action manipulation pipeline. Through extensive hardware experiments, we demonstrate that our pipeline is robust to large intra-category shape variation and can accomplish manipulation tasks requiring centimeter level precision.

Acknowledgements

The authors thank Ethan Weber (instance segmentation training data generation) and Pat Marion (visualization) for their help. This work was supported by: National Science Foundation, Award No. IIS-1427050; Draper Laboratory Incorporated, Award No. SC001-0000001002; Lockheed Martin Corporation, Award No. RPP2016-002; Amazon Research Award.

References

1. Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *arXiv preprint arXiv:1808.00177*, 2018.
2. Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312. ACM, 1996.

3. Ross Finman, Thomas Whelan, Michael Kaess, and John J Leonard. Toward lifelong object segmentation from change detection in dense rgb-d maps. In *Mobile Robots (ECMR), 2013 European Conference on*, pages 178–185. IEEE, 2013.
4. Peter R Florence, Lucas Manuelli, and Russ Tedrake. Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. *Conference on Robot Learning (CoRL)*, 2018.
5. Wei Gao and Russ Tedrake. Filterreg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11095–11104, 2019.
6. Marcus Gualtieri, Andreas ten Pas, and Robert Platt. Pick and place without geometric object models. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7433–7440. IEEE, 2018.
7. Marcus Gualtieri, Andreas Ten Pas, Kate Saenko, and Robert Platt. High precision grasp pose detection in dense clutter. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 598–605. IEEE, 2016.
8. Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
9. Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Scalable deep reinforcement learning for vision-based robotic manipulation. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 651–673. PMLR, 29–31 Oct 2018.
10. Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
11. Jeffrey Mahler, Matthew Matl, Vishal Satish, Michael Danielczuk, Bill DeRose, Stephen McKinley, and Ken Goldberg. Learning ambidextrous robot grasping policies. *Science Robotics*, 4(26):eaau4984, 2019.
12. Jeffrey Mahler, Florian T Pokorny, Brian Hou, Melrose Roderick, Michael Laskey, Mathieu Aubry, Kai Kohlhoff, Torsten Kröger, James Kuffner, and Ken Goldberg. Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1957–1964. IEEE, 2016.
13. Jeremy Maitin-Shepard, Marco Cusumano-Towner, Jinna Lei, and Pieter Abbeel. Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *2010 IEEE International Conference on Robotics and Automation*, pages 2308–2315. IEEE, 2010.
14. P. Marion, P. R. Florence, L. Manuelli, and R. Tedrake. Label fusion: A pipeline for generating ground truth labels for real rgb-d data of cluttered scenes. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3235–3242, May 2018.
15. Pat Marion. Director: A robotics interface and visualization framework, 2015.
16. Francisco Massa and Ross Girshick. maskrcnn-benchmark: Fast, modular reference implementation of Instance Segmentation and Object Detection algorithms in PyTorch. <https://github.com/facebookresearch/maskrcnn-benchmark>, 2018. Accessed: [Insert date here].
17. Stephen Miller, Mario Fritz, Trevor Darrell, and Pieter Abbeel. Parametrized shape models for clothing. In *2011 IEEE International Conference on Robotics and Automation*, pages 4861–4868. IEEE, 2011.
18. Stephen Miller, Jur Van Den Berg, Mario Fritz, Trevor Darrell, Ken Goldberg, and Pieter Abbeel. A geometric approach to robotic laundry folding. *The International Journal of Robotics Research*, 31(2):249–267, 2012.

19. Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.
20. Diego Rodriguez, Corbin Cogswell, Seongyong Koo, and Sven Behnke. Transferring grasping skills to novel instances by latent space non-rigid registration. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.
21. Caner Sahin and Tae-Kyun Kim. Category-level 6d object pose recovery in depth images. In *European Conference on Computer Vision*, pages 665–681. Springer, 2018.
22. Tanner Schmidt, Richard Newcombe, and Dieter Fox. Self-supervised visual descriptor learning for dense correspondence. *IEEE Robotics and Automation Letters*, 2(2):420–427, 2017.
23. Max Schwarz, Christian Lenz, Germán Martín García, Seongyong Koo, Arul Selvam Periyasamy, Michael Schreiber, and Sven Behnke. Fast object learning and dual-arm coordination for cluttered stowing, picking, and packing. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3347–3354. IEEE, 2018.
24. Daniel Seita, Nawid Jamali, Michael Laskey, Ron Berenstein, Ajay Kumar Tanwani, Prakash Baskaran, Soshi Iba, John Canny, and Ken Goldberg. Robot bed-making: Deep transfer learning using depth sensing of deformable fabric. *arXiv preprint arXiv:1809.09810*, 2018.
25. Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018.
26. Russ Tedrake and the Drake Development Team. Drake: A planning, control, and analysis toolbox for nonlinear dynamical systems, 2016.
27. Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *Conference on Robot Learning (CoRL)*, 2018.
28. Jur Van Den Berg, Stephen Miller, Ken Goldberg, and Pieter Abbeel. Gravity-based robotic cloth folding. In *Algorithmic Foundations of Robotics IX*, pages 409–424. Springer, 2010.
29. He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2642–2651, 2019.
30. Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision*, pages 75–82. IEEE, 2014.
31. Andy Zeng, Shuran Song, Kuan-Ting Yu, Elliott Donlon, Francois R Hogan, Maria Bauza, Daolin Ma, Orion Taylor, Melody Liu, Eudald Romo, et al. Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.
32. Andy Zeng, Kuan-Ting Yu, Shuran Song, Daniel Suo, Ed Walker, Alberto Rodriguez, and Jianxiong Xiao. 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)*, pages 1386–1383. IEEE, 2017.