


We can learn to predict object affordances from synthetic data generated with domain randomization



Code + Video + PDF


Learning to Segment Object Affordances on Synthetic Data for Task-oriented Robotic Handovers

Albert Christensen, Daniel Lehotsky, Marius Willemose Jørgensen and Dimitrios Chrysostomou

Introduction

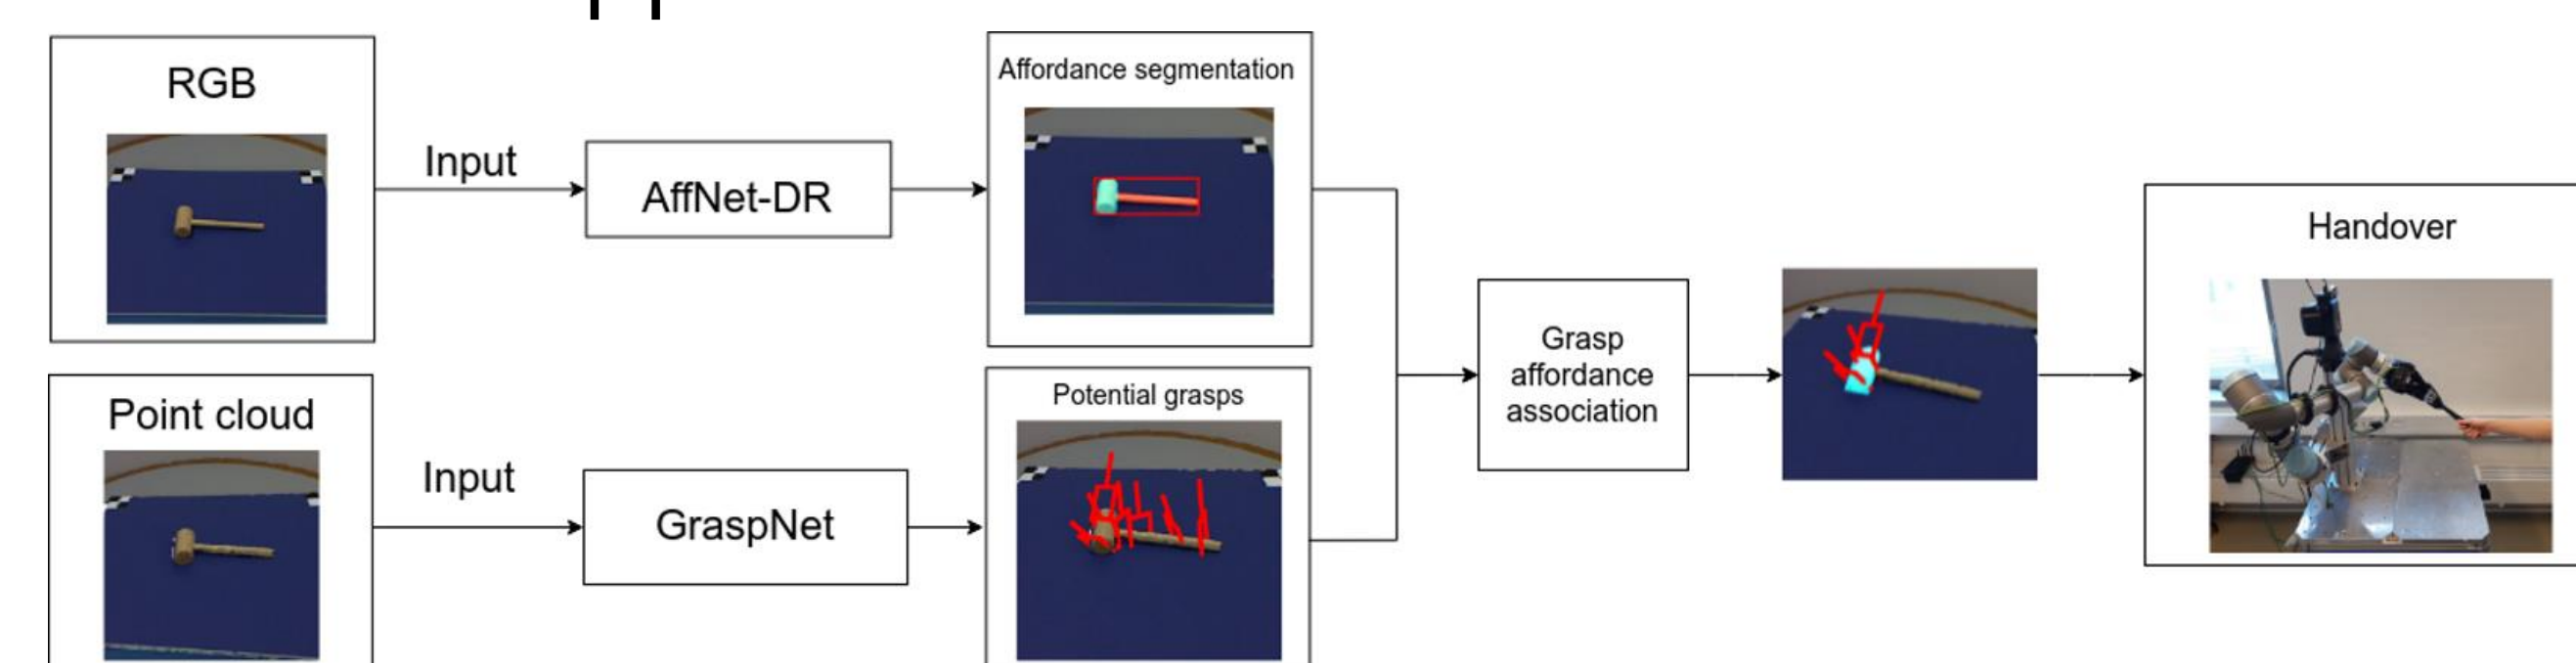
- Appropriate grasps for robot-to-human handovers can be generated by predicting object affordances.
- Predicting object affordances requires expensive pixel-level annotated datasets.
- We use domain randomization to overcome the sim2real gap issue when training on cheap synthetic data.

Method




A synthetic dataset with domain randomization was generated to overcome the sim2real gap.

Affordance prediction combined with a grasp sampler can find task-oriented grasps for a robot-to-human handover application.



Results

Our method generalizes to the real-world UMD dataset.



We outperform previous methods that use synthetic data

	Real-world data		Synthetic data	
	AffNet [1]	Baseline	AffNet-DA [2]	Our method
Grasp	0.731	0.482	0.473	0.611
Cut	0.762	0.575	0.599	0.604
Scoop	0.793	0.647	0.332	0.639
Contain	0.833	0.859	0.83	0.710
Pound	0.836	0.655	0.224	0.804
Support	0.821	0.519	0.541	0.578
W-grasp	0.814	0.848	0.821	0.785
Average	0.799	0.655	0.546	0.676

Real world robotic handover application results

Object	Affordance	Success rates	
		Baseline	Our method
Hammer	Grasp	40 %	90 %
	Pound	0 %	90 %
Spoon	Grasp	50 %	90 %
	Scoop	20 %	70 %
Ladle	Grasp	80 %	80 %
	Contain	90 %	80 %
Turner	Grasp	90 %	70 %
	Contain	70 %	90 %
Average success rate		55 %	82.5 %