

SKT-Hang: Hanging Everyday Objects via Object-Agnostic Semantic Keypoint Trajectory Generation

Chia-Liang Kuo¹ Yu-Wei Chao² Yi-Ting Chen¹
¹ National Yang Ming Chiao Tung University ² NVIDIA

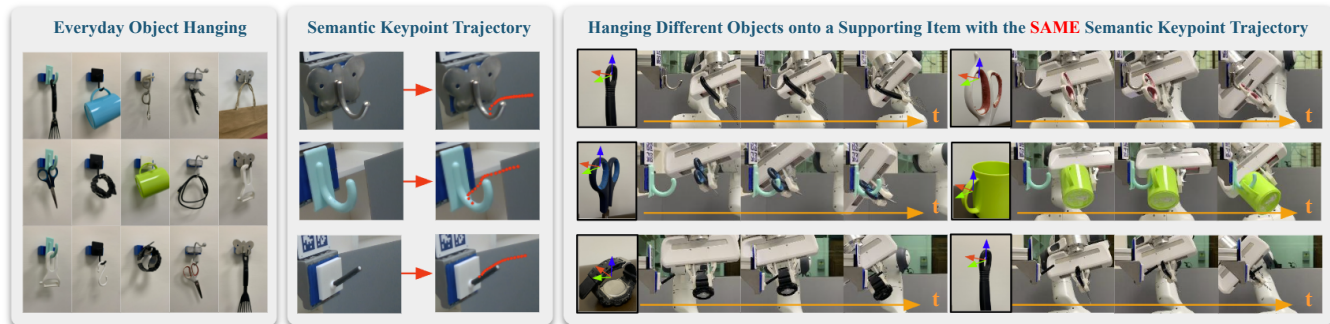


Fig. 1: The task of robot hanging a grasped object onto a supporting item involves a diverse range of grasped objects and supporting items with various shapes and geometric structures. In this work, we introduce *Semantic Keypoint Trajectory* (SKT), an actionable representation that specifies **where** and **how** to hang grasped objects onto a supporting item. Our experiments show that SKT is object-agnostic, enabling hanging various objects onto the same supporting item.

Abstract—We study the problem of hanging a wide range of grasped objects on diverse supporting items. Hanging objects is a ubiquitous task that is encountered in numerous aspects of our everyday lives. However, both the objects and supporting items can exhibit substantial variations in their shapes and structures, bringing two challenging issues: (1) determining the task-relevant geometric structures across different objects and supporting items, and (2) identifying a robust action sequence to accommodate the shape variations of supporting items. To this end, we propose *Semantic Keypoint Trajectory* (SKT), an object-agnostic representation that is highly versatile and applicable to various everyday objects. We also propose *Shape-conditioned Trajectory Deformation Network* (SCTDN), a model that learns to generate SKT by deforming a template trajectory based on the task-relevant geometric structure features of the supporting items. We conduct extensive experiments and demonstrate substantial improvements in our framework over existing robot hanging methods in the success rate and inference time. Finally, our simulation-trained framework shows promising hanging results in the real world. For videos and supplementary materials, please visit our project webpage: <https://hcis-lab.github.io/SKT-Hang/>.

I. INTRODUCTION

Hanging is a common routine in various aspects of our daily lives, taking place in situations like households, logistics, and stationery stores. Equipping robots with the skill to hang objects in unconstrained settings can significantly alleviate labor shortages in these domains. In this paper, we focus on how to enable robots to hang a wide range of objects in hand onto arbitrary supporting items, such as hooks or racks. This is challenging since the robot needs to reason about the hanging structure and determine the corresponding actions while being robust to the diversity in the hanged objects and supporting items.

Existing algorithms [1], [2], [3], [4] apply category-level representations, such as semantic keypoint or dense correspondence, to identify the hanging parts such as mug handles or contact points on supporting items. These category-level representations are subsequently leveraged to define the desired target poses, demonstrating their utility in the task of hanging mugs on racks [1], [3], [4]. However, mere knowledge of hanging poses is insufficient when dealing with diverse objects and supporting items that require fine-grained intermediate actions (see Fig. 1).

To determine action sequences for hanging diverse objects and supporting items, OmniHang [2] proposes a two-stage solution. First, they detect contact points between objects and supporting items to identify suitable hanging poses for objects. Second, they employ a sampling-based motion planner to generate collision-free trajectories for hanging. However, sampling-based motion planners require exhaustive search, resulting in significant computational burdens to determine a feasible trajectory for each object-supporting item pair. Furthermore, predicting collisions under partial observability poses a significant challenge in generating feasible hanging trajectories.

In this work, we present a new representation called *Semantic Keypoint Trajectory* (SKT) and a novel framework for generating SKT to address the challenges faced in the existing methods. SKT is an actionable representation that simultaneously models the hanging part of the supporting item and the hanging movements of the objects' keypoints. Fig. 1 illustrates the concept of SKT. The design of SKT is inspired by the observation that keypoints can be defined across objects such that they have similar movements when hanging onto the same supporting item. In our experiments,

we validate that SKT has the capacity to facilitate the hanging of a wide range of objects onto the same supporting item. It is worth noting that SKT can effectively eliminate the need for replanning for every object-supporting item pair, saving the computational resources for other prominent tasks.

We introduce a novel learning framework for SKT generation: *Shape-Conditioned Trajectory Deformation Network* (SCTDN). SCTDN takes the partial point cloud of a supporting item as input and generates the corresponding SKT by deforming a retrieved template SKT based on the task-relevant geometric structure features. We consider the following two critical design choices. First, we suggest deforming a template trajectory rather than predicting it from scratch, based on the observation that similar SKTs correspond to similar hanging parts of supporting items. We propose an unsupervised clustering algorithm to construct a template trajectory database. Second, we extract features of task-relevant geometric structures when deforming SKTs. This approach is inspired by the trajectory generation with affordance guidance introduced in [5].

We demonstrate the effectiveness of SCTDN with a promising hanging success rate for a wide range of objects and supporting items compared to state-of-the-art algorithms [6], [2], [5]. We also show efficiently generated action sequences compared to [2]. Moreover, given a supporting item, the SKTs predicted by SCTDN can enable hanging various objects compared to a strong baseline [5]. Specifically, we apply all baselines to hang 50 different grasped objects onto 60 unseen supporting items, resulting in a total of 3,000 hanging trials. All methods must determine the corresponding hanging actions for diverse pairs of objects and supporting items under partial observability, making our evaluation setting challenging. We perform thorough ablation studies to justify our design choices, i.e., the selection of deformation over prediction from scratch and the importance of integrating task-relevant geometric structures. Finally, we show promising results of hanging grasped objects onto various supporting items in the real world.

The contributions of this paper are as follows:

- We present *Semantic Keypoint Trajectory*, a novel actionable and object-agnostic representation that simultaneously models the hanging structure of the supporting item and the movements of the objects' keypoints.
- We introduce *Shape-conditioned Trajectory Deformation Network (SCTDN)*, a novel framework for generating semantic keypoint trajectories through reasoning over the geometry of the hanging component and performing trajectory deformation.
- We perform comprehensive experiments and demonstrate significant improvements in both the success rate and inference time, surpassing state-of-the-art object hanging algorithms.
- We present compelling results of hanging grasped objects onto various supporting items in the real world.

II. RELATED WORK

A. Keypoint in Robot Manipulation

Keypoint has gained significant popularity in the realm of generalized robot manipulation [7], [1], [8], [6], [9], [4], [10]. These approaches have yielded promising results in manipulation tasks such as grasping [7], tool use [8], [6], object placement [1], [3], mug hanging [1], [4], and peg-hole insertion [6]. Recent studies have also delved into integrating keypoints with reinforcement learning or imitation learning to determine robot actions for long-horizon tasks such as object pushing [11], [12], pick and place [11], shoe flipping [11], and cable insertion [9]. Inspired by the success of keypoint in modeling task-relevant geometric structures across different objects, we introduce the Semantic Keypoint Trajectory as a novel actionable representation for the hanging task. We demonstrate its effectiveness in offering sequential guidance for the hanging of various objects onto a variety of supporting items.

B. Actionable Representations for Affordance

In recent years, the research community has been exploring actionable representations, such as the combination of affordance and trajectory [13], [5], [14], [15], [16], optical flow [17], and scene flow [18], [19], for various downstream applications, including articulated object manipulation [13], [5], [14], [16], [18], hand motion prediction [15], fabric manipulation [17], and tool use [19]. The goal is to bridge the perception-action gap, enhancing the ability to transfer manipulation skills to unseen configurations. A crucial difference between our proposed representation and existing works lies in our emphasis on the trajectory of the task-relevant geometric structures of an object, rather than the hand pose or gripper pose of a robotic arm.

Among these works, our proposed framework exhibits similarities to the *object-centric actionable visual priors* proposed in VAT-Mart [5]. VAT-Mart jointly learns the interaction point given a partial point cloud and the corresponding trajectory for articulated object manipulation. In our work, we introduce a trajectory deformation-based approach, instead of a trajectory prediction framework. Empirically, we demonstrate this strategy leads to more effective and efficient trajectory generation for diverse hanging tasks.

III. SEMANTIC KEYPOINT TRAJECTORY AND DATASET

A. Semantic Keypoint Trajectory (SKT)

A semantic keypoint trajectory is defined as $\mathcal{T} = \{\xi_0, \xi_1, \dots, \xi_{T-1}\}$, where \mathcal{T} is a sequence of 6D poses defined in a supporting item's coordinate frame, $\xi_i \in \text{SE}(3)$ is the i 'th waypoint, and T is the trajectory length. It is worth noting that, ξ_0 is the end waypoint attached to the supporting item.

Rotation Definition. We systematically define the rotation of a semantic keypoint $\xi_i \in \text{SE}(3)$ as follows (see Fig. 2 (c)): We define the forward direction of an object hung onto the supporting item as the x-axis of a semantic keypoint ξ_i . The y-axis of a semantic keypoint ξ_i is the cross product

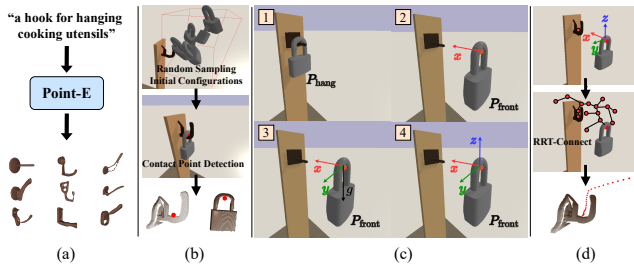


Fig. 2: The overall data collection pipeline. (a) Shape generation via text-to-3D framework [21]. (b) Contact point collection via forward simulation. (c) Rotation extraction via the object’s forward direction and the gravity direction. (d) Semantic keypoint trajectory generation via RRT-Connect [22].

between the x-axis of a semantic keypoint and the unit vector of gravity. The z-axis is subsequently defined as the cross product of the unit vectors of the x-axis and y-axis. There are two notable observations that elucidate the basis for defining rotations in this specific manner. First, in the hanging task, we observed that the direction of object movement is usually perpendicular to the plane formed by the hanging part on an object, such as the “hanging ring” at the tail of a tool or the mug handle. Second, this design choice offers a significant advantage because it allows for the simplification of the trajectory representation during training, retaining only the positional component that eliminates the complexity of regressing a sequence of $SO(3)$ rotations. During inference, we use adjacent 3D waypoints to identify the moving direction (i.e., the x-axis). Empirically, we demonstrate the design choice can favorably hang a wide range of objects to diverse supporting items.

Task execution. First, the relative transform between the robot’s end-effector pose T^{ee} and the grasped object’s $SE(3)$ keypoint T^{kpt} is recorded as $T_{\text{kpt}}^{\text{ee}} = (T^{\text{kpt}})^{-1}T^{\text{ee}}$. We assume $T_{\text{kpt}}^{\text{ee}}$ remains fixed after an object is stably grasped. This assumption is aligned with [6]. To obtain T^{kpt} for each object, one can use off-the-shelf methods such as [6], [20]. The robot first moves the end-effector to the initial waypoint $\xi_{T-1} \times T_{\text{kpt}}^{\text{ee}}$, then performs a hanging manipulation from $\xi_{T-1} \times T_{\text{kpt}}^{\text{ee}}$ to $\xi_0 \times T_{\text{kpt}}^{\text{ee}}$ in reverse order of \mathcal{T} .

B. Data Collection Pipeline

To learn to predict SKT, we propose an automated data collection pipeline within a simulation environment. We make this choice due to the challenges and costs associated with collecting a substantial number of supporting items, along with their corresponding semantic keypoints and SKTs in the real world. The data collection pipeline is illustrated in Fig. 2.

First, we apply Point-E [21], a text-to-3D framework, to collect a diverse set of supporting items (see Fig. 2(a)). To obtain the semantic keypoints, we automatically collect the contact points of the supporting items and the hanging objects as semantic keypoints via forward simulation (see Fig. 2(b)), as demonstrated in [2], [20].

Fig. 2(c) demonstrates how we determine an object’s semantic keypoint rotation. Using a horizontal supporting item during the contact point collection process, we first obtain the position P_{hang} of the object hanging pose. We then

move the object in front of the support item at the position P_{front} and calculate the x-axis (i.e., forward direction) of the keypoint’s rotation matrix as the unit vector of $(P_{\text{hang}} - P_{\text{front}})$. Next, we compute the y-axis of the keypoint’s rotation matrix as the cross-product of the x-axis and the gravity unit vector. Finally, we derive the z-axis as the cross-product of the x-axis and y-axis. Given the availability of complete shapes of supporting items in simulation, we collect SKTs for a supporting item by employing the RRT-Connect path planner [22] with the collision detection API in the Pybullet simulator [23]. It is noteworthy that, we utilize a single object as the reference object to collect the SKTs for all the supporting items, as illustrated in Fig. 2 (d). We select this object as the pivot object because its shape is relatively simple, enabling the RRT-Connect path planner to generate paths more efficiently.

In summary, our dataset comprises 429 supporting items and 50 objects with semantic keypoints. We collected 50 semantic keypoint trajectories for each supporting item and captured 42 point clouds from various camera viewpoints. Our training set consists of 309 supporting items, while we reserved 60 supporting items for validation and another 60 for testing.

IV. PROBLEM FORMULATION

We consider the task of a robot hanging a grasped object onto a supporting item. In this task, our primary objective is to predict the semantic keypoint trajectory, given the partial point cloud of a supporting item $\mathcal{P}^S \in \mathbb{R}^{N \times 3}$ as input. The parameter N is the number of points of the supporting item. This trajectory then serves as a guide for objects to execute a sequence of actions for hanging. The output is a position-only semantic keypoint trajectory $\hat{\mathcal{T}}^S \in \mathbb{R}^{T \times 3}$ for the supporting item, where T is the length of the trajectory. During inference, the position-only trajectory is augmented to a $SE(3)$ trajectory based on the method described in III-A.

V. METHODOLOGY

A. Overview of SCTDN

We propose a Shape-Conditioned Trajectory Deformation Network (SCTDN), a novel learning framework for the hanging task. It leverages the strategy of deforming a template trajectory based on shape-conditioned features to generate the target semantic keypoint trajectory (SKT) for the partial point cloud of the input supporting item.

Fig. 3 offers an overview of the SCTDN architecture, which encompasses three core modules:

Classification Module. This module predicts the shape category c^S of the input point cloud \mathcal{P}^S . It enables the retrieval of the appropriate template SKT $\mathcal{T}^{\text{temp}, c^S}$ from the template trajectory database. This design choice is motivated by the observation that similar SKTs correspond to similar hanging parts of the supporting items. For example, when using curved hooks, the SKTs tend to follow curved paths, while straight hooks typically produce straight SKTs. Inspired by this observation, we implement unsupervised categorization to group supporting items based on their SKTs and create a

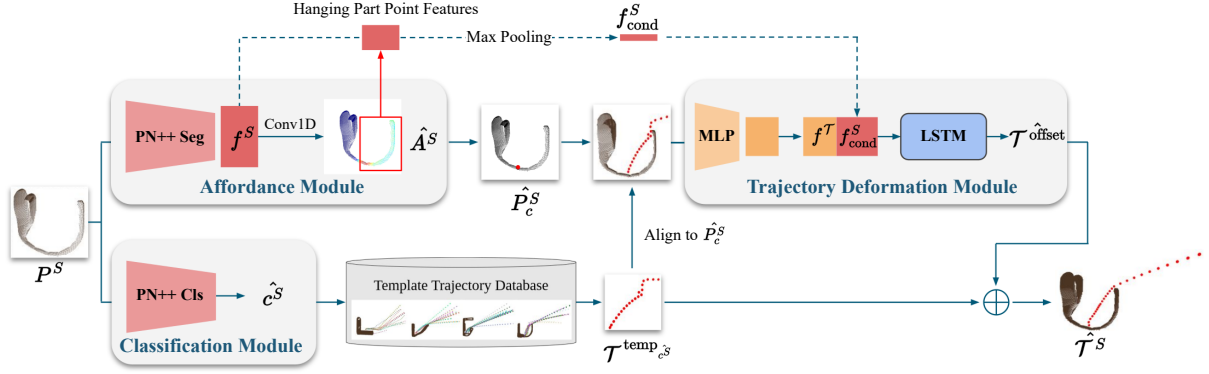


Fig. 3: The architecture of the proposed Shape-Conditioned Trajectory Deformation Network (SCTDN).

template trajectory database. Please refer to section V-B for more details.

Affordance Module. This module predicts the affordance map $A^S \in \mathbb{R}^{N \times 1}$ for the input point cloud \mathcal{P}^S . An affordance map is a heatmap that specifies the probability of each point belonging to the hanging part of a supporting item’s point cloud. The point with the maximal probability is the contact point $\hat{\mathcal{P}}_c^S \in \mathbb{R}^3$, which enables the alignment of the first waypoint of the template trajectory.

Trajectory Deformation Module. This module predicts the waypoint offsets $\mathcal{T}^{\text{offset}} \in \mathbb{R}^{T \times 3}$ to deform the template trajectory $\mathcal{T}^{\text{temp}_{c^S}}$ based on the shape-conditioned feature f_{cond}^S . We obtain the shape-conditioned feature f_{cond}^S by max-pooling the point features corresponding to the hanging part of the input point cloud from the Affordance Module. The reason for including the shape-conditioned feature in trajectory deformation is based on the observation that there is a significant correlation between the SKT and the hanging part of a supporting item.

Inference. Given the partial point cloud \mathcal{P}^S as input, the SCTDN first retrieves a template trajectory $\mathcal{T}^{\text{temp}_{c^S}} = \{\xi_0^{\text{temp}_{c^S}}, \xi_1^{\text{temp}_{c^S}}, \dots, \xi_{T-1}^{\text{temp}_{c^S}}\} \in \mathbb{R}^{T \times 3}$ associated with the shape category \hat{c}^S predicted by the classification module from the template trajectory database. Then, the model determines the contact point $\hat{\mathcal{P}}_c^S$ by selecting the index of the point with the highest probability in \hat{A}^S , predicted by the Affordance Module. We align the first waypoint of the template trajectory $\mathcal{T}^{\text{temp}_{c^S}}$ to the contact point $\hat{\mathcal{P}}_c^S$ by adding $\xi_0^{\text{temp}_{c^S}} - \hat{\mathcal{P}}_c^S$ to each waypoint. Subsequently, the trajectory deformation module takes the aligned $\mathcal{T}^{\text{temp}_{c^S}}$ and a shape-conditioned feature f_{cond}^S obtained from the affordance module as input and generate the waypoint offsets $\mathcal{T}^{\text{offset}} = \{\xi_0^{\text{offset}}, \xi_1^{\text{offset}}, \dots, \xi_{T-1}^{\text{offset}}\} \in \mathbb{R}^{T \times 3}$. The final trajectories is defined as $\hat{\mathcal{T}}^S = \{\{\xi_i^{\text{temp}_{c^S}} + \xi_i^{\text{offset}}, \} | i \in [0, 1, \dots, T-1]\}$.

B. Template Trajectory Database and Affordance Maps

Template Trajectory Database Creation. First, we flatten each trajectory \mathcal{T} in the training set and reduce its dimensionality to 2 using principal component analysis. Following that, we apply the K-means clustering algorithm to cluster the lower dimensional vectors into groups (see Fig. 4 (b)). Finally, the category label $c^S \in [0, K-1]$ is assigned to




Fig. 4: The data post-processing pipeline. (a) Affordance map generation pipeline. (b) Unsupervised categorization to group the supporting items by the corresponding trajectories.

each supporting item based on the closest center. With the above process, we build a template trajectory database for the development of the proposed framework. This database contains the SKTs corresponding to the supporting item that is closest to each center.

Affordance Maps Generation. To facilitate the training of the affordance prediction module, we devise an automatic post-processing pipeline for generating the ground truth affordance maps. Fig. 4 (a) illustrates the generation process. Initially, we identify the points on the partial point cloud closest to the contact point and each waypoint within the SKT. By leveraging mixed Gaussian distributions, we generate two distinct probability maps for each point in a point cloud - the contact map $\mathcal{M}_{\text{cp}}^S$, and the segmentation map $\mathcal{M}_{\text{seg}}^S$. The final affordance map is then rendered by employing a weighted sum of the two maps, denoted as $A^S = \alpha \mathcal{M}_{\text{cp}}^S + \beta \mathcal{M}_{\text{seg}}^S$. Practically, we set $\alpha = 0.5$ and $\beta = 0.5$, and consider the variance σ^2 as a hyperparameter (in our case, it’s set to 0.5 cm).

C. Implementation Details

In our experiments, we configure the point number $N=1,000$, and the category number $K=5$, and the trajectory length $T \in \{10, 20, 40\}$. We identify the hanging part as the point with a probability greater than 0.1 in \hat{A}^S .

Classification Module. We apply a *classification* PointNet++ [24] network to extract the global feature $f^g \in \mathbb{R}^{512}$, then input it to a two-layer MLP network ($512 \rightarrow 256 \rightarrow K$) to predict the class \hat{c}^S . The loss used for this module is the standard cross-entropy loss, denoted as L_c .

Affordance Module. We apply a *segmentation* PointNet++ [24] network to extract the per-point feature $f^S \in$

TABLE I: The success rate evaluation of each approach. † denotes our re-implementation of the baseline methods.

	Inference Time (sec)	All (%)	Easy (%)	Normal (%)	Hard (%)	Very Hard (%)
K-PAM 2.0† [6]	-	40.1	76.0	54.1	21.6	8.8
OmniHang† [2]	14.29	40.1	62.9	47.7	34.3	15.3
VAT-Mart† [5]	0.007	55.2	84.5	82.0	42.0	12.3
Modified VAT-Mart †	0.016	72.1	78.8	89.6	57.7	62.4
SCTDN (Ours)	0.009	83.7	85.5	91.7	79.7	77.7

$\mathbb{R}^{N \times 512}$ of the input point cloud. Subsequently, given the f^S , we apply a Conv1D layer along with the sigmoid function to predict the affordance map \hat{A}^S . We utilize the standard L2 loss as the affordance loss, denoted as L_a .

Trajectory Deformation Module. During the training stage, the inputs for this module consist of a template SKT $\mathcal{T}^{\text{temp}_{c^S}}$ correspond to the input shape category c^S and a shape-conditioned feature f_{cond}^S derived from the per-point features f^S . First, we apply a three-layer MLP network ($T \times 3 \rightarrow 64 \rightarrow 32 \rightarrow 32$) to encode each waypoint in $\mathcal{T}^{\text{temp}_{c^S}}$, resulting in waypoint features $f^{\mathcal{T}} \in \mathbb{R}^{T \times 32}$. Subsequently, we concatenate the shape-conditioned feature f_{cond}^S with each waypoint feature in $f^{\mathcal{T}}$. These combined features serve as inputs to two LSTM [25] blocks, ultimately yielding predicted waypoint offsets $\hat{\mathcal{T}}^{\text{offset}}$. These offsets are used to deform the template trajectory $\mathcal{T}^{\text{temp}_{c^S}}$ for the target supporting item. The loss function for this head is a standard L2 loss, denoted as $L_{\mathcal{T}}$.

Training Loss. The total loss is the weighted sum of the three losses: $L_{\text{total}} = L_{\mathcal{T}} + 0.1 \times L_a + 0.1 \times L_c$.

VI. EXPERIMENT

In this section, we present a series of experiments aimed at answering the following questions: (1) Does our proposed method effectively generate suitable trajectories for hanging a diverse set of grasped objects onto novel supporting items with various shapes compared to existing methods? (2) Do shape-conditioned information and trajectory deformation play crucial roles in generating feasible actions for hanging? (3) Does semantic keypoint trajectory (SKT) facilitate the hanging of diverse objects while being object-agnostic? (4) Can our simulation-trained models transfer to real-world scenarios?

A. Experimental Setup

The environment of a hanging task involves a supporting item and an object grasped by a robot arm. The robot’s objective is to determine a sequence of actions for hanging a grasped object onto a supporting item. For evaluation, we assess the performance of all methods by measuring their success rates and inference times across 3,000 hanging processes involving 50 objects and 60 unseen supporting items (refer to Fig. 5(a)) in PyBullet [23]. These supporting items are manually categorized into four levels of difficulty: **Easy**, **Normal**, **Hard**, and **Very Hard** based on the complexity of the hanging part on them, with 15 assigned to each level.

Baselines. We compare our method with four different baselines: kPAM 2.0 [6], OmniHang [2], VAT-Mart [5], and a modified version of VAT-Mart. For kPAM 2.0, given the

ground truth $SE(3)$ keypoint and the target hanging pose for each object, the robot moves the object from its initial pose to the target pose. For OmniHang [2], given the ground truth hanging pose for each object, we incorporate the collision estimation network they proposed into RRT-Connect [22] to plan a trajectory for hanging. In the case of VAT-Mart and the modified version, given the $SE(3)$ keypoint for each object, we closely follow their implementation of the trajectory proposal network to generate 10 different SKTs for hanging each object onto a supporting item. The trial is considered successful if one of these trajectories successfully hangs the object onto the supporting item. For the modified VAT-Mart, we utilize the shape-conditioned point features introduced in V-A, instead of the contact point feature used in VAT-Mart. For SCTDN, given the $SE(3)$ keypoint for each object, we use one predicted SKT for evaluation.

B. Hanging Success Rate and Inference Time

We report the success rate of robot hanging and the corresponding inference time in Table I. Our method demonstrates promising performance, achieving an average success rate of 83.7%, compared to all baselines. In particular, the proposed method shows superior performance in challenging cases, i.e., **Hard** and **Very Hard**. The modified VAT-Mart with the shape-conditioned point features significantly outperforms the original VAT-Mart, especially in the **Very Hard** level. Our experiments corroborate that the methods based on SKT outperform kPAM and OmniHang by a substantial margin. While the success rate of OmniHang improves for challenging levels compared to kPAM, it still faces significant challenges in detecting collisions under partial observability. It is worth noting that, all learning-based methods are capable of generation trajectories within 20 milliseconds, while OmniHang, on average, requires 14.29 sec for planning.

C. Ablation Study of the SCTDN Architecture

In this experiment, we study the contribution of the shape-conditioned features and trajectory deformation module. The results are shown in Table II. The model achieves a 58.2% success rate when we only align the retrieved template trajectory to the predicted contact point, without considering the deformation module. We show that the deformation module boosts the hanging success rates significantly. When we use the contact point feature for the deformation module, we observe a 15.3% performance gain. On the other hand, a 25.5% performance improvement is obtained with the shape-conditioned feature introduced in V-A. The results indicate the importance of the trajectory deformation module and

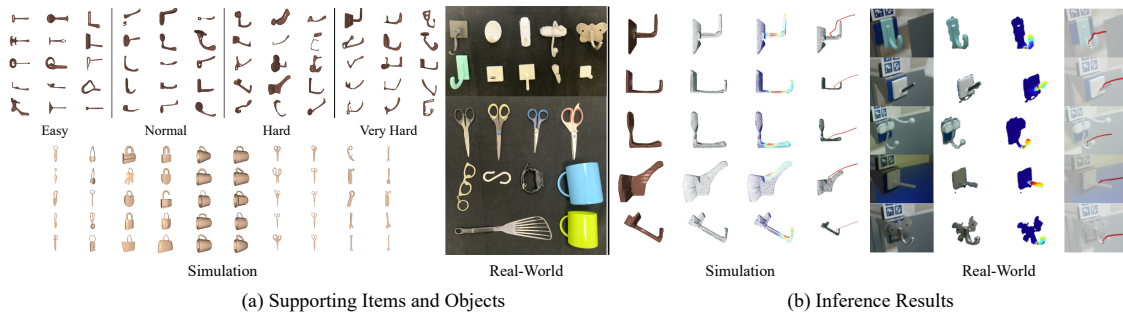


Fig. 5: The supporting items and objects for evaluations and the visualization of point clouds, affordance maps, and the semantic keypoint trajectories for each supporting item.

TABLE II: Ablation study of our network designs.

Contact Point	Contact Point Feature	Part Feature	Deformation	Success Rate (%)
✓				58.2
✓	✓		✓	73.5
✓		✓	✓	83.7

support that it is crucial to model the geometric structures of the supporting item’s hanging part.

D. One Semantic Keypoint Trajectory for Various Objects

In this experiment, we assess whether an SKT can be applied to various objects for each supporting item. We apply modified VAT-Mart and SCTDN to predict SKT for each supporting item. We compare the overall success rates of hanging a reference object and five additional object categories (each with ten instances) onto 60 different supporting items. Note that the reference object is used to generate trajectories for all supporting items, discussed in Sec. III-B. The results reported in Table III indicate that SKT enables similar success rates of hanging different objects compared to the reference object. This greatly enhances the versatility and applicability of SKT to a wide range of everyday objects. Moreover, the SCTDN demonstrates a promising capability to generate SKT, compared to modified VAT-Mart.

E. Real-World Evaluation

We conducted real-world experiments to evaluate the transferability of our simulation-trained models to real-world scenarios. Our experimental setup included a Franka Emika robot arm, an Intel D-435 camera, an object grasped by the robot arm, and a supporting item. In the real-world evaluation, we used a total of 10 objects and 10 supporting items, as shown in Fig. 5(a). To handle noise in real-world point clouds, we defined the foreground region of the supporting item using an AprilTag and applied DBSCAN [26] to remove noisy points. Alternatively, one can use image segmentation models [27] to obtain a segmented point cloud.

The hanging success rate evaluation involved the robot hanging all 10 objects onto the 10 supporting items, resulting in 100 hanging processes. The overall success rate for hanging was 92%, indicating the effectiveness of our approach in guiding various objects onto different supporting items in real-world scenarios. The right part of Fig. 5(b) visualizes the predicted Semantic Keypoint Trajectories (SKTs) in simulation and real-world settings. For additional qualitative results and demonstrations, please refer to our video submission and the [webpage](#).

TABLE III: The hanging success rate of reference objects and other objects. † denotes our re-implementation of the baseline methods.

	Modified VAT-Mart †	SCTDN (Ours)
Reference Object (%)	71.7	91.7
Mug (%)	78.0	88.8
Cooking Utensil (%)	70.7	81.2
Scissor (%)	67.9	75.3
Tool (%)	61.5	78.0
Others (%)	68.8	84.3

F. More Ablation Studies of SCTDN and VAT-Mart

We conduct additional ablation studies to compare the hanging success rates of SCTDN and VAT-Mart-based methods across various trajectory lengths $T \in \{10, 20, 40\}$ and different waypoint dimensions (position-only and position-and-rotation waypoints) within the SKT. For comprehensive discussions and detailed table information, please visit our project webpage.

VII. CONCLUSION

In this paper, we propose *Semantic Keypoint Trajectory*, a novel representation for hanging everyday objects. This representation is actionable and interpretable, specifying both where and how to hang while being object-agnostic. This enables the robot to efficiently learn the complex hanging task. To predict such a representation, we introduce a shape-conditioned trajectory deformation network by leveraging the insight that crucial geometric parts play a significant role in trajectory generation. By deforming existing trajectories, we are able to obtain a semantic keypoint trajectory adapted to the shape of novel supporting items. We conducted empirical evaluations through robot executions to hang various objects on a diverse range of supporting items, demonstrating the effectiveness of our proposed approach in both the simulation and the real world. Our work holds potential for broader application in other relevant tasks.

VIII. ACKNOWLEDGEMENT

The work is sponsored in part by the Higher Education Sprout Project of the National Yang Ming Chiao Tung University and Ministry of Education (MOE), the Yushan Fellow Program Administrative Support Grant, and the National Science and Technology Council (NSTC) under grants 110-2222-E-A49-001-MY3, 110-2634-F-002-051, 111-2634-F-002-022-, Mobile Drive Technology Co., Ltd (MobileDrive), and Industrial Technology Research Institute Mechanical and Mechatronics Systems Lab.

REFERENCES

- [1] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, “kpam: Keypoint affordances for category-level robotic manipulation,” in *The International Symposium of Robotics Research*. Springer, 2019, pp. 132–157.
- [2] Y. You, L. Shao, T. Migimatsu, and J. Bohg, “OmniHang: Learning to hang arbitrary objects using contact point correspondences and neural collision estimation,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5921–5927.
- [3] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, “Neural descriptor fields: Se (3)-equivariant object representations for manipulation,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6394–6400.
- [4] C. Pan, B. Okorn, H. Zhang, B. Eisner, and D. Held, “Tax-pose: Task-specific cross-pose estimation for robot manipulation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1783–1792.
- [5] R. Wu, Y. Zhao, K. Mo, Z. Guo, Y. Wang, T. Wu, Q. Fan, X. Chen, L. Guibas, and H. Dong, “VAT-Mart: Learning visual action trajectory proposals for manipulating 3D articulated objects,” in *International Conference on Learning Representations (ICLR)*, 2022.
- [6] W. Gao and R. Tedrake, “kPAM 2.0: Feedback control for category-level robotic manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2962–2969, 2021.
- [7] P. R. Florence, L. Manuelli, and R. Tedrake, “Dense Object Nets: Learning dense visual object descriptors by and for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2018, pp. 373–385.
- [8] Z. Qin, K. Fang, Y. Zhu, L. Fei-Fei, and S. Savarese, “KETO: Learning keypoint representations for tool manipulation,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7278–7285.
- [9] M. Vecerik, J.-B. Regli, O. Sushkov, D. Barker, R. Pevceciciute, T. Rothörl, R. Hadsell, L. Agapito, and J. Scholz, “S3K: Self-supervised semantic keypoints for robotic manipulation via multi-view consistency,” in *Conference on Robot Learning*. PMLR, 2021, pp. 449–460.
- [10] Z. Xue, Z. Yuan, J. Wang, X. Wang, Y. Gao, and H. Xu, “Useek: Unsupervised se (3)-equivariant 3d keypoints for generalizable manipulation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1715–1722.
- [11] P. Florence, L. Manuelli, and R. Tedrake, “Self-supervised correspondence in visuomotor policy learning,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 492–499, 2019.
- [12] L. Manuelli, Y. Li, P. Florence, and R. Tedrake, “Keypoints into the future: Self-supervised correspondence in model-based reinforcement learning,” *arXiv preprint arXiv:2009.05085*, 2020.
- [13] K. Mo, L. J. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani, “Where2act: From pixels to actions for articulated 3d objects,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6813–6823.
- [14] Z. Xu, Z. He, and S. Song, “Universal manipulation policy network for articulated objects,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2447–2454, 2022.
- [15] S. Liu, S. Tripathi, S. Majumdar, and X. Wang, “Joint hand motion and interaction hotspots prediction from egocentric videos,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 3282–3292.
- [16] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, “Affordances from human videos as a versatile representation for robotics,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 778–13 790.
- [17] T. Weng, S. M. Bajracharya, Y. Wang, K. Agrawal, and D. Held, “Fabricflownet: Bimanual cloth manipulation with a flow-based policy,” in *Conference on Robot Learning*. PMLR, 2022, pp. 192–202.
- [18] B. Eisner, H. Zhang, and D. Held, “Flowbot3d: Learning 3d articulation flow to manipulate articulated objects,” *Robotics Science and Systems (RSS)*, 2022.
- [19] D. Seita, Y. Wang, S. J. Shetty, E. Y. Li, Z. Erickson, and D. Held, “Toolflownet: Robotic manipulation with tools via predicting tool flow from point clouds,” in *Conference on Robot Learning (CoRL)*. PMLR, 2023, pp. 1038–1049.
- [20] K. Takeuchi, I. Yanokura, Y. Kakiuchi, K. Okada, and M. Inaba, “Automatic hanging point learning from random shape generation and physical function validation,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4237–4243.
- [21] A. Nichol, H. Jun, P. Dhariwal, P. Mishkin, and M. Chen, “Point-E: A system for generating 3D point clouds from complex prompts,” *arXiv preprint arXiv:2212.08751*, 2022.
- [22] J. J. Kuffner and S. M. LaValle, “RRT-connect: An Efficient Approach to Single-query Path Planning,” in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2000, pp. 5921–5927.
- [23] E. Coumans and Y. Bai, “PyBullet: a Python module for physics simulation for games, robotics and machine learning,” <https://pybullet.org>, 2016–2021.
- [24] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “PointNet++: Deep hierarchical feature learning on point sets in a metric space,” *Advances in neural information processing systems*, vol. 30, 2017.
- [25] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [26] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Conference on Knowledge Discovery and Data Mining (KDD)*, vol. 96, no. 34, 1996, pp. 226–231.
- [27] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023.