## One-shot Imitation Learning via Interaction Warping



## Motivation

- Manipulation policies should generalize over many object instances in pick and place tasks.

- We aim to learn from a single demonstration by using a shape warping prior.

grasp demo

placement demo

Coherent Point Drift (CPD, Myronenko et al., 2009)


Initialization


Iteration 10


Iteration 20

## Interaction Warping (IW)



Method Overview

1. Learn a latent space of object shapes for each class.

We use 10 example objects from ShapeNet for each object class. We pick a canonical object and fit a PCA to a
dataset of CPD warps.

2. Register interaction points from a demonstration onto the canonical object.

We automatically extract interaction points as contact and nearby points from a demonstration. We attach them to the warped canonical object

3. Warp interaction points to novel object instances. Solve for pick or place pose by matching the points.


## IW Contributions

1. Learning of a low-dimensional latent space of meshes.
2. Joint inference of object shape and pose using gradient descent with many random restarts
3. Transfer from a single demonstration by warping interaction points attached to object meshes.

## Results

- Place success rates in three simulated tasks. Baselines: Relational Neural Descriptor Fields (Simeonov et al. 2022) and TAX-Pose (Pan et al., 2022).

|  | \# | \# Train. <br> Method | Mug on Tree <br> Demo | Beswh on Mug | Bottle in Container |  |  |  |
| :--- | :--- | :--- | :--- | :---: | :---: | :---: | :---: | :---: |
| Upright | Arbitrary | Upright | Arbitrary | Upright | Arbitrary |  |  |  |
| R-NDF [17] | 1 | 200 | 60.0 | 51.0 | 69.0 | 68.0 | 19.0 | 8.0 |
| TAXPose [2] | 1 | 200 | 61.0 | 41.0 | 16.0 | 9.0 | 4.0 | 1.0 |
| IW (Ours) | 1 | 10 | $\mathbf{8 6 . 0}$ | $\mathbf{8 3 . 0}$ | $\mathbf{8 2 . 0}$ | $\mathbf{8 4 . 0}$ | $\mathbf{6 2 . 0}$ | $\mathbf{6 0 . 0}$ |

Table 1: Success rates of predicted target poses of objects in simulation. Upright and Arbitrary refer to the starting pose of the manipulated object. Measured over 100 trials with unseen object pairs.

- Real-world experiments with a wide range of objects.
Mug on tree Bowl on mug Bottle in box


Table 2: Success rates of real-world pick-and-place experiments with a single demonstration. The manipulated object (e of a mug) starts in an arbitrary pose (we use a stand to get a range of poses) and the target object (e.g. a mug-tree) starts in an arbitrary upright pose.

- Mesh and grasp prediction in the wild - Detic (Zhou et al., 2022) + SAM (Kirillov et al., 2023) + IW.


