One-shot Imitation Learning via Interaction Warping

Ondrej Biza[1], Skye Thompson[2], Kishore Reddy Pagidi[1], Abhinav Kumar[1], Elise van der Pol[3,*], Robin Walters[1,*], Thomas Kipf[4,*], Jan-Willem van de Meent[5,†], Lawson L.S. Wong[1,†], Robert Platt[1,†]

[1] Northeastern University [2] Brown University [3] Microsoft Research [4] Google DeepMind [5] University of Amsterdam [*] Equal contribution [†] Equal advising

Motivation

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

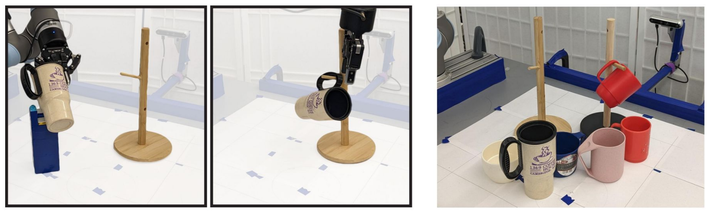

- We aim to **learn from a single demonstration** by using a shape warping prior.

grasp demo

placement demo

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

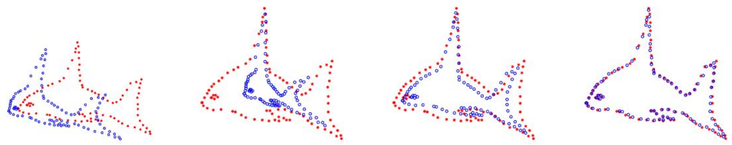

Initialization Iteration 10 Iteration 20 Iteration 40

- Warps source point cloud to match the target.
- Each point can move independently.
- Nearby points are regularized to move coherently.
- Implemented as a Gaussian mixture model with a frequency basis regularization.

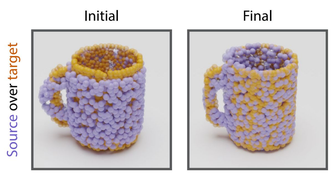

Interaction Warping (IW)

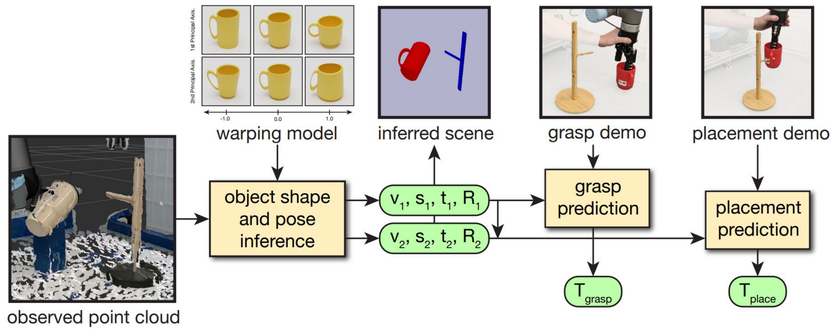

Method Overview

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.

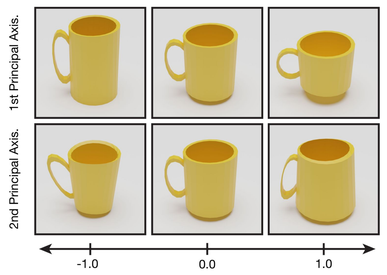

2. **Register interaction points** from a demonstration onto the canonical object.

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

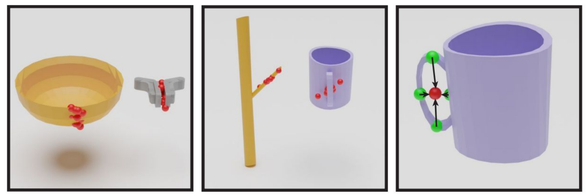

3. **Warp interaction points** to novel object instances. Solve for pick or place pose by matching the points.

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).

Method	# Demo	# Train. Meshes	Mug on Tree Upright	Mug on Tree Arbitrary	Bowl on Mug Upright	Bowl on Mug Arbitrary	Bottle in Container Upright	Bottle in Container Arbitrary
R-NDF [17]	1	200	60.0	51.0	69.0	68.0	19.0	8.0
TAX-Pose [2]	1	200	61.0	41.0	16.0	9.0	4.0	1.0
IW (Ours)	1	10	**86.0**	**83.0**	**82.0**	**84.0**	**62.0**	**60.0**

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

Method	Mug on Tree Pick	Mug on Tree Pick&Place	Bowl on Mug Pick	Bowl on Mug Pick&Place	Bottle in Container Pick	Bottle in Container Pick&Place	Mean Pick	Mean Pick&Place
NDF[1] [5]	93.3	26.7	75.0	33.3	20.0	6.7	62.8	22.2
R-NDF [17]	64.0	12.0	37.5	37.5	26.7	20.0	42.7	23.2
IW (Ours)	**96.0**	**92.0**	**87.5**	**83.3**	**86.7**	**83.3**	**90.1**	**86.2**

Table 2: Success rates of real-world pick-and-place experiments with a single demonstration. The manipulated object (e.g. 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.