

Learning Distributional Models for Relative Placement in Robotics

Jenny Wang

CMU-RI-TR-24-24

May, 2024

School of Computer Science

The Robotics Institute

Carnegie Mellon University

Pittsburgh, Pennsylvania

Thesis Committee:

David Held, *Chair*

Oliver Kroemer

Shubham Tulsiani

Nikolaos Gkanatsios

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics*

Copyright ©2024 Jenny Wang

Abstract

Relative placement tasks are an important category of tasks in which one object needs to be placed in a desired pose relative to another object. Previous work has shown success in learning relative placement tasks from just a small number of demonstrations when using relational reasoning networks with geometric inductive biases. However, such methods cannot flexibly represent multimodal tasks, like a mug hanging on any of n racks. We propose a method that incorporates additional properties that enable learning multimodal relative placement solutions, while retaining the provably translation-invariant and relational properties of prior work. We show that our method is able to learn precise relative placement tasks with only 10-20 multimodal demonstrations with no human annotations across a diverse set of objects within a category.

Acknowledgements

This could not have come to be without all the support from the people around me.

First, I would like to thank my advisor, David Held, for his relentless support throughout it all. His guidance brought about a certain rigor and methodology to the work that allows me to point back at it and truly believe, "Yes, that led to that." The TAX-Pose authors and my fellow lab mates Chuer Pan, Brian Okorn, Harry Zhang, and Ben Eisner were very helpful in catching me up and getting this project going. I would also like to thank my committee members Oliver Kroemer, Shubham Tulsiani, and Nikolaos Gkanatsios for their advice on this work.

I would also like to thank my lab mates for the advice and technical help they have given me. In particular, I thank Ben Eisner for helping me set up compute and hearing me out, Thomas Weng for all the times I needed an Ubuntu live USB, and Wenxuan Zhou for reaching out at the beginning. Additionally, I thank Octavian Donca for coming into the project strong and encouraging me to move forward. There are also TJ Vitchutripop, Yishu Li, Bowen Jiang, Pranay Gupta, Zhanyi Sun, Angela Rodriguez-Izquierdo, Chuer Pan, Lifan Yu and so many more who I thank for being a good vibe and creating awesome times.

Furthermore, I'd like to thank some good friends I've made here at CMU who aren't in the lab—Gabrielle Ohlson, Kavya Puthuveetil, Eliot Xing, Je-han Yang, Zulekha Karachiwalla, and everyone in between, thank you for creating fun memories with me in our random shenanigans. And, to cap everything off, I would like to thank my mom, my dad, and my brother for all the support during this program and leading all the way up to it.

Thank you.

Funding

This material is based upon work supported by the United States United States Air Force and DARPA under Contract No. FA8750-18-C-0092, NIST under Grant No. 70NANB23H178, and the Uber Presidential Fellowship.

Contents

1	Introduction	1
1.1	Overview	1
1.2	Contributions of this thesis	2
1.3	Related Publications	2
2	Background	3
2.1	Action Representations	3
2.2	Relative Placement Tasks	3
2.3	Multimodal Action Sets	4
3	Applying Variational Models to Relative Placement	5
3.1	Introduction	5
3.2	Problem Statement	5
3.3	Method	7
3.4	Experiments	13
3.5	Results	14
4	Conclusions and Future Work	17
5	Appendix	18
5.1	Training Details	18
5.2	Task Details	20
5.3	Additional experiments	21
5.4	Interpretability of the the learned latent distribution	22
5.5	Proof for rotational invariance for the featurization	23

List of Tables

3.1	Task success rates for the original NDF 1 rack mug hanging task and multimodal variants. We evaluate on unseen mugs initialized to arbitrary configurations in space. <i>1 Rack:</i> Our model performs on par with TAX-Pose [14] in the unimodal placement task and outperforms the other baselines. <i>2 Racks:</i> We extend TAXPose to a multimodal task and train our model for random 2-rack scenes. We evaluate on 50 random 2-rack configurations. <i>3 Racks (Generalization):</i> We further evaluate the 2 rack model to test generalization to more racks than seen at training time. Three racks are placed in 50 random configurations. We report placement task success.	14
3.2	Ablations. We show ablations of our model for the multimodal 2-racks mug-hanging task and also for generalization to the unseen 3-racks task. We report placement success, and also rotation error ($^{\circ}$) and translation error (m) relative to the closest demonstration. Using a discrete spatial z distribution over points in the point cloud performs better than using a continuous z vector; using a learned prior performs better than using a non-learned prior (Normal or uniform), especially for generalization to “3 Racks.”	16

List of Figures

3.1	Our method’s learned prior $p_\phi(z X)$ learns a distribution over modalities for the task, for example placing a mug on the left rack or the right rack. During inference time, this allows the model to predict a diverse set of ways to perform the task. . . .	6
3.2	Model overview. We train an encoder $q_\psi(z Y)$ to map from demonstration placements Y into a spatially grounded latent variable z . Using a dense and SE(3)-invariant representation of z , \mathcal{F}_p , we condition the TAX-Pose decoder, $p_\theta(Y X, z)$, to predict the cross-pose $\hat{\mathbf{T}}_{AB}$ that reconstructs Y from the SE(3) transformed action and anchor point clouds, X_A and X_B . Additionally, we train a learned prior $p_\phi(z X)$ to map from points X to points in the latent space, capturing the different modes of the demonstrations. At inference time, we decode samples from this learned prior to generate valid cross-poses for the objects in the observation X	8
3.3	Visualization of a successful relative placement for hanging a mug with an arbitrary pose on one of many racks in a random 3 rack scene. The mug is initialized to an arbitrary pose in the air and must be grasped first by the robot.	13
3.4	Samples from our model trained on various multimodal mug-on-rack environments cover a wide distribution of placements that correspond to task success. Top row- one sample, denoted by orange. Bottom row- multiple samples denoted by orange, red, and green.	15
5.1	We show that TAX-PoseD learns to cover a wide variety of ways to complete a task. Blue- an anchor object depicting the scene. Red, Green, and Orange- manipulated objects at various poses sampled from TAX-PoseD.	22
5.2	Our method learns disentangled representations for relative placement tasks. Empirically, the latent distribution assigns high probability to the rack upon which the mug is hanged in the demonstration.	22
5.3	We show visualizations of the learned latent space across the Multi-Peg Rack, Multi-Peg Multi-Rack, and Bowl Grasping tasks. Empirically, the high probability points tend to correspond spatially to placements of the manipulation object.	23

Chapter 1

Introduction

1.1 Overview

Many robotic manipulation tasks can be framed as relative placement tasks. For example, hanging a mug on a mug rack requires placing the mug in a position relative to one of the pegs of the rack. Even complex, long-horizon tasks such as organizing a cluttered table can be framed as a series of relative placements: first, predict an $SE(3)$ transformation that stacks one book on top of another, then predict a transformation that puts the pencil in the pencil box, and then predict a transformation that centers the keyboard in front of the monitor.

Previous work such as TAX-Pose [14] has shown that, for relative placement tasks, using network architectures that explicitly reason about object relationships helps the network to generalize significantly better across object poses and instances. However, this previous work outputs only a single relative placement prediction for each observation. In multi-modal settings, this leads to predictions which are the mean of valid placement modes, which may be incorrect. For example, suppose that a set of demonstrations place a mug on any of n racks in the scene. The average of these demonstrations will be a point in the middle of the racks, which is not a valid placement. In contrast, many tasks are defined by a distribution of relationships: a robot may be tasked to grasp anywhere along the rim of a bowl, place a fork on the left of any of the plates (e.g. when setting the table), or grasp any one of a cabinet’s drawers.

To address these challenges, we present TAX-PoseD, a Distributional variant of TAX-Pose [14]. Our method predicts task-specific object relationships from just a few demonstrations and no human annotations, while robustly accounting for multimodal demonstration distributions.

1.2 Contributions of this thesis

In this thesis, we take steps towards understanding how to leverage geometric biases present in relative placement tasks for task-specific relative pose prediction. We analyze four practical implementations for distributional models and show how these design choices impact a neural network’s prediction precision and also generalization across task variants.

Our core technical contributions include:

- A method for efficiently learning distributional relative placement tasks; our approach extends TAX-Pose [14] to handle multimodal, distributional demonstrations.
- A novel spatially-grounded architecture for a cVAE [21] that represents the latent variable distribution as a categorical distribution over 3D points; this results in a grounded and interpretable latent space that avoids the smoothing effect commonly found in cVAEs, leading to significantly improved performance for multi-modal placement tasks.

We evaluate our method on challenging multimodal tasks and evaluate its generalization across diverse objects within a category. We demonstrate that our method is both interpretable and achieves strong performance on distributional relative placement tasks.

1.3 Related Publications

Most ideas described in this thesis have appeared in the following publication:

- J. Wang, O. Donca, and D. Held. Learning Distributional Demonstration Spaces for Task-Specific Cross-Pose Estimation. International Conference on Robotics and Automation (ICRA), 2024.

Chapter 2

Background

2.1 Action Representations

Many action representations have been explored that enable robots to increase their learning efficiency when learning to solve manipulation tasks. For example, representations can be per-point [4, 12, 27] or consist of keypoints [17]. Additionally, architectures that leverage local geometry have been shown to provide useful priors for learning from point cloud data [10, 16, 24, 28]. Our proposed method combines the strengths of dense per-point representations and geometric inductive biases while also reasoning over distributions for relative placement tasks.

2.2 Relative Placement Tasks

Many tasks can be defined as a sequence of relative placement tasks, such as hanging a mug on a rack or putting a dish into a microwave. These tasks are “relative” placements because the desired pose of the first object is not based on a fixed set of coordinates in the world frame but rather it is defined relative to some other object in the scene. A number of papers have developed methods designed for this class of tasks [14, 18–20]. This thesis builds upon TAX-Pose [14], which achieves strong performance on unimodal tasks such as hanging a single mug on a single rack; however, as we show in this work, TAX-Pose performs poorly on multimodal tasks such as when there are multiple mug racks available. Our method handles such symmetries by learning a distribution over relative placements. Symmetries are also addressed by concurrent work [20], which uses diffusion models. However, diffusion models are often slow during inference; in contrast, our method can operate at high inference speeds because we require only a single forward pass through our network.

2.3 Multimodal Action Sets

Actions that successfully complete a task may occupy multiple regions of an agent’s action space. When learning from demonstrations, an architecture that does not model the multimodality of the demonstration space may tend to output the average over correct actions, which leads to failures. Implicit representations have been explored for modeling the multimodality of a task [5]. Multi-modal action reasoning can be implemented as spatial action maps [26], GANs [15], particle filters [25], diffusion models [2,20], and VAEs [1,13]. We build upon the VAE literature and apply it to modeling modes of the demonstration space for relative placement tasks.

Chapter 3

Applying Variational Models to Relative Placement

3.1 Introduction

We present TAX-PoseD, a method for relative placement tasks that accounts for multimodality over task solutions by sampling from a learned latent distribution over modes in the solution set. Using a Conditional Variational Autoencoder (cVAE) [21], our encoder maps a successful task demonstration to a distribution from which we sample a latent z , and our decoder maps z to a reconstruction of the same scene. During inference time, we sample a latent z from a prior and then decode it to a (hopefully) successful placement pose. However, we find that a traditional cVAE architecture of using a latent z vector defined in an arbitrary vector space fails to learn meaningful modes for this task, leading to poor performance. In contrast, we propose a spatially-grounded latent space over points in the point cloud which leads to significantly improved performance. Below, we describe our approach in more detail.

3.2 Problem Statement

We consider the task of relative placement, for which an object must be placed in some configuration relative to another object. For example, for the mug hanging task, a mug must be placed in some pose relative to a mug rack.

Following Pan et al [14], success for a relative placement task can be defined as follows: Let $\mathbf{T}_{\mathcal{A}}^*$ and $\mathbf{T}_{\mathcal{B}}^*$ be a pair of SE(3) poses for objects \mathcal{A} and \mathcal{B} , respectively (in some fixed world frame), for which the task is considered successful. For example, $\mathbf{T}_{\mathcal{A}}^*$ can be the pose for a mug, and $\mathbf{T}_{\mathcal{B}}^*$ can be the pose for a mug rack, for which the mug is on the mug rack. Given a pair of poses $\mathbf{T}_{\mathcal{A}}^*$

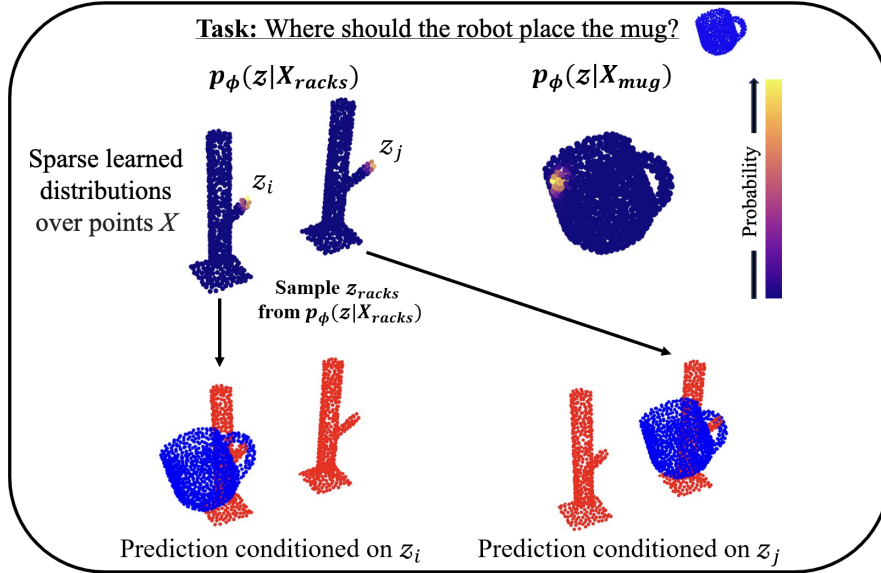


Figure 3.1: Our method’s learned prior $p_\phi(z | X)$ learns a distribution over modalities for the task, for example placing a mug on the left rack or the right rack. During inference time, this allows the model to predict a diverse set of ways to perform the task.

and \mathbf{T}_B^* that complete the task, for relative placement tasks, the task is also considered successful when any $SE(3)$ transformation \mathbf{T} is applied to both objects \mathcal{A} and \mathcal{B} . In other words, for a relative placement task, if the task is successful for poses \mathbf{T}_A^* and \mathbf{T}_B^* , then the task is also successful for poses $\mathbf{T} \cdot \mathbf{T}_A^*$ and $\mathbf{T} \cdot \mathbf{T}_B^*$. Following this reasoning, suppose that objects \mathcal{A} and \mathcal{B} are in poses \mathbf{T}_A and \mathbf{T}_B , respectively, and suppose that \mathbf{T}_A^* and \mathbf{T}_B^* are defined as above. Then poses \mathbf{T}_A and \mathbf{T}_B must satisfy the following Boolean function to be considered a successful relative placement:

$$\begin{aligned} \text{RelPlace}(\mathbf{T}_A, \mathbf{T}_B) &= \text{SUCCESS} & (3.1) \\ &\text{iff } \exists \mathbf{T} \in SE(3) \\ &\text{s.t. } \mathbf{T}_A = \mathbf{T} \cdot \mathbf{T}_A^* \text{ and } \mathbf{T}_B = \mathbf{T} \cdot \mathbf{T}_B^*. \end{aligned}$$

In practice, we are usually not given the transforms \mathbf{T}_A and \mathbf{T}_B but rather we are given the point clouds for these objects \mathbf{P}_A and \mathbf{P}_B . In order to solve a relative placement task, we need to estimate an $SE(3)$ transformation $\mathbf{T}_{AB} := f(\mathbf{P}_A, \mathbf{P}_B)$ that transforms object \mathcal{A} in such a manner that satisfies RelPlace (Eqn. 3.1). In other words, we need to find the transform \mathbf{T}_{AB} such that $\text{RelPlace}(\mathbf{T}_{AB} \cdot \mathbf{T}_A, \mathbf{T}_B) = \text{SUCCESS}$. In TAX-Pose [14], this transformation \mathbf{T}_{AB} is referred to as the *cross-pose* between objects \mathcal{A} and \mathcal{B} .

Often, there is a set of valid solutions that all solve the task, for example when a plate should

be placed within any of the placemats on a table. Thus, we extend the definition of task success given by Equation 3.1 to the distributional setting. For distributional relative placement tasks, suppose that object \mathcal{B} is in pose \mathbf{T}_B^* . Then let $\{\mathbf{T}_{\mathcal{A},1}^*, \dots, \mathbf{T}_{\mathcal{A},N}^*\}$ be a set of N poses for object \mathcal{A} that satisfy the relative placement task. Because this is a relative placement task, then the task is also successful for poses $\mathbf{T} \cdot \mathbf{T}_{\mathcal{A},i}^*$ and $\mathbf{T} \cdot \mathbf{T}_B^*$ (for objects \mathcal{A} and \mathcal{B} respectively) for all $i \in [1, N]$ for any SE(3) transformation \mathbf{T} . Thus, suppose that objects \mathcal{A} and \mathcal{B} are in poses \mathbf{T}_A and \mathbf{T}_B , respectively. Then these poses must satisfy the following Boolean function to be considered a successful relative placement for a *distributional* relative placement task:

$$\begin{aligned} \text{RelPlace}_D(\mathbf{T}_A, \mathbf{T}_B) &= \text{SUCCESS} & (3.2) \\ \text{iff } \exists i \in [1, N], \exists \mathbf{T} \in \text{SE}(3) \\ \text{s.t. } \mathbf{T}_A &= \mathbf{T} \cdot \mathbf{T}_{\mathcal{A},i}^* \text{ and } \mathbf{T}_B = \mathbf{T} \cdot \mathbf{T}_B^*. \end{aligned}$$

In practice, we are not usually given a comprehensive set of poses $\{\mathbf{T}_{\mathcal{A},1}^*, \dots, \mathbf{T}_{\mathcal{A},N}^*\}$ for object \mathcal{A} that solve a relative placement task relative to some pose \mathbf{T}_B^* for object \mathcal{B} . Instead, we are often given a set of demonstration pairs $(\mathbf{T}_{\mathcal{A},j}^*, \mathbf{T}_{\mathcal{B},j}^*)$ for $j \in [1, M]$ for a set of M demonstrations. From these demonstration pairs, we aim to learn a generative function $f_D(\mathbf{P}_A, \mathbf{P}_B)$ such that sampling from this function gives a solution to the distributional relative placement task, where \mathbf{P}_A and \mathbf{P}_B are point clouds for objects \mathcal{A} and \mathcal{B} respectively. In other words, sampling from $f_D(\mathbf{P}_A, \mathbf{P}_B)$ gives an SE(3) transformation \mathbf{T}_{AB} such that $\text{RelPlace}_D(\mathbf{T}_{AB} \cdot \mathbf{T}_A, \mathbf{T}_B) = \text{SUCCESS}$. We refer to any such transformation \mathbf{T}_{AB} as a cross-pose, and the function $f_D(\mathbf{P}_A, \mathbf{P}_B)$ outputs a distribution over cross-poses.

For this thesis, we also assume that the demonstrations are given by segmented point clouds, in which we have a segmentation between the objects \mathcal{A} and \mathcal{B} in the demonstration.

3.3 Method

3.3.1 Relative Placement as a Variational Problem

Conditional Variational AutoEncoders (cVAEs) [21] are a variant of VAEs [9] that condition the encoder and decoder on additional inputs to control the data generation process. Given an input variable X and output variable Y , the tractable variational lower bound of a cVAE (the ELBO) is written as the following [21]:

$$\begin{aligned} \log p_\theta(Y | X) &\geq \mathbb{E}_{q_\psi(z|X,Y)} [\log p_\theta(Y | X, z)] & (3.3) \\ &\quad - KL(q_\psi(z | X, Y) || p_\phi(z | X)) \end{aligned}$$

with the prior $p_\phi(z | X)$ and an encoder $q_\psi(z | X, Y)$ for latent variable z . The ELBO (Eq. 3.3) can be thought of as the sum of a reconstruction loss (first term) and a regularization loss (second

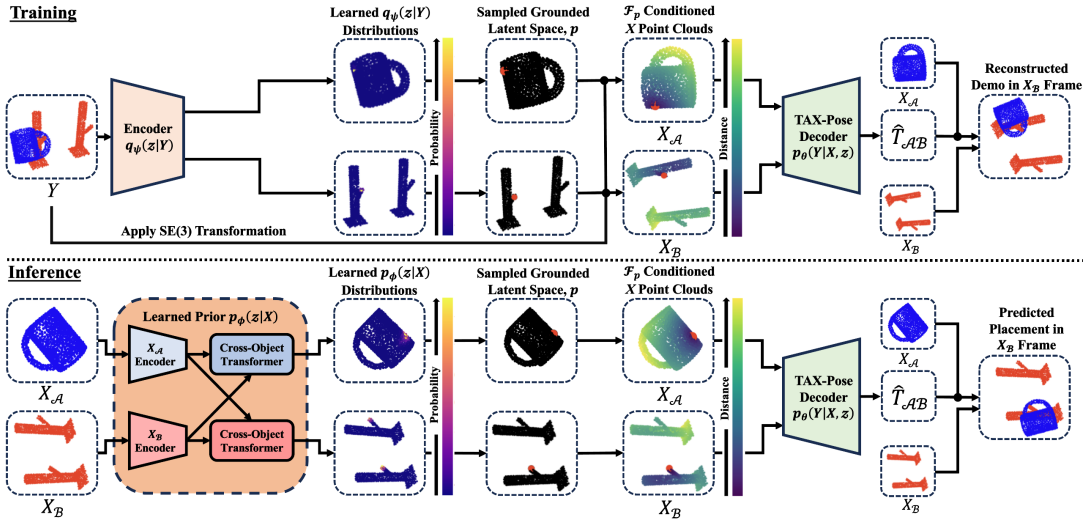


Figure 3.2: Model overview. We train an encoder $q_\psi(z | Y)$ to map from demonstration placements Y into a spatially grounded latent variable z . Using a dense and $SE(3)$ -invariant representation of z , \mathcal{F}_p , we condition the TAX-Pose decoder, $p_\theta(Y | X, z)$, to predict the cross-pose \hat{T}_{AB} that reconstructs Y from the $SE(3)$ transformed action and anchor point clouds, X_A and X_B . Additionally, we train a learned prior $p_\phi(z | X)$ to map from points X to points in the latent space, capturing the different modes of the demonstrations. At inference time, we decode samples from this learned prior to generate valid cross-poses for the objects in the observation X .

term). In the relative placement setting, Y is the point cloud of objects in a demonstration pose for a relative placement task, and X is an observation of these objects in arbitrary poses. During training, we obtain X by transforming the demonstration objects Y with arbitrary transforms sampled from $SE(3)$. Our goal is to learn a decoder $p_\theta(Y | X, z)$ that, conditioned on a latent z and the observation X , predicts the object configuration in the demonstration Y .

As we note below, this straight-forward implementation of a cVAE does not work well for learning to imitate demonstrations of relative placement tasks. Below, we discuss the modifications we make to the cVAE framework for our task. For example, rather than predicting a point cloud Y directly, we instead predict a cross-pose T_{AB} that can be applied to the objects in the observation X to move them into the demonstration configuration. The robot can then move object \mathcal{A} by transform T_{AB} to move it into the goal pose. Further, our experiments show that a continuous latent variable z is not appropriate for relative placement tasks that have discrete multimodalities, such as placing a mug on one of k distinct racks. Thus, we instead use a discrete, spatially grounded latent space, as described below.

3.3.2 Learning a Spatially-Grounded Latent Space for Demonstrations

Encoder:

First, we learn an encoder $q_\psi(z | X, Y)$ that compresses demonstration point clouds Y into a latent space z . Since we focus on relative placement tasks, the demonstration point cloud Y indicates an arrangement of objects that achieve a successful relative placement, and the latent encoding z is a compressed form of this demonstration.

As explained in Section 3.2, for relative placement tasks, an object \mathcal{A} must be placed in some configuration relative to another object \mathcal{B} . The demonstration point clouds Y show examples of successful relative placements of objects \mathcal{A} and \mathcal{B} . During training, we obtain the input X by transforming the objects \mathcal{A} and \mathcal{B} in the demonstration point cloud Y with arbitrary transformations sampled from $SE(3)$. Thus, to encode the demonstration Y we do not need knowledge of X . We can thereby simplify the encoder $q_\psi(z | X, Y)$ as just $q_\psi(z | Y)$ since z is conditionally independent of X given Y , as shown in Figure 3.2. More details about the encoder implementation can be found in Appendix 5.1.

Decoder:

For the decoder $p_\theta(Y | X, z)$, we must decode the latent z and the object point cloud X into the demonstration point cloud Y . However, training a decoder to predict the demonstration point cloud Y directly would require the decoder to output unnecessarily complex details about the scene. Instead, we train the decoder to predict a transform $\mathbf{T}_{AB} \in SE(3)$ that would move the objects in X to their corresponding poses in the demonstration Y (referred to as a “cross-pose”, see Section 3.2 for details). Further, this predicted cross-pose \mathbf{T}_{AB} is more useful than predicting a point cloud Y since we can then command the robot to transform the objects in X by transform \mathbf{T}_{AB} to achieve the task. To make this more clear, we will write the decoder below as $p_\theta(\mathbf{T}_{AB} | X, z)$, to demonstrate that it outputs a distribution over transforms \mathbf{T}_{AB} . In practice, the decoder will be implemented as a latent-conditioned version of TAX-Pose [14]; more details about the decoder implementation can be found in Appendix 5.1.

In practice, during inference, we will first sample a latent z from the prior $p_\phi(z | X)$. Conditioned on this sampled z and the input X , we will then use the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$ to output a (deterministic) transform $\hat{\mathbf{T}}_{AB}$. We will then command the robot to move object \mathcal{A} by the sampled transform to achieve the relative placement task. Thus, combining the prior distribution with the deterministic decoder together gives a distribution over transforms \mathbf{T}_{AB} .

Reconstruction loss:

As shown in Eqn. 3.3, the encoder and decoder are supervised using a reconstruction loss and a regularization loss. In more detail, during training we first encode the demonstration point cloud Y into a distribution over latent z using the encoder $q_\psi(z | Y)$. After sampling z from this distribution, we decode the latent z and the input X using the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$ into a predicted transform $\hat{\mathbf{T}}_{AB}$. The reconstruction loss term $\mathbb{E}_{q_\psi(z|X,Y)} [\log p_\theta(\mathbf{T}_{AB} | X, z)]$ corresponds to the difference between the demonstration point cloud Y and the reconstructed scene after applying the predicted transform \mathbf{T}_{AB} to the point cloud for object \mathcal{A} .

In more detail, let $X_{\mathcal{A}}$ and $X_{\mathcal{B}}$ be the point clouds for object \mathcal{A} and \mathcal{B} respectively. Then let us define the reconstructed point cloud $\hat{Y} = (\hat{\mathbf{T}}_{AB} \cdot X_{\mathcal{A}}) \cup X_{\mathcal{B}}$ as the result of applying the predicted transform $\hat{\mathbf{T}}_{AB}$ to point cloud $X_{\mathcal{A}}$ and then concatenating point cloud $X_{\mathcal{B}}$. We will train the encoder and decoder such that, if the encoder $q_\psi(z | Y)$ and the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$ were lossless, then \hat{Y} would be identical to the original demonstration point cloud Y . The reconstruction loss computes the distance between the demonstration point cloud Y and the reconstructed point cloud \hat{Y} , using similar losses as in TAXPose [14] (see Appendix 5.1 for details).

Latent space:

As mentioned above, we encode the demonstration point cloud Y into a latent z using the encoder $q_\psi(z | Y)$. However, we found it difficult to train such a model if we represent the latent z as an arbitrary vector in \mathbb{R}^d like in traditional VAEs (see ablations in subsection 3.5.1). We believe that this is because of the “smoothing” effect of cVAEs; if there are two discrete mug racks, some portion of the latent z space should model placing the mug on rack 1, and some portion of the latent z space should model placing the mug on rack 2. However, because the latent space is smooth, there must be points that lie in the transition region between these spaces, and the model would have difficulty mapping such points to a valid output. Further, if the next scene has 3 racks, then the latent space must be divided further, creating training difficulties with consistently mapping the latent space to the set of valid solutions.

We avoid these issues by training a *spatially-grounded* latent space. We represent the latent variable z as a sample from a discrete N -dimensional distribution, where N is the number of points in the point cloud X ; in other words, the latent variable z is represented as a categorical distribution over the points in the scene. Thus, the prior $p_\phi(z | X)$ as well as the encoder $q_\psi(z | Y)$ both represent multinomial distributions over all of the points in the point clouds X and Y respectively. In practice, we actually represent the latent space by two separate categorical distributions, one over the points in $X_{\mathcal{A}}$ and another over the points in $X_{\mathcal{B}}$ (and similarly for the points in Y); we thus sample a latent point $p_{\mathcal{A}}$ among the points in object \mathcal{A} and another latent point $p_{\mathcal{B}}$ among the points in object \mathcal{B} , and we condition the decoder on both sampled latent points (see the next

section for details).

The spatially-grounded latent distribution has a nice intuitive interpretation: the sampled latent points z indicates which of the demonstration modes is selected. For example, for the task of hanging a mug on one of N racks, the latent point in X_B empirically tends to indicate which of the racks will be used for hanging (see a visualization of this latent space in Figure 3.1. Thus, this spatially grounded latent space is able to interpretably select a mode from the demonstration distribution by conditioning on a learned distribution over points from X_A and X_B . Because this distribution is discrete, we avoid the “smoothing” issue normally found with cVAEs mentioned in the previous paragraph. In order to back-propagate through these discrete latent distributions, we perform differentiable sampling with Gumbel Softmax [8].

Architecturally, we design the prior $p_\phi(z | X)$ as well as the encoder $q_\psi(z | Y)$ to predict a value over each point in the point clouds X and Y respectively, as shown in Figure 3.2. We then normalize these values to obtain a discrete probability distribution over the points in the point clouds.

Latent conditioning:

As mentioned previously, we use a modified version of TAX-Pose [14] for the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$. However, we must modify the TAX-Pose model to condition on the latent points $z = (p_A, p_B)$. One option for doing so would be to represent z as the discrete value selected from the N -dimensional multinomial distribution over points. However, the discrete value would correspond to the index of the selected point, which has no particular meaning to the network, which uses a permutation-invariant point cloud-based architecture [24].

Another option is to use the 3D location of the points (p_A, p_B) directly; however, we would like the model to be translation-equivariant, so we would like to avoid explicitly inputting the 3D locations of the latent points. Another alternative is to concatenate to each input a binary indicator of whether that point was selected as the latent value z . However, such a sparse input is very difficult for the decoder to process.

Instead, we create a dense and SE(3)-invariant representation of z , as follows: For each point $p \in X_A$, we compute a feature $\mathcal{F}_p := \|p - p_A\|_2$, and likewise for X_B , which indicates the distance between each point p and the latent point p_A or p_B . This feature is both translationally- and rotationally-invariant, in that if we translate or rotate the entire point cloud X_A , then \mathcal{F}_p will remain the same (see Appendix 5.5 for the proof). We concatenate \mathcal{F}_p to the input of each point in X in order to condition the decoder on the latent variable z . A visualization of \mathcal{F}_p is depicted in Figure 3.2.

The TAX-Pose [14] model that we use for the decoder achieves translation invariance by mean-centering the points in X_A and X_B . However, for larger multi-object scenes such as those with multiple racks explored in this work, mean-centering doesn’t lead to a consistent point cloud ob-

ervation. To address this, for the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$, we center the points in X_A (and X_B respectively) on the sampled points p_A and p_B . Empirically, we find that these architectural choices lead to significantly improved performance for relational object placement tasks.

3.3.3 Generalizing to Arbitrary Configurations for Inference

During cVAE training, we also include a regularization loss $KL(q_\psi(z | X, Y) || p_\phi(z | X))$ as shown in the second term of Equation 3.3. This regularization term allows us to sample from the prior $p_\phi(z | X)$ at inference time and generate a latent z that is similar to the distribution of encoded latent z values $q_\psi(z | X, Y)$. Normally, the prior distribution $p_\phi(z | X)$ is set to be a Normal Gaussian distribution; in our case, the latent space z is discrete, so we could set the prior to a uniform distribution.

However, a uniform prior would not meaningfully encode the modes of the demonstration space. As shown in Figure 3.1, the prior $p_\phi(z | X)$ should map the input to a distribution that encodes the modes of the demonstrations. For example, in a scene with multiple mug racks, the distribution over z would encode the different options of selecting a mug rack. A uniform prior z over all points in the scene would encourage an equal probability mass on points that are not mug racks at all, and thus the latent z would not meaningfully encode the demonstration modes. We observe that using a uniform prior leads to poor generalization in practice. Thus, we train the encoder and decoder without an explicit regularization loss.

Instead, we train a learned prior $p_\phi(z | X)$ to map from the observation X to points in the latent space over z that encode the different modes of the demonstrations. We train the prior using only the regularization loss from Equation 3.3:

$$\mathcal{L} := KL(q_\psi(z | X, Y) || p_\phi(z | X)) \quad (3.4)$$

We apply a stop gradient on this loss so that we do not back-propagate through the encoder $q_\psi(z | X, Y)$, which we separately (and simultaneously) train using the reconstruction loss to map from a demonstration pose Y to a point in the latent space z to capture each demonstration mode. Because the learned prior $p_\phi(z | X)$ does not condition over the demonstration Y , it must marginalize over all demonstrations in the dataset. Thus, for the example of hanging a mug with multiple racks, the encoder $q_\psi(z | X, Y)$ will map from a demonstration Y to a distribution over the points on the rack that the mug is hung on in the demonstration Y , whereas the learned prior $p_\phi(z | X)$ will map from the input X to a distribution over *all* of the mug racks as shown in Figure 3.1. Further implementation details on the learned prior can be found in Appendix 5.1 on the website.

The inference procedure for our method can be found in Figure 3.2 (bottom). We first input the observation X into a learned prior $p_\phi(z | X)$ which outputs a distribution over the latent space z . We sample a z from this distribution, which we convert to a dense rotationally-invariant feature \mathcal{F}_p as described above. Finally, we input X and our representation of z into the modified




Figure 3.3: Visualization of a successful relative placement for hanging a mug with an arbitrary pose on one of many racks in a random 3 rack scene. The mug is initialized to an arbitrary pose in the air and must be grasped first by the robot.

TAX-Pose [14] decoder $p_{\theta}(\mathbf{T}_{AB} | X, z)$ to output the predicted cross-pose $\hat{\mathbf{T}}_{AB}$. We then send this predicted cross-pose to our robot system, which uses motion planning [22] to move the action object \mathcal{A} such that it is transformed by transformation $\hat{\mathbf{T}}_{AB}$. If the learned prior and decoder are trained well, then this should complete the relative placement task.

3.4 Experiments

We evaluate our method on the NDF [19] simulated mug placement tasks in addition to multi-modal variants of the original mug-hanging task. In all these tasks, the goal for the robot is to move an action object to a pose relative to the rest of the scene. The robot is given only 10 demonstrations, from which we extract a point cloud recording of the final configuration from each demonstration. The tasks are implemented in Pybullet [3], in which a simulated robot is placed on a table, surrounded by four depth cameras. We increase the difficulty of the original mug-hanging NDF task to require reasoning under multimodality: instead of one rack, the environment contains multiple racks in various configurations.

We implement our method for each relative placement task with two sequential cross-pose estimation steps: the cross-pose between the gripper and the action object (e.g. the mug) for grasping, followed by the cross-pose between the object and the rest of the scene (e.g. the mug racks) for placing. We use motion planning to move the gripper into the desired positions for both parts of the task; after the gripper rigidly grasps the mug, the mug-on-rack cross-pose defines the

desired target pose of the gripper. In all of our evaluations, the object (e.g. mug) is initialized to an arbitrary initial pose in space. A placement is marked as a success when the object is placed stably on one of the target placement locations.

3.5 Results

We present task success evaluations for a variety of tasks with object instances and configurations that are unseen during training. We first evaluate our model on the “1 Rack” mug hanging task from NDF [19], which requires no multimodality. In “1 Rack Place”, our model performs on par with TAX-Pose [14] (Table 3.1) and outperforms other baselines in placement. Details of DON [6] and NDF [19] can be found in prior work [19].

Next, we show how our formulation extends TAX-Pose [14] to multimodal mug-hanging tasks of increasing difficulty in Table 3.1. We train and evaluate the TAX-Pose model and our method on the “2 Racks” task, in which the task is to hang the mug on either of two racks that each have an arbitrary pose on a tabletop. We see a large improvement over the original TAX-Pose method in the 2-rack case in Table 3.1. We also test our model’s ability to generalize to more racks than the number seen during training: despite only seeing two racks during training, our model is able to retain strong performance for the “3 Racks” task. We observe that the TAX-Pose baseline tends to incorrectly predict mug placements in between the racks. Visually, we see that the learned prior $p_\phi(z | X)$ distribution will cover all of the racks available during inference, as shown in Figure 3.1.

	1 Rack Place	2 Racks Place	(Generalization) 3 Racks Place
DON [6]	0.45	-	-
NDF [19]	0.75	-	-
TAX-Pose [14]	0.84	0.07	0.14
TAX-PoseD (Ours)	0.82	0.57	0.59

Table 3.1: **Task success rates for the original NDF 1 rack mug hanging task and multimodal variants.** We evaluate on unseen mugs initialized to arbitrary configurations in space. *1 Rack:* Our model performs on par with TAX-Pose [14] in the unimodal placement task and outperforms the other baselines. *2 Racks:* We extend TAXPose to a multimodal task and train our model for random 2-rack scenes. We evaluate on 50 random 2-rack configurations. *3 Racks (Generalization):* We further evaluate the 2 rack model to test generalization to more racks than seen at training time. Three racks are placed in 50 random configurations. We report placement task success.




Figure 3.4: Samples from our model trained on various multimodal mug-on-rack environments cover a wide distribution of placements that correspond to task success. Top row- one sample, denoted by orange. Bottom row- multiple samples denoted by orange, red, and green.

3.5.1 Ablations

In prior work, VAE variants often represent the latent z as a global vector [7] sampled from a Normal distribution. Our method, in contrast, represents z as a discrete distribution over points in the point cloud, which we refer to below as “spatial z ” in Table 3.2. We provide ablations to our method for which we replace the encoder $q_\psi(z | X, Y)$ with one that outputs a continuous z vector like in prior work (“continuous z ”). We condition the decoder with such a z by following the architecture of the goal-conditioned TAX-Pose-GC [14]. We also provide ablations for regularizing the encoder $q_\psi(z | X, Y)$ to a Normal or Uniform prior. In Table 3.2, we show that the continuous z ablations fail to predict precise placements for both multimodal environments, and the Uniform prior ablation for discrete spatial z struggles to generalize to the “3 Racks” environment. In contrast, our method is able to account for this multimodality and generalization.

round-mode=places,round-precision=3						
	2 Racks			3 Racks		
	Place \uparrow	ϵ_R \downarrow	ϵ_t \downarrow	Place \uparrow	ϵ_R \downarrow	ϵ_t \downarrow
TAX-Pose [14]	0.07	103	0.45	0.14	71.8	0.33
TAX-PoseD ablations:						
- continuous z , Normal prior	0.33	70.2	0.28	0.18	66.2	0.31
- continuous z , learned prior	0.41	43.3	0.24	0.29	44.5	0.29
- spatial z , uniform prior	0.60	14.1	0.05	0.18	25.3	0.25
- Ours: spatial z, learned prior	0.57	13.2	0.05	0.59	16.0	0.06

Table 3.2: **Ablations.** We show ablations of our model for the multimodal 2-racks mug-hanging task and also for generalization to the unseen 3-racks task. We report placement success, and also rotation error ($^\circ$) and translation error (m) relative to the closest demonstration. Using a discrete spatial z distribution over points in the point cloud performs better than using a continuous z vector; using a learned prior performs better than using a non-learned prior (Normal or uniform), especially for generalization to “3 Racks.”

Chapter 4

Conclusions and Future Work

We presented TAX-PoseD, a data-driven model for relative placement tasks that accounts for multimodality in task specification by learning a latent distribution over points in the object point clouds. We show that our method efficiently learns distributional relative placement tasks with only 10 demonstrations. Our method extends the VAE literature to practically implement an interpretable, spatially-grounded latent space. We achieve high task success on the mug placing task with a varying number of racks, which requires multimodal reasoning, with the use of $SE(3)$ -invariant modules for reasoning about multimodal relative placement.

We would like to point to future directions of this work. Future work may tackle more diverse scenes and wider distributions of demonstrations. For example, more placement strategies can be demonstrated for the same objects. Also, there can be work in relaxing the requirement to segment the manipulated object from the scene, or even modeling relationships between multiple objects instead of a pair. We also hope to explore a variety of data augmentation techniques corresponding to the distribution of points in the point cloud. In the future, this formulation can also be extended to handle deformable objects.

Chapter 5

Appendix

5.1 Training Details

5.1.1 Loss functions

We train the encoder and decoder with a reconstruction loss; the decoder outputs a transformation $\hat{\mathbf{T}}_{AB}$ which we apply to the action object point cloud X_A to obtain a predicted point cloud \hat{Y} ; we then compare this predicted point cloud to the demonstration ground-truth point cloud Y to obtain the loss. We train these networks with similar losses as were used in TAX-Pose [14] losses, i.e. a weighted sum of a Point Displacement Loss (weight=1), a Direct Correspondence Loss (weight=0.1), and a Correspondence Consistency Loss (weight=1).

We also train the learned prior, $p_\phi(z | X)$, to match the distribution predicted by the encoder, $q_\psi(z | X, Y)$. In practice, while this can be implemented as the KL loss between the encoders' outputs $p_\phi(z | X)$ and $q_\psi(z | X, Y)$, empirically we find using the Jensen-Shannon Divergence (JSD) [11] to improve the learned prior performance. For our spatially grounded latent space defined as a categorical distribution over the points in the scene, the JSD is defined as:

$$JSD(q_\psi, p_\phi) = \frac{1}{2}KL\left(q_\psi \parallel \frac{q_\psi + p_\phi}{2}\right) + \frac{1}{2}KL\left(p_\phi \parallel \frac{q_\psi + p_\phi}{2}\right) \quad (5.1)$$

Because our network outputs separate distributions over the action object points X_A and the anchor object points X_B , we take the sum of the JSD over the action object distribution and the JSD over the anchor object distribution. We use a learning rate of $\alpha = 1e-4$ and gradient clipping of $1e-3$ for all models.

5.1.2 Network Architecture for Encoder

The encoder $q_\psi(z | Y)$ is implemented as a single-scale PointNet++ [16] directly encoding the points of the action and anchor objects in Y into per-point features, which are then normalized into probability distributions with a softmax. In more detail: we separately normalize the action and anchor object points to create two distributions: $q_\psi(z_A | Y)$ and $q_\psi(z_B | Y)$, from which we sample the spatially-grounded latent points p_A and p_B . We use a single-scale PointNet++ for encoding demonstrations with $q_\psi(z | Y)$ to capture the local relationship between the action object and the single demonstration mode in the anchor object of Y .

5.1.3 Network Architecture for Learned Prior

The learned prior $p_\phi(z | X)$ is implemented as two DGCNN [24] networks that each separately encode the action and anchor point clouds, followed by a Transformer [23] that performs a cross-object comparison across the action and anchor objects. The output of the transformer are per-point features, which are then each passed through a point-wise MLP and normalized into probability distributions (separately for the action and anchor points) with a softmax.

5.1.4 Network Architecture for Decoder

As previously described, we use a modified version of TAX-Pose as the decoder $p_\theta(\mathbf{T}_{AB} | X, z)$. We modify TAXPose to condition on the latent z as described in Sec.???. Another minor modification that we make is as follows: within TAX-Pose’s cross-correspondence estimator, two point-wise MLPs are used to map cross-object point embeddings into correspondence residuals and per-point weights. Empirically, we find that using a single point-wise MLP to jointly predict the correspondence residuals and per-point weights improves the performance of our method.

5.1.5 Dataset generation

To generalize to multiple racks in our mug-hanging experiments, we create a training set with multi-rack demonstrations. To do so, we modify the environment used for 1-rack mug hanging to generate multi-rack demonstrations. We create a distractor rack by copying the points from the rack in the 1-rack demonstration with a random SE(3) transformation applied; we then concatenate this distractor rack to the scene point cloud. We ensure that the scene is physically-plausible by performing axis-aligned rectangular prism collision checking between the original demonstration and the distractor rack. When training the encoder $q_\psi(z | Y)$, we do not apply any transformations to the points in the demonstration point cloud Y . However, when we train the learned prior $p_\phi(z | X)$, we rotate both the original rack and the distractor rack to enable the learned prior to be robust to new rack poses that might be encountered at test-time.

To generate an arbitrary initial configuration X from a demonstration Y , we uniformly sample two random $SE(3)$ transformations and apply them to the points corresponding to objects \mathcal{A} and \mathcal{B} , respectively, to produce a point cloud for each object in an arbitrary pose: $X_{\mathcal{A}}$ and $X_{\mathcal{B}}$. These two point clouds together form the point cloud $X = X_{\mathcal{A}} \cup X_{\mathcal{B}}$. We found it helpful to center the $X_{\mathcal{A}}$ and $X_{\mathcal{B}}$ input to the decoder by the sampled point from the latent distribution $p_{\mathcal{A}} \sim p(z_{\mathcal{A}} | X_{\mathcal{A}})$ and $p_{\mathcal{B}} \sim p(z_{\mathcal{B}} | X_{\mathcal{B}})$, respectively. The inputs to the encoders $q_{\psi}(z | X, Y)$ and $p(z | Y)$ are mean centered.

5.1.6 Synthetic occlusions

We apply planar and ball synthetic occlusions to the action object as data augmentations simulating partially-occluded scenes. Specifically, the planar synthetic occlusion chooses a random 3D plane to be between a point in the point cloud and the point cloud center, then removes all points on one side of the plane. Also, the ball synthetic occlusion chooses a random point in the point cloud then removes all points within a certain radius of that chosen point. These occlusions can roughly simulate occlusions from a limited number of camera perspectives and various object interactions. In addition, when downsampling the demonstration point cloud, we apply a random downsampling 50% of the time and apply furthest point downsampling 50% of the time to capture more randomness in the distribution of points.

5.1.7 Motion Planning

The model’s cross-pose predictions guide an OMPL motion planner for the robot to produce the executed trajectory, which takes into account obstacles and joint limits.

5.2 Task Details

Details about the different environments that we use for evaluation can be found in Table 5.3.

Environment	Description
Mug-on-Rack (1 rack) / 1 Rack	Original NDF task of hanging a mug on a mug rack. The rack is at a fixed location on the table and there are a variety of unseen mugs at test time.
2 Racks	The two mug racks are randomly rotated with respect to the z axis and randomly translated across the table surface. We evaluate on a fixed test set of 50 rack configurations.
3 Racks	The three mug racks are randomly rotated with respect to the z axis and randomly translated across the table surface. We evaluate on a fixed test set of 50 rack configurations.

5.3 Additional experiments

Here, we show additional experiments, applying our method on a variety of other tasks.

Environment	Description
Mug Grasping	A two-fingered gripper is to be positioned to a pre-grasp pose on the rim of a variety of mugs. After the pose is predicted, the gripper closes to grasp the mug.
Bowl Grasping	Similarly to Mug Grasping, the goal is to predict the pre-grasp pose on the rim of a bowl. This task has inherent rotational symmetry.
Bottle Grasping	Similarly to Mug Grasping, the goal is to predict the pre-grasp pose on the neck of the bottle. This task also has inherent rotational symmetry.
Book-on-Bookshelf	A rectangular book is to be placed on a partially-filled bookshelf, which has multiple gaps between already placed books.
Multi-Peg Rack	A rack has multiple randomly generated pegs, and the robot may hang the mug on any one of the pegs.
Multi-Peg Multi-Racks	Multi-Peg Rack is extended to the multi-rack setting, and the robot may hang the mug on any one of the pegs, on any one of the racks.

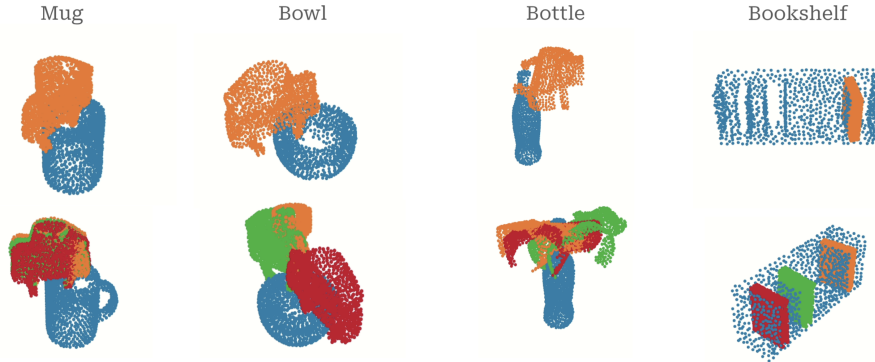


Figure 5.1: We show that TAX-PoseD learns to cover a wide variety of ways to complete a task. Blue- an anchor object depicting the scene. Red, Green, and Orange- manipulated objects at various poses sampled from TAX-PoseD.

5.4 Interpretability of the the learned latent distribution

Our method learns disentangled representations of the distributional demonstration spaces for a task. When we view encoder $q_\psi(z | X, Y)$ for one mode of demonstration, we empirically observe that its representation is spatially distinct from the representation for another mode of demonstration (Figure 5.2). In addition, the encoder $p_\phi(z | X)$ that maps across all modes of demonstrations in $q_\psi(z | X, Y)$ interpretably represents task symmetries across distinct racks in the mug hanging task (Figure 3.1).

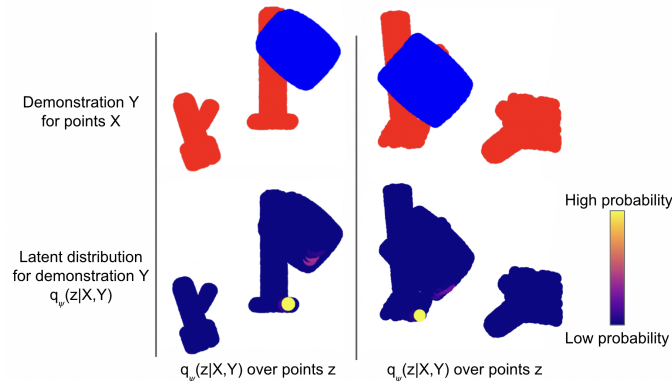


Figure 5.2: Our method learns disentangled representations for relative placement tasks. Empirically, the latent distribution assigns high probability to the rack upon which the mug is hung in the demonstration.




Figure 5.3: We show visualizations of the learned latent space across the Multi-Peg Rack, Multi-Peg Multi-Rack, and Bowl Grasping tasks. Empirically, the high probability points tend to correspond spatially to placements of the manipulation object.

5.5 Proof for rotational invariance for the featurization

Here, we prove that our latent conditioning described in ?? is $SE(3)$ -invariant.

Let there be an arbitrary transformation $T \in SE(3)$ that rigidly transforms an object in an arbitrary pose X to the same object in a demonstration pose Y . Transformation T can be rewritten as a rotation R and translation component t .

$$X := TY$$

Let p_x be a point in point cloud X where the feature is being computed, and let z_x be the point sampled from the latent distribution. Define p_y and z_y similarly for point cloud Y . Then, our method’s featurization is invariant to $SE(3)$ transformations with respect to the demonstration Y :

$$\begin{aligned}
 \mathcal{F} &:= \|p_x - z_x\|_2 = \|T^{-1}p_y - T^{-1}z_y\|_2 \\
 &= \|R^T(p_y - t - z_y + t)\|_2 \\
 &= \|p_y - z_y\|_2
 \end{aligned} \tag{5.2}$$

Note that without the L2 norm, the featurization would not be rotationally-invariant.

Bibliography

- [1] A. Allshire, R. Martín-Martín, C. Lin, S. Manuel, S. Savarese, and A. Garg. Laser: Learning a latent action space for efficient reinforcement learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6650–6656. IEEE, 2021.
- [2] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [3] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- [4] B. Eisner, H. Zhang, and D. Held. Flowbot3d: Learning 3d articulation flow to manipulate articulated objects. *arXiv preprint arXiv:2205.04382*, 2022.
- [5] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit behavioral cloning. In *Conference on Robot Learning*, pages 158–168. PMLR, 2022.
- [6] P. R. Florence, L. Manuelli, and R. Tedrake. Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. *arXiv preprint arXiv:1806.08756*, 2018.
- [7] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2016.
- [8] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [9] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [10] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [11] J. Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.
- [12] 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*, pages 6813–6823, 2021.

- [13] A. Mousavian, C. Eppner, and D. Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2901–2910, 2019.
- [14] 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*, pages 1783–1792. PMLR, 2023.
- [15] T. Pandeva and M. Schubert. Mmgan: Generative adversarial networks for multi-modal distributions. *arXiv preprint arXiv:1911.06663*, 2019.
- [16] 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*, 30, 2017.
- [17] 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)*, pages 7278–7285. IEEE, 2020.
- [18] A. Simeonov, Y. Du, Y.-C. Lin, A. R. Garcia, L. P. Kaelbling, T. Lozano-Pérez, and P. Agrawal. Se (3)-equivariant relational rearrangement with neural descriptor fields. In *Conference on Robot Learning*, pages 835–846. PMLR, 2023.
- [19] 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)*, pages 6394–6400. IEEE, 2022.
- [20] A. Simeonov, A. Goyal, L. Manuelli, L. Yen-Chen, A. Sarmiento, A. Rodriguez, P. Agrawal, and D. Fox. Shelving, stacking, hanging: Relational pose diffusion for multi-modal rearrangement. *arXiv preprint arXiv:2307.04751*, 2023.
- [21] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015.
- [22] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. <https://ompl.kavrakilab.org>.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [24] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019.
- [25] F. Wirnshofer, P. S. Schmitt, G. von Wichert, and W. Burgard. Controlling contact-rich manipulation under partial observability. In *Robotics: Science and Systems*, 2020.
- [26] J. Wu, X. Sun, A. Zeng, S. Song, J. Lee, S. Rusinkiewicz, and T. Funkhouser. Spatial action maps for mobile manipulation. *arXiv preprint arXiv:2004.09141*, 2020.
- [27] R. Wu, Y. Zhao, K. Mo, Z. Guo, Y. Wang, T. Wu, Q. Fan, X. Chen, L. Guibas, and H. Dong. Vatmart: Learning visual action trajectory proposals for manipulating 3d articulated objects. *arXiv preprint arXiv:2106.14440*, 2021.

- [28] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021.