FLEX: Full-Body Grasping Without Full-Body Grasps



Purva Tendulkar



Dídac Surís



Carl Vondrick

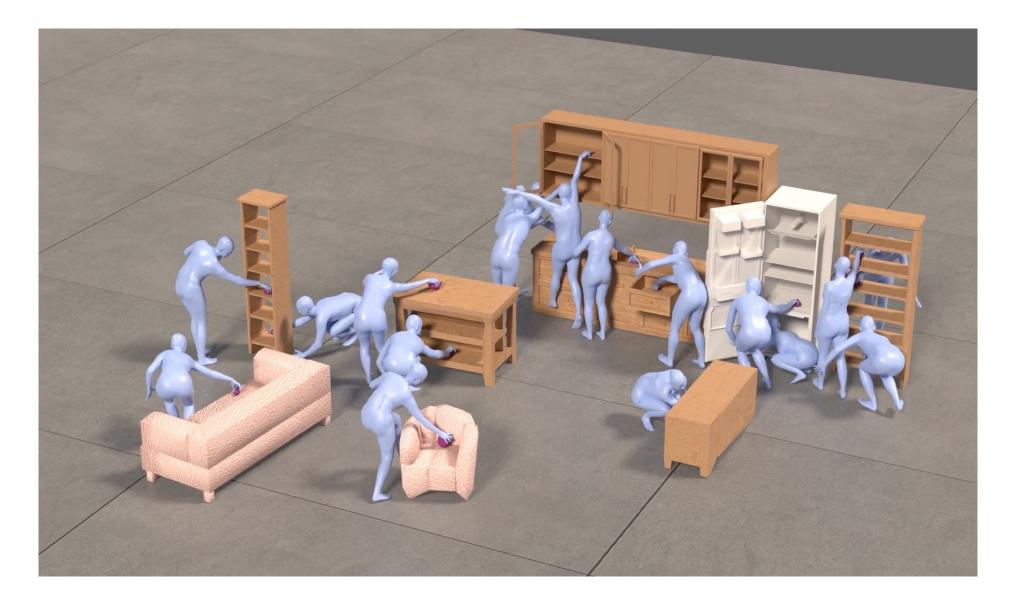
CVPR 2023

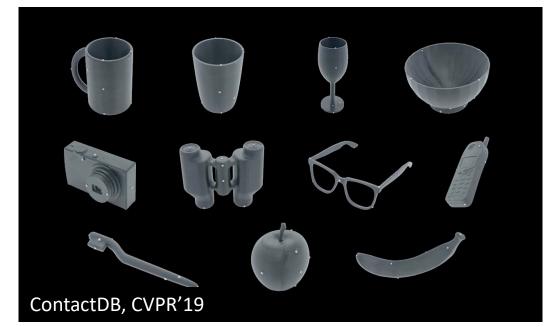
flex.cs.columbia.edu

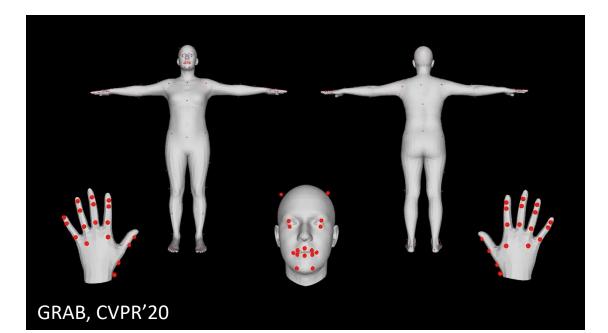
Poster: THU-PM-055

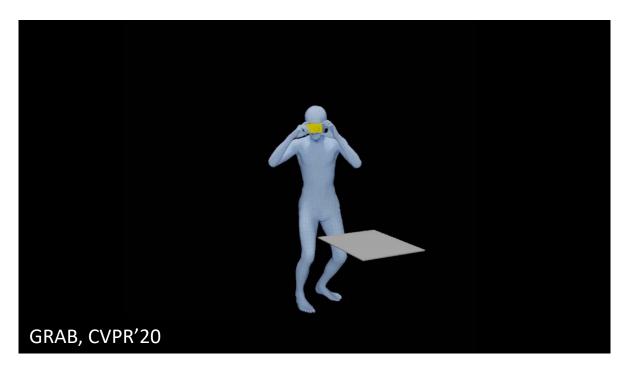
COLUMBIA UNIVERSITY

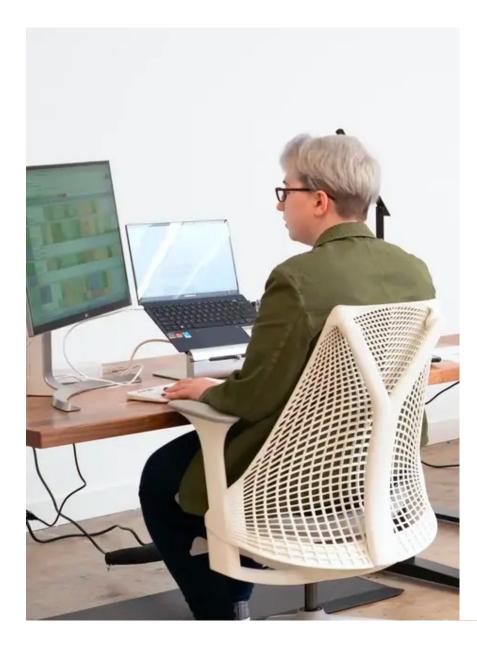
Task: Full-Body Grasping in Scenes



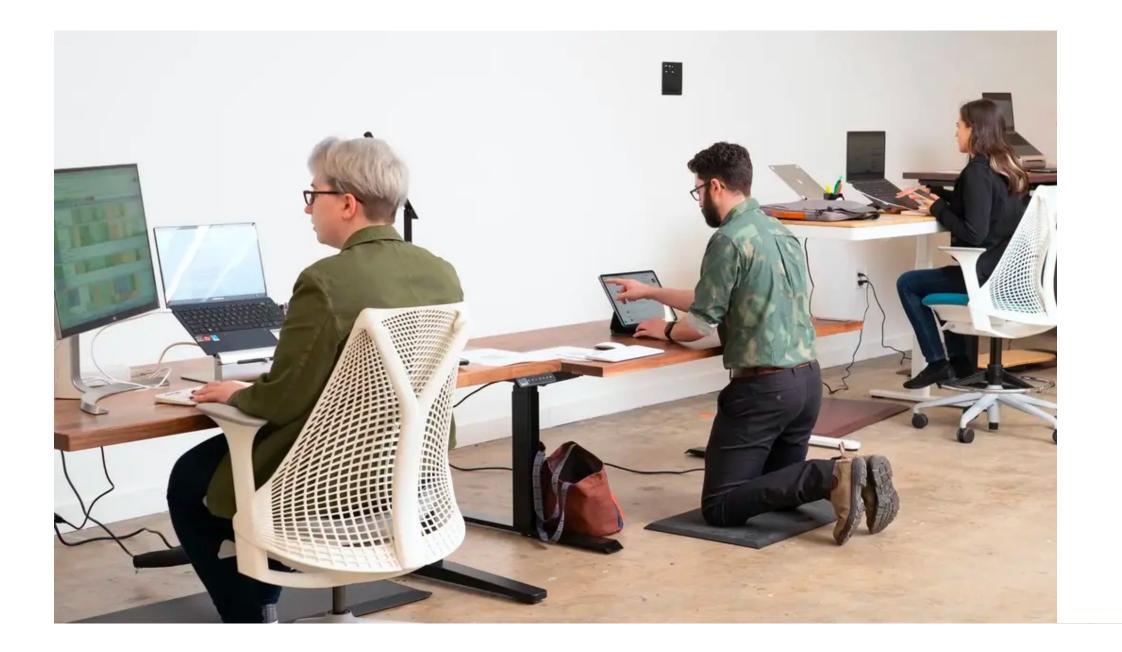




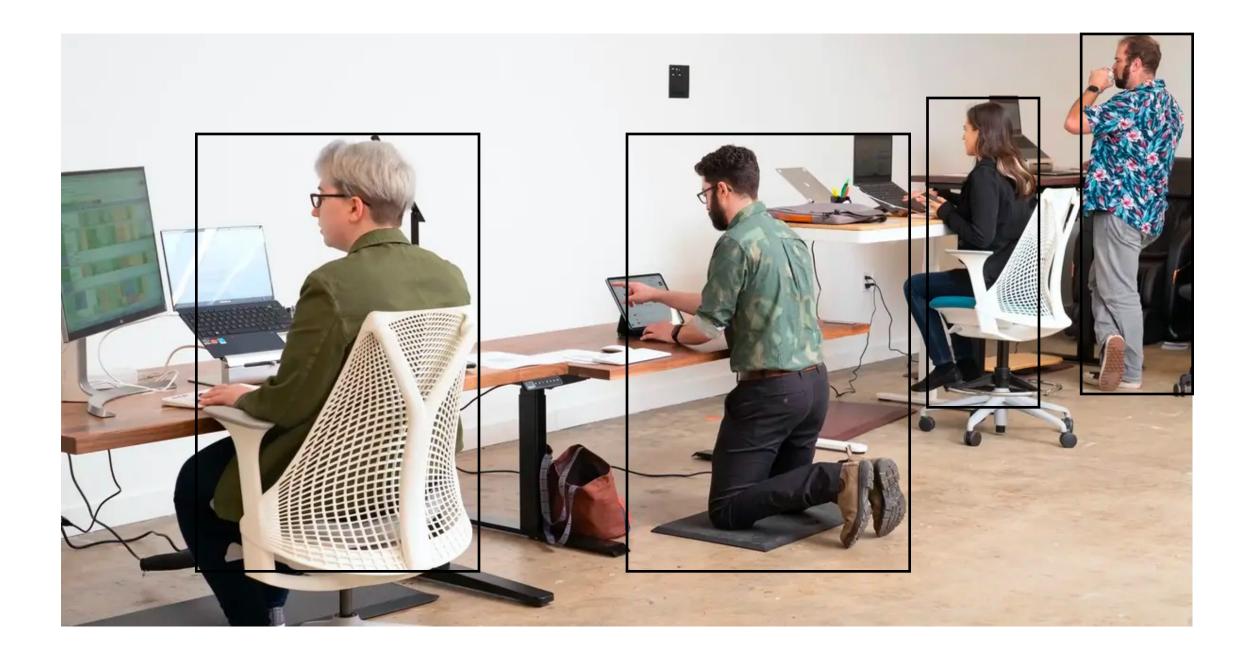


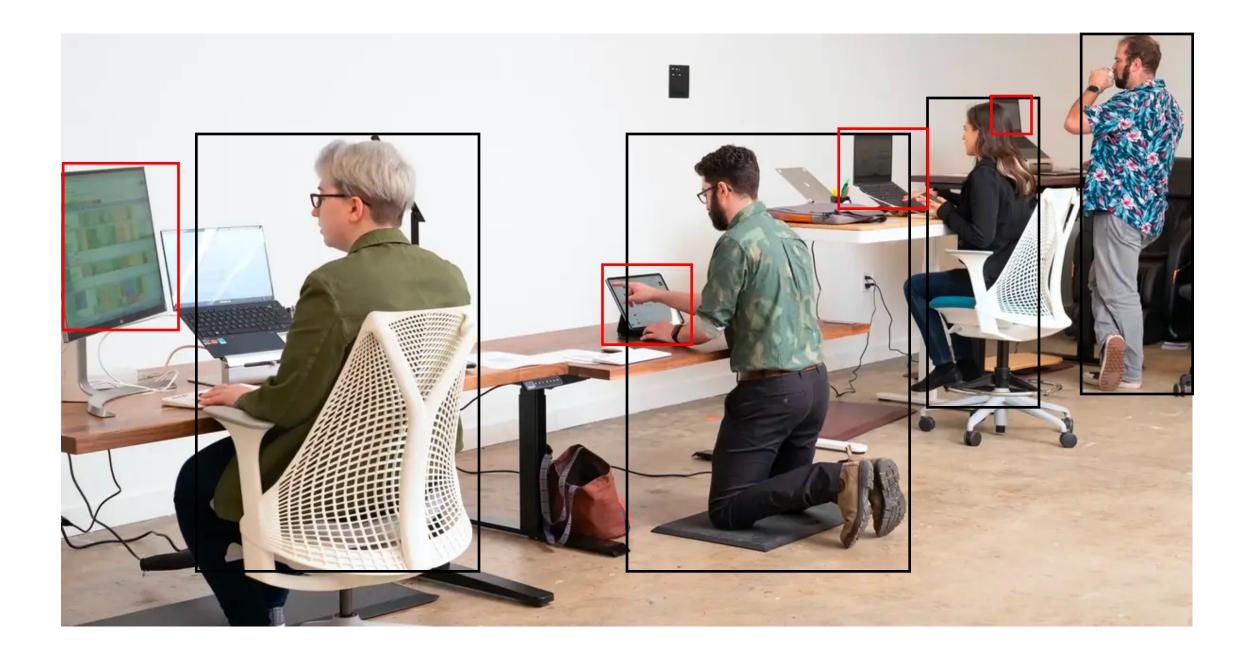




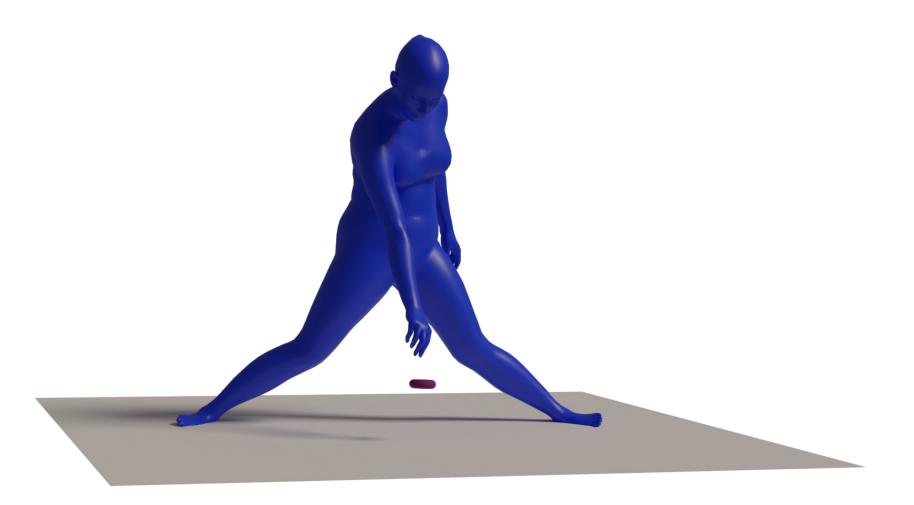




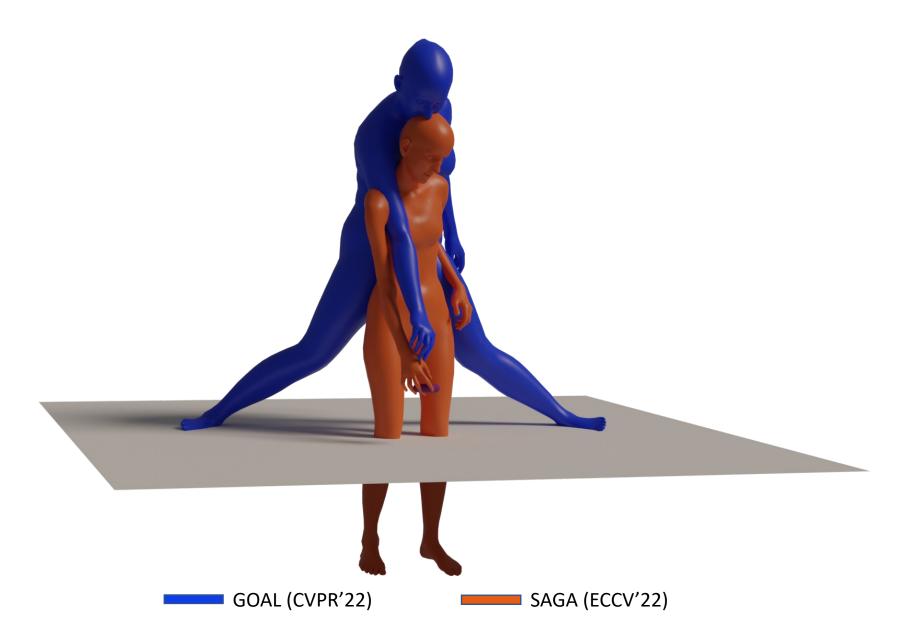




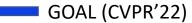




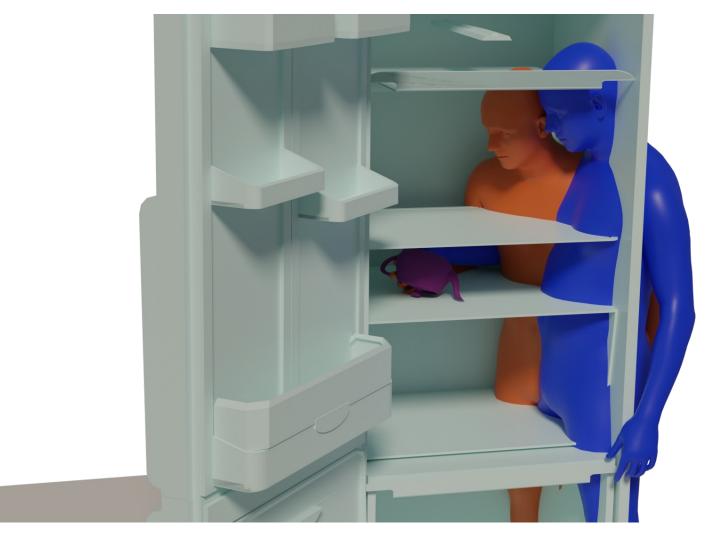






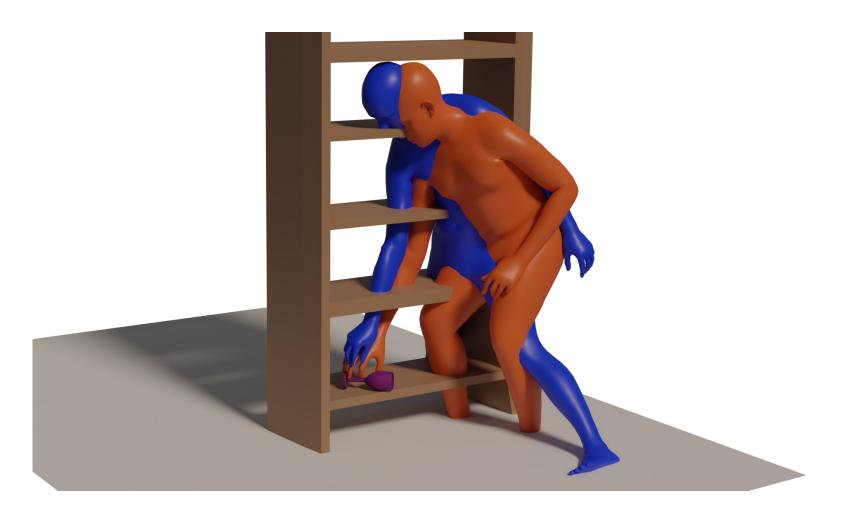


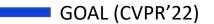














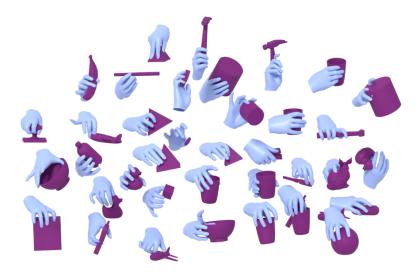
Do we *really* need full-body grasping data to solve full-body grasping?

MPI Limits





Large-scale MoCap Datasets

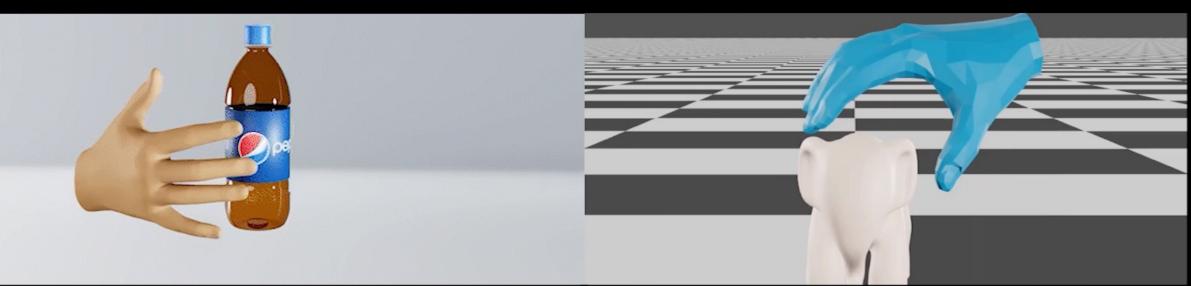






GRAB CVPR '20

GraspTTA ICCV '21



Grasp'D ECCV '22

TOCH ECCV '22



Large-scale MoCap Datasets



Hand Grasping Datasets & Methods

Prior over body poses



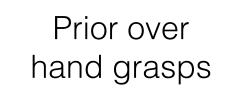
Large-scale MoCap Datasets



Hand Grasping Datasets & Methods Prior over hand grasps

Prior over body poses





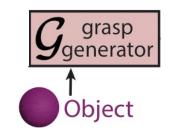


Large-scale MoCap Datasets

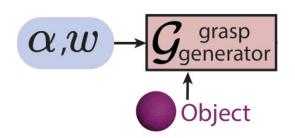


Hand Grasping Datasets & Methods

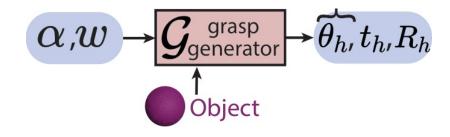




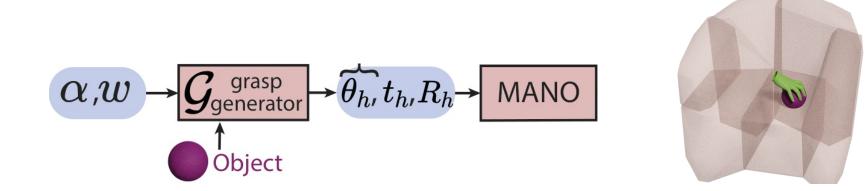


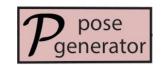


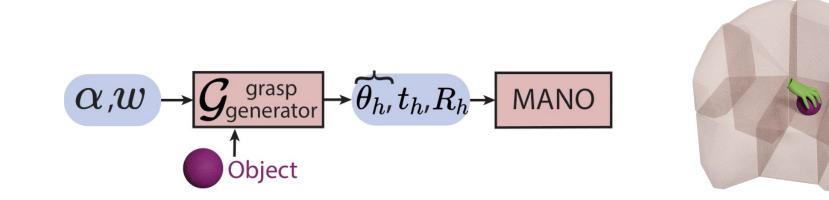


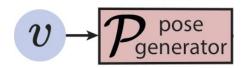


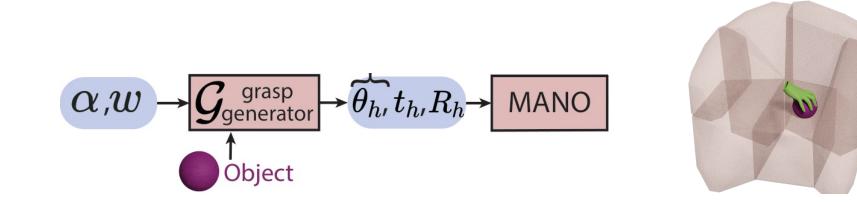


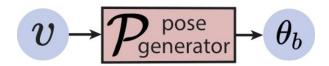


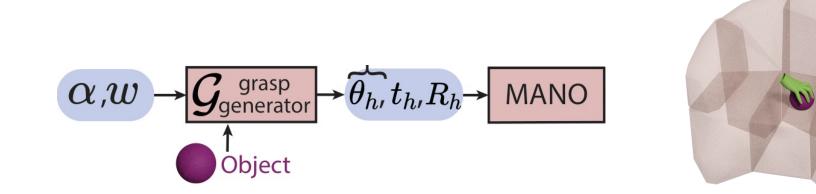


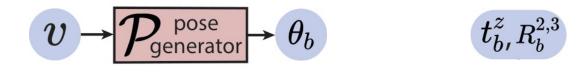


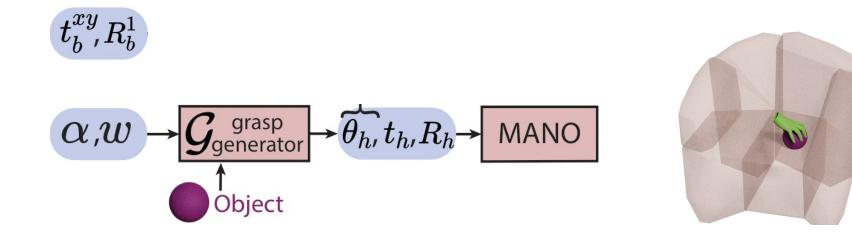


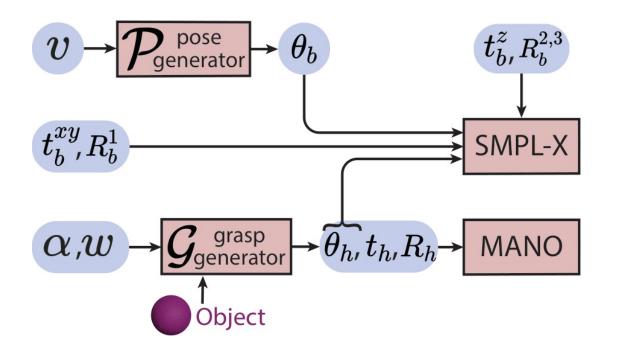




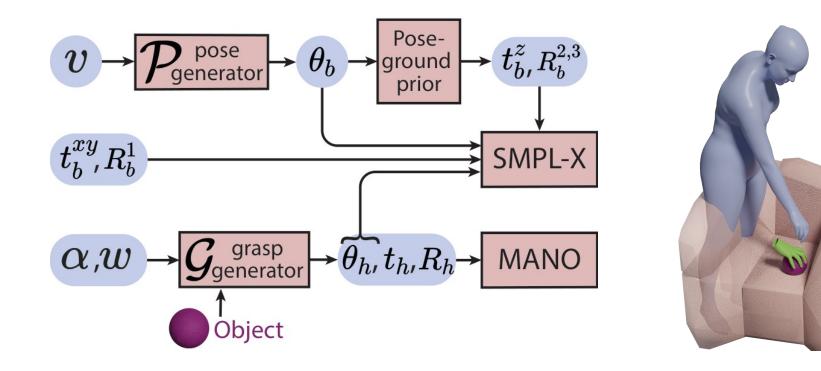


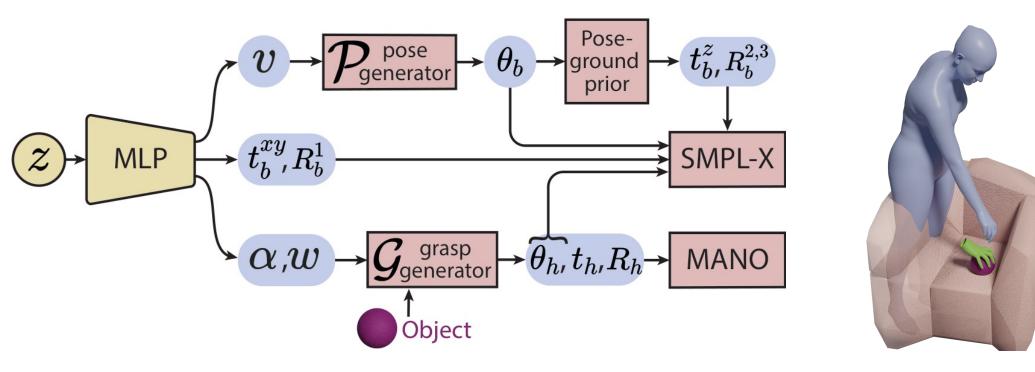


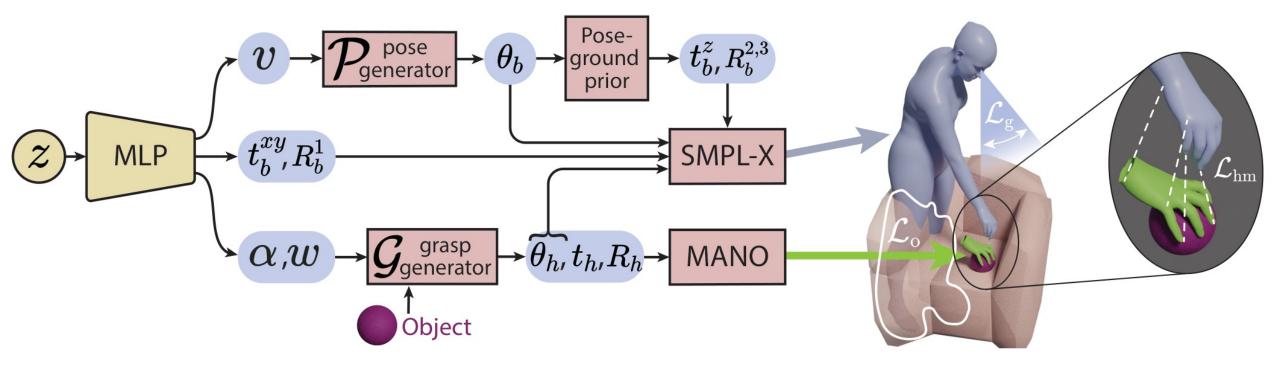


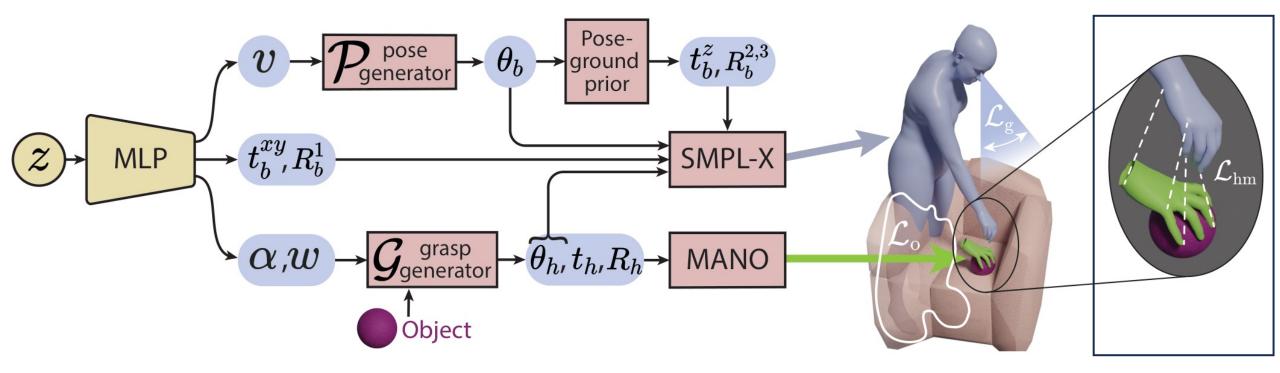








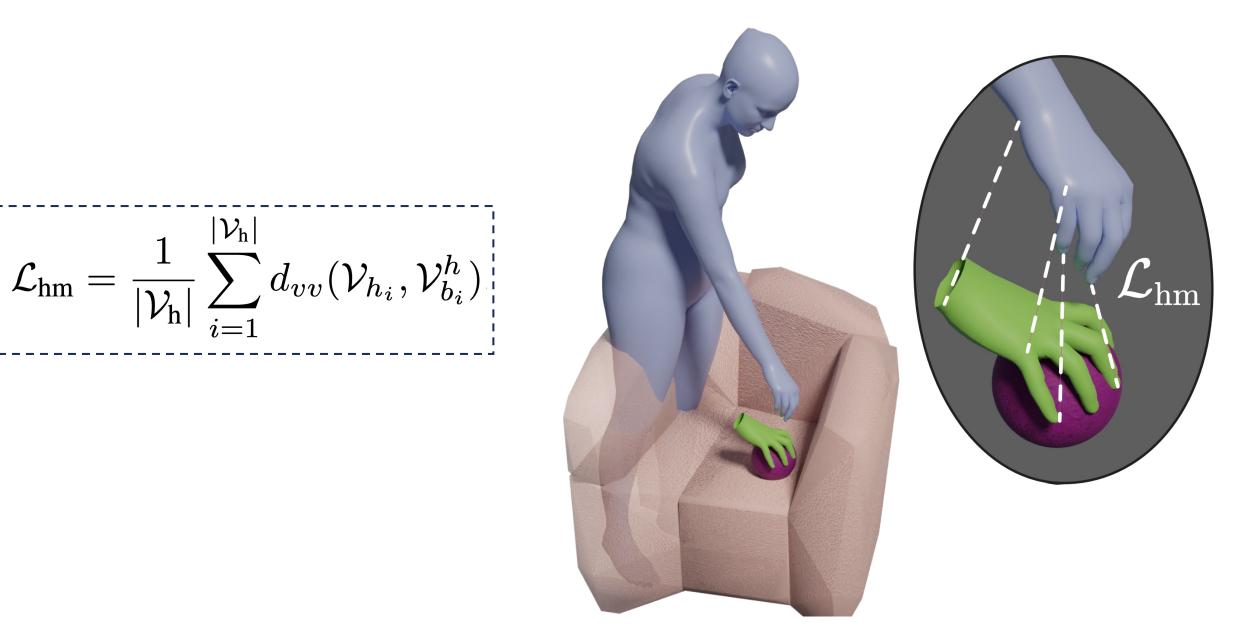


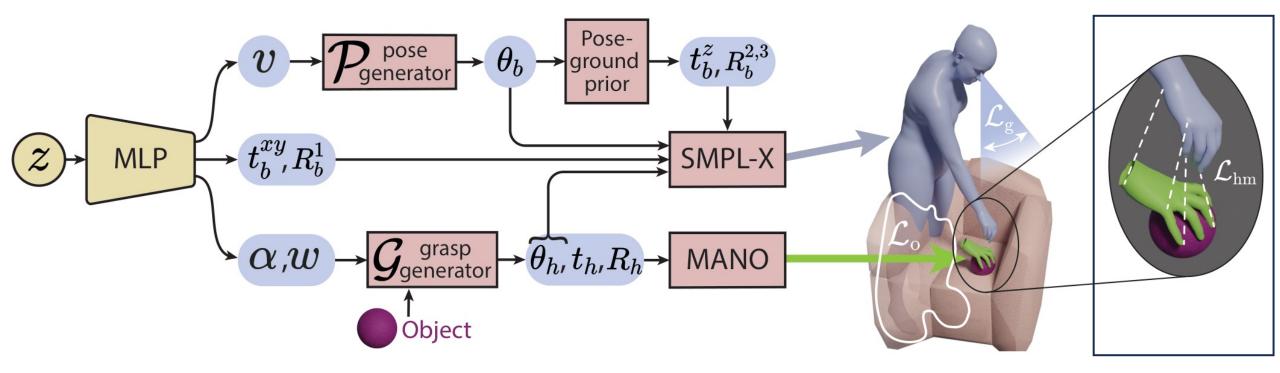


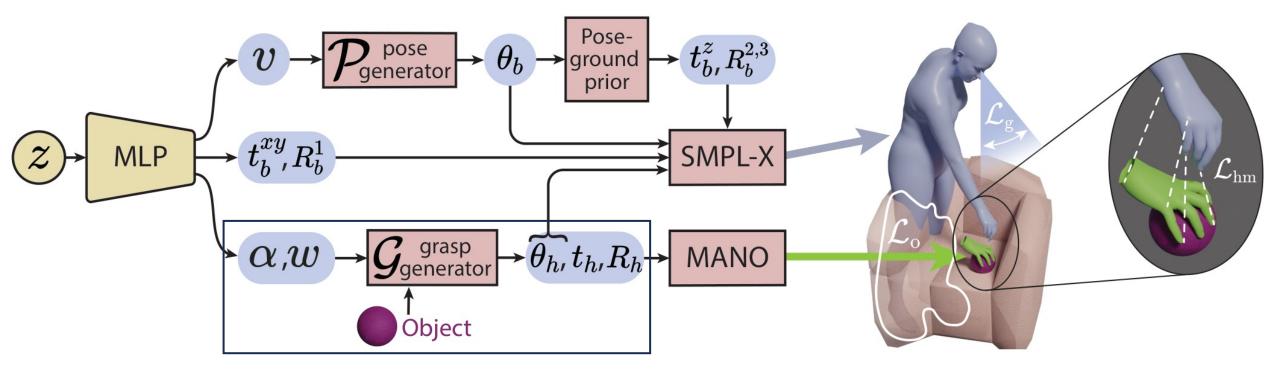
Hand Matching Loss



Hand Matching Loss



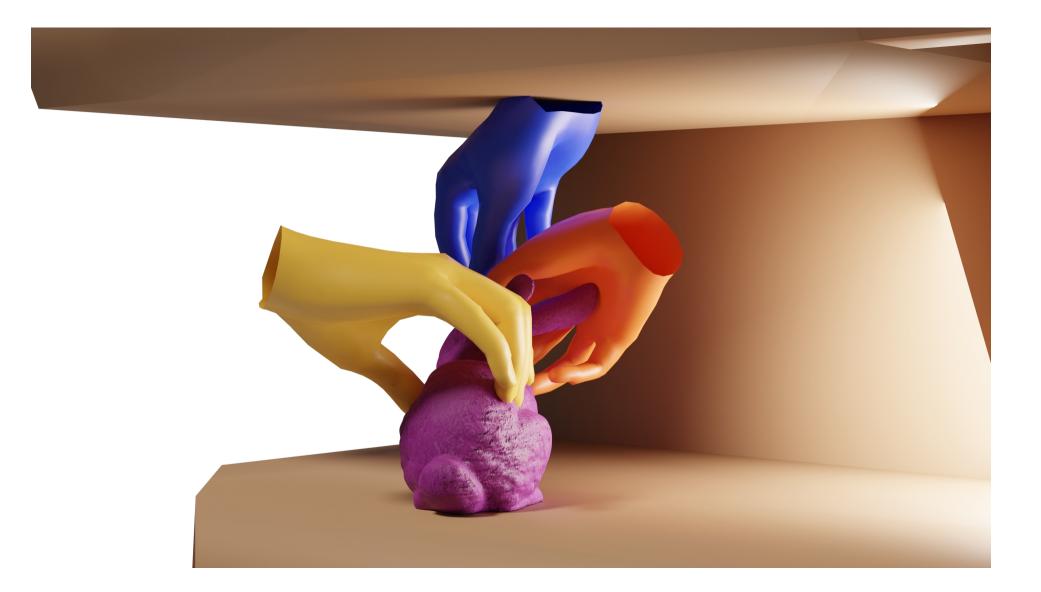




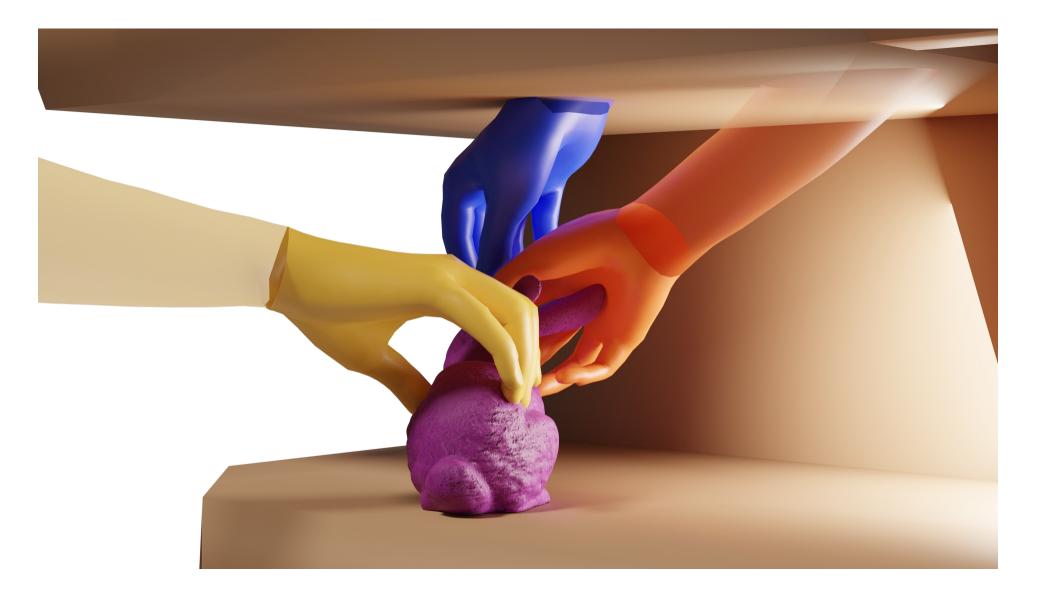
Hand Grasp Search

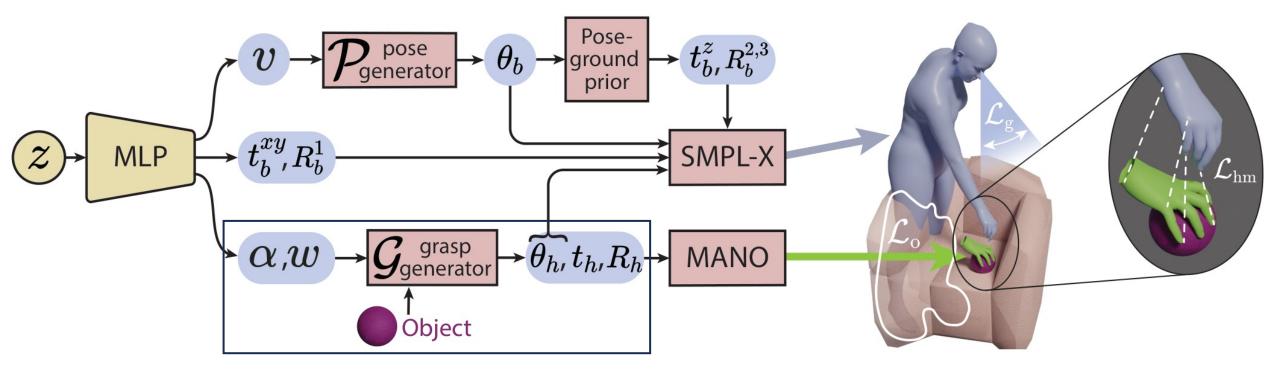


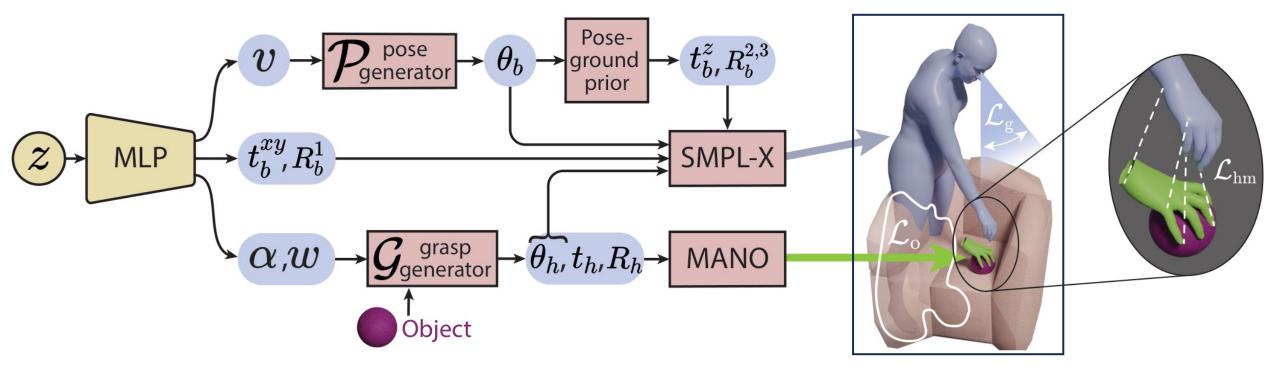
Hand Grasp Search



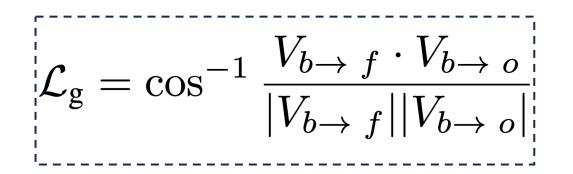
Hand Grasp Search



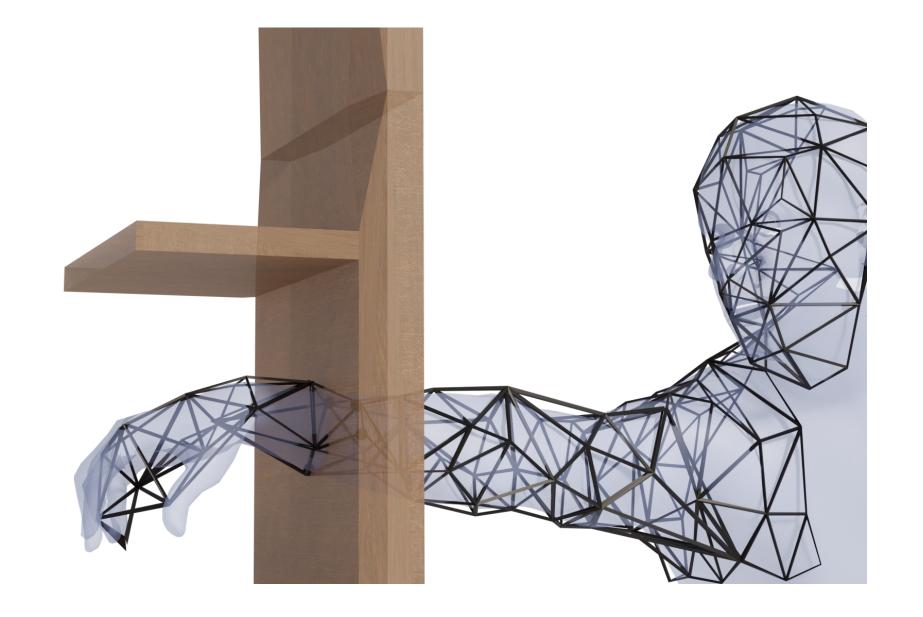


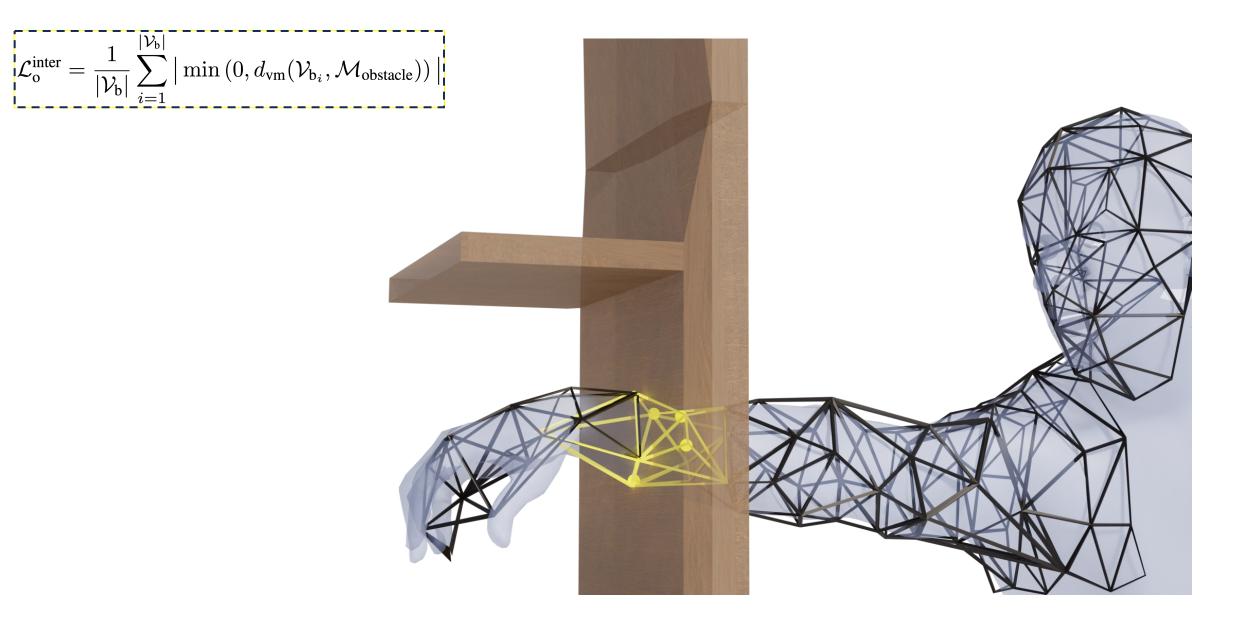


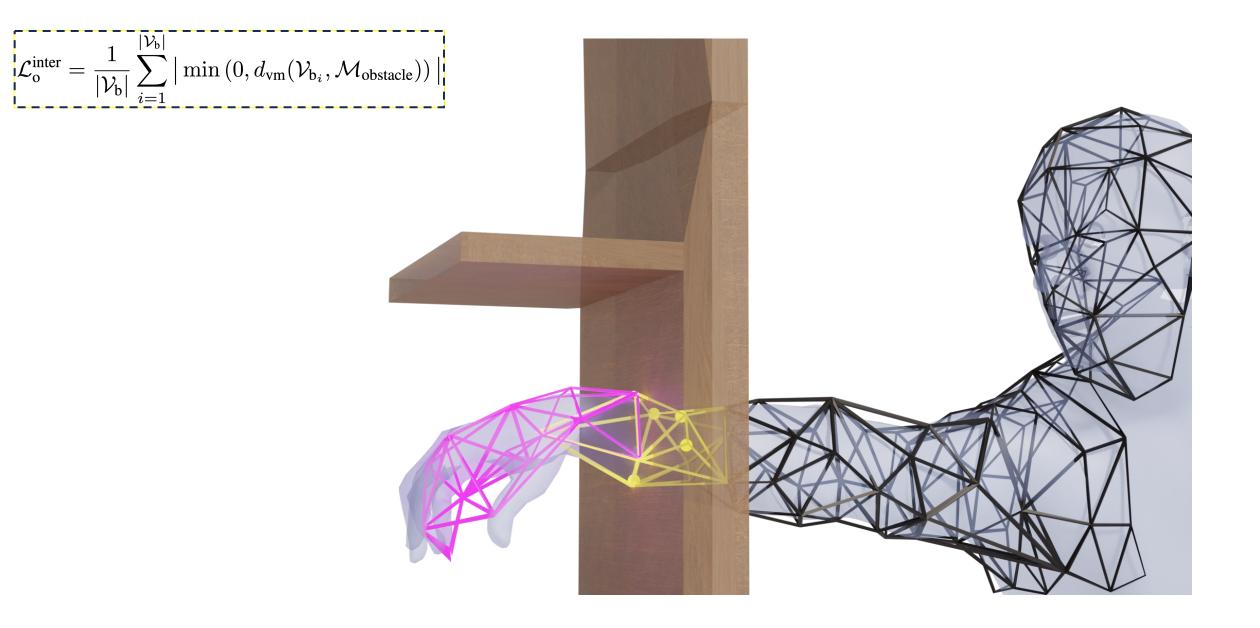
Gaze Loss

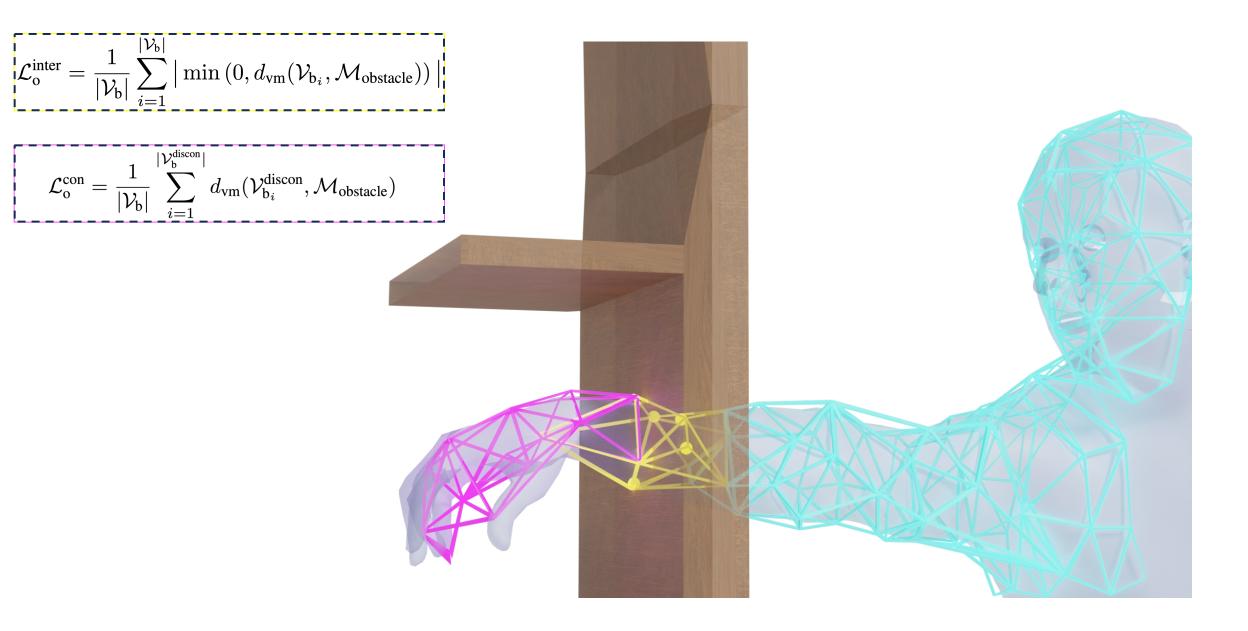


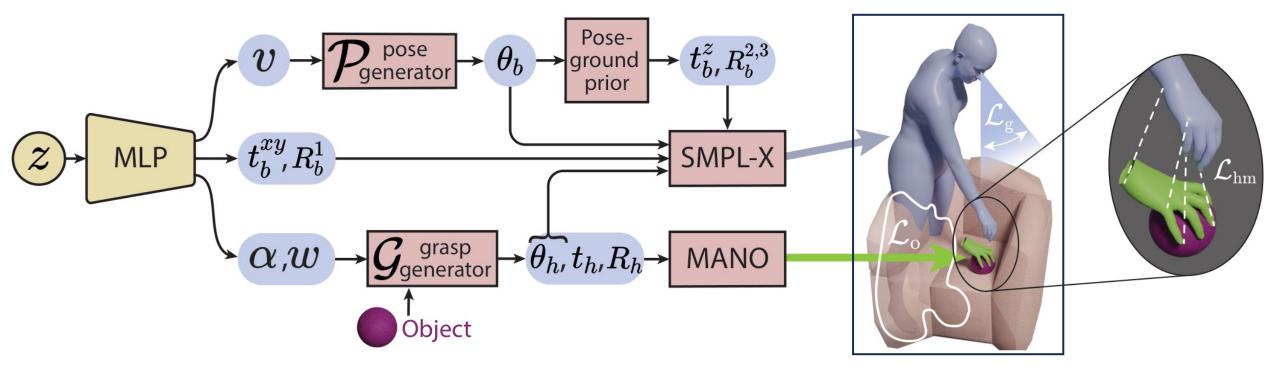


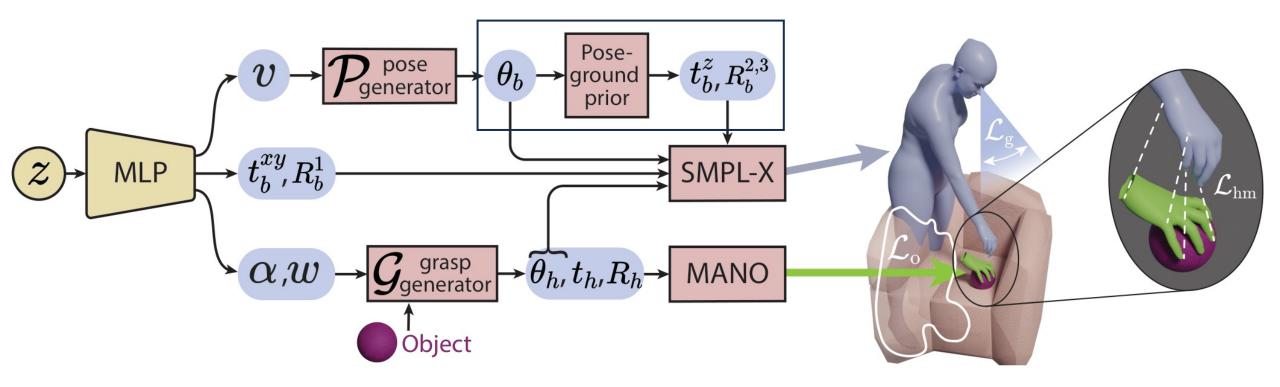


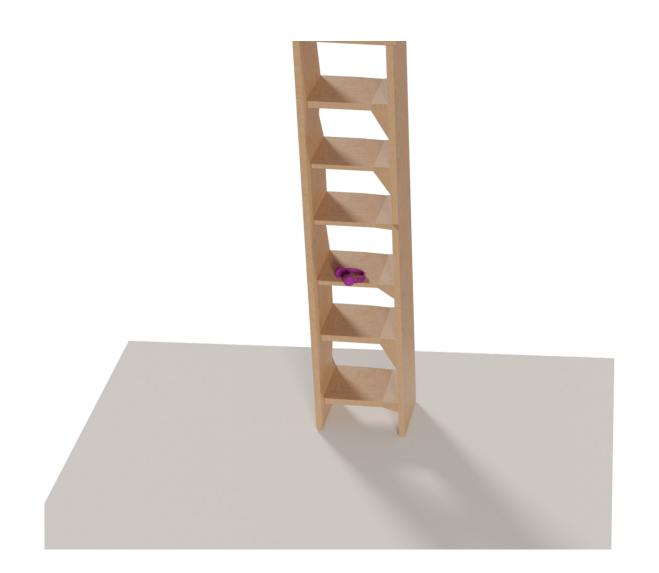


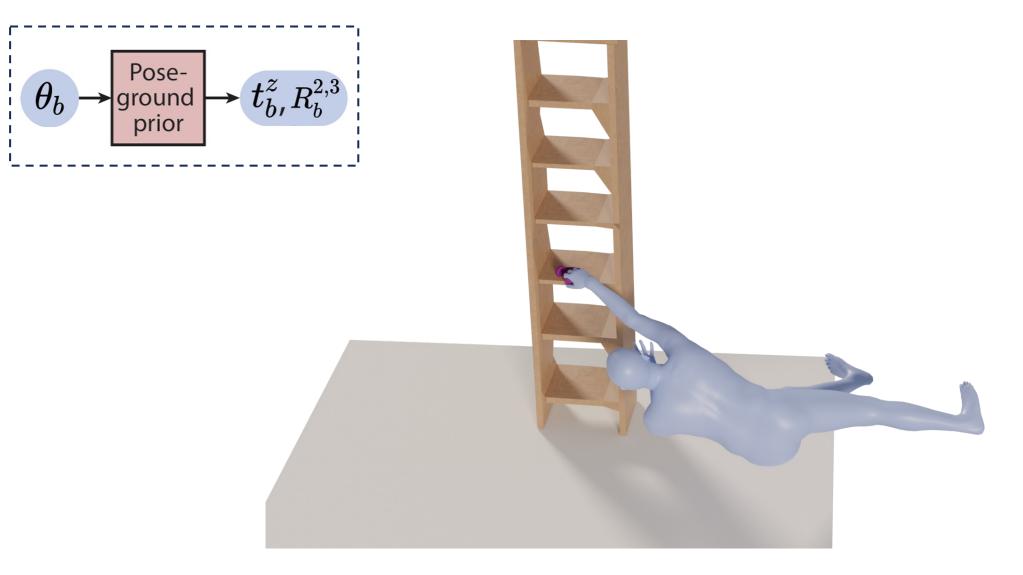


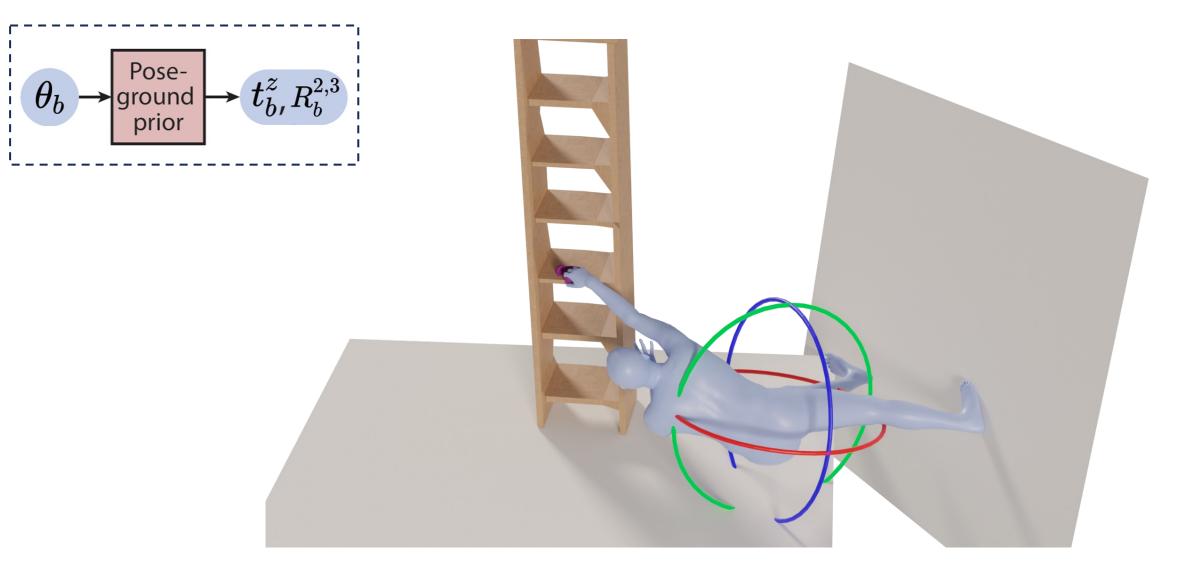


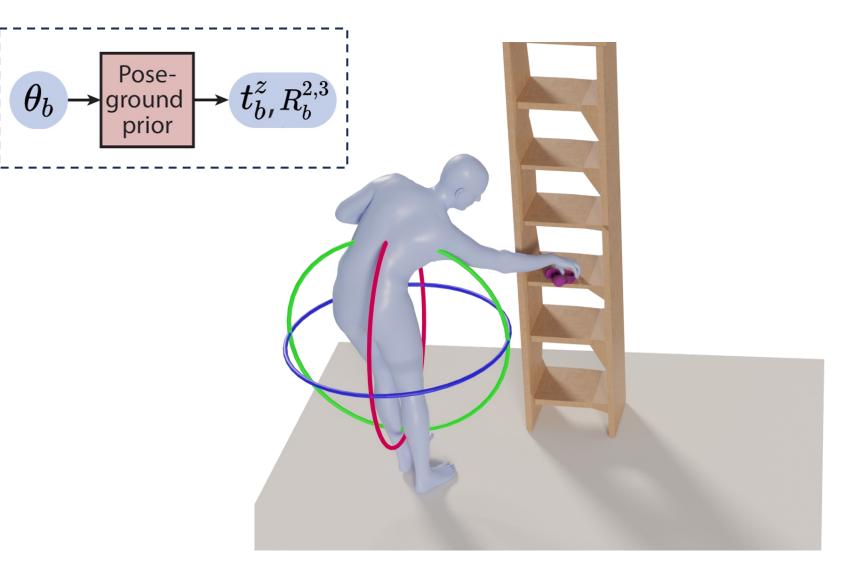




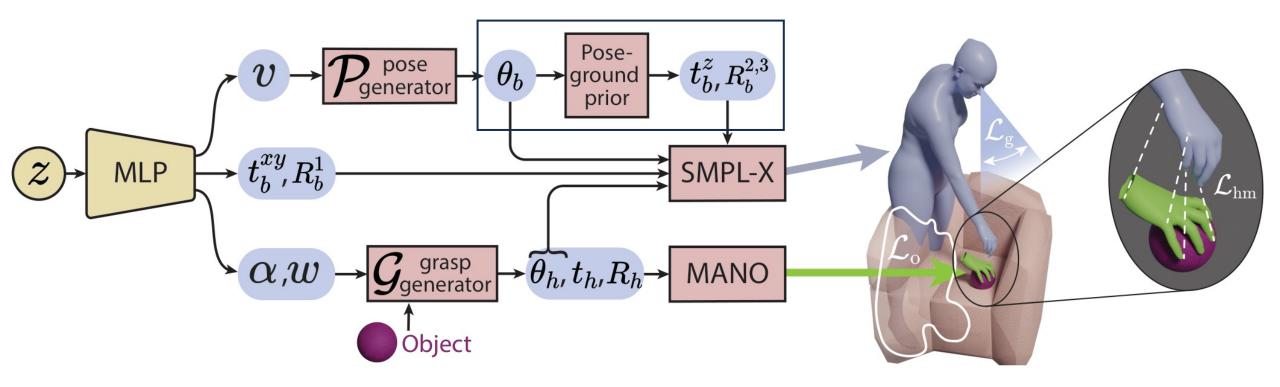


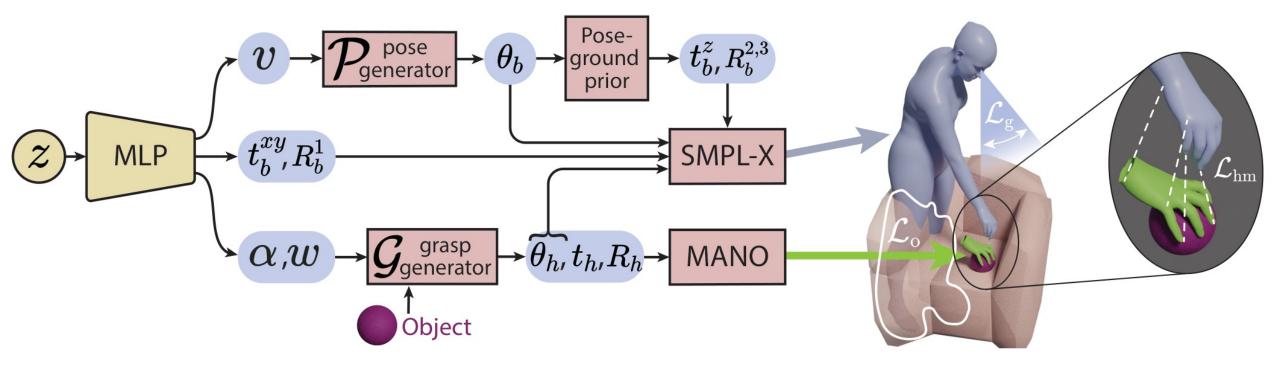


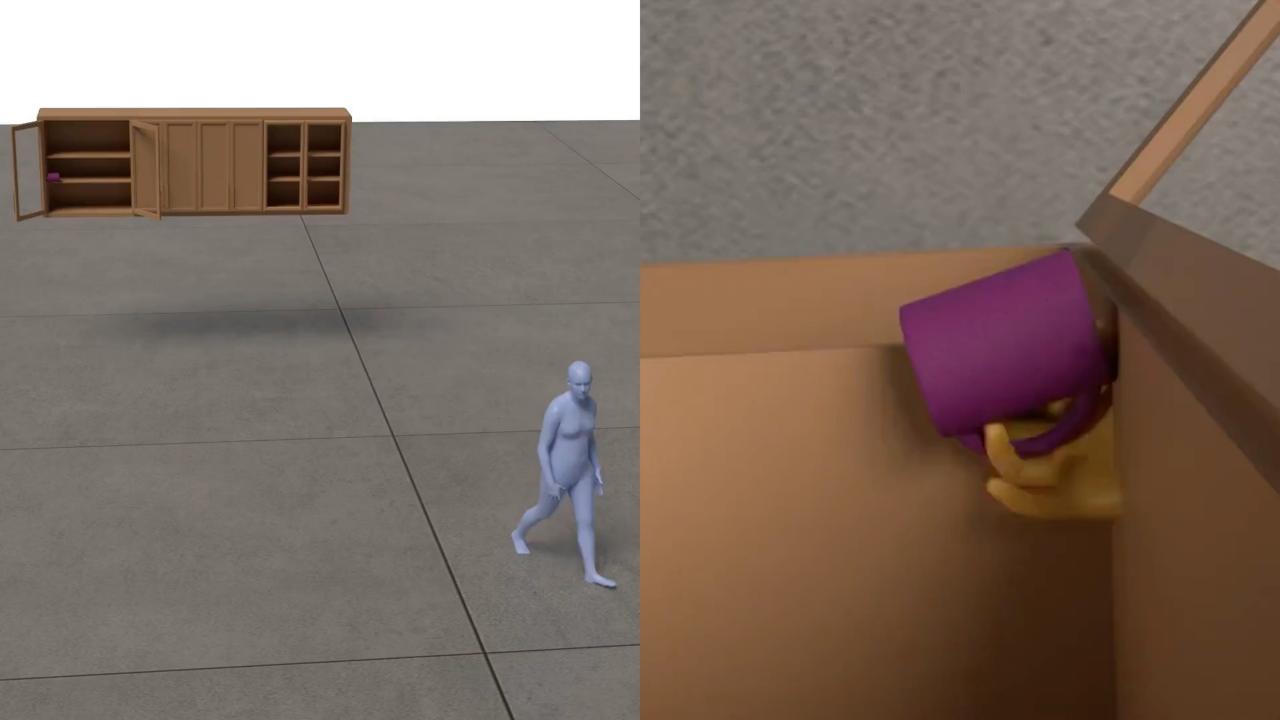




- 2-Layer MLP
- MSE Regression
- Trained on 12.7 million examples of AMASS
- Frozen during
 inference







Dataset: ReplicaGrasp



Dataset: ReplicaGrasp

- 50 everyday objects from GRAB (e.g., wineglass, cellphone)
- 48 receptacles from ReplicaCAD (e.g., couch, drawer)
- Simulated using Habitat simulator to be in one upright and random configurations
- Total 4.8K instances





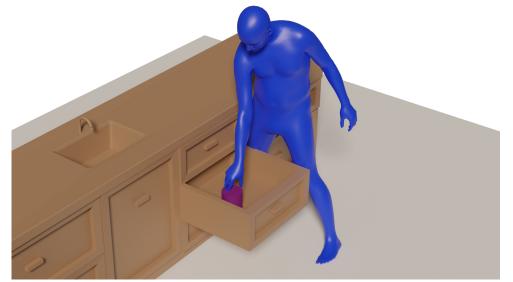








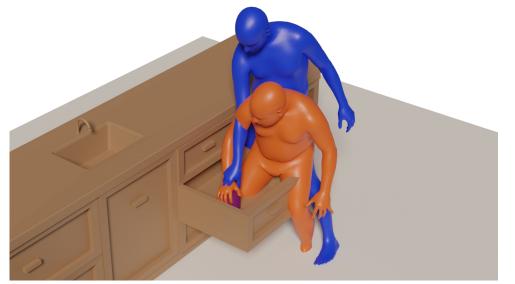












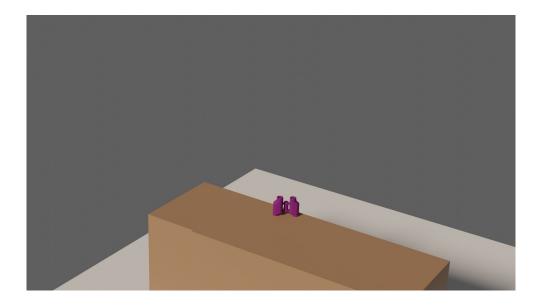
SAGA

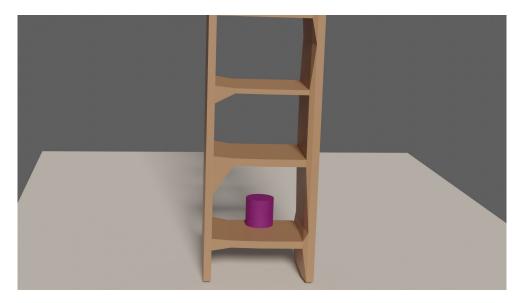




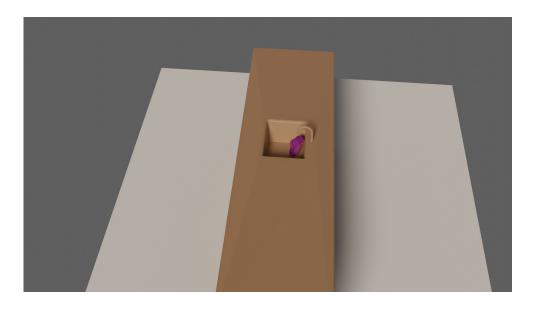


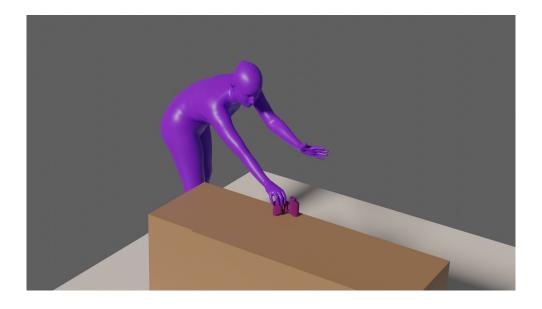






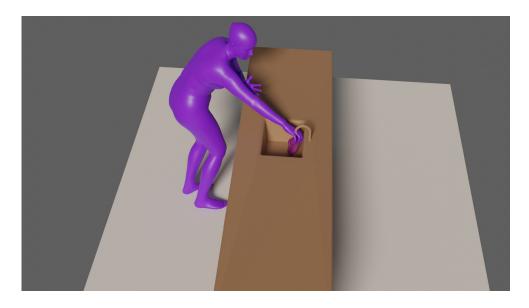


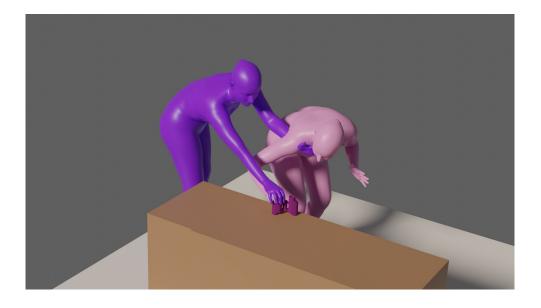






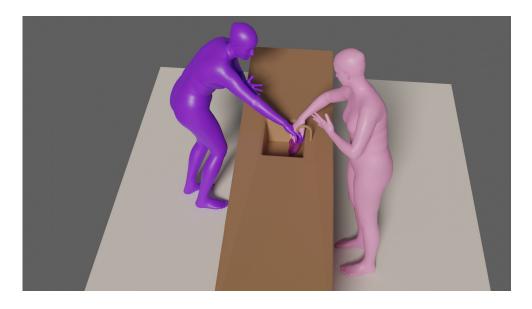












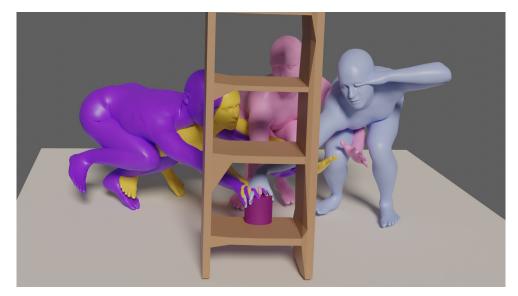


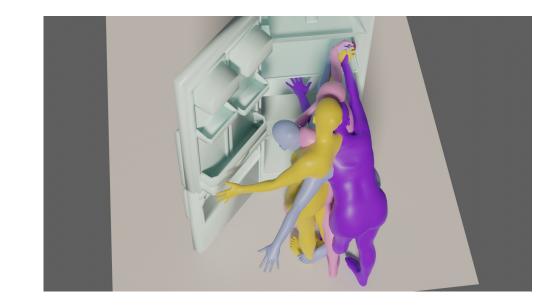


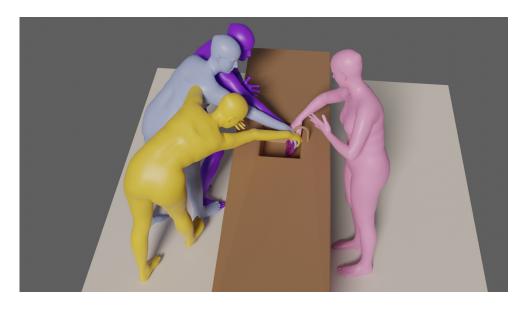




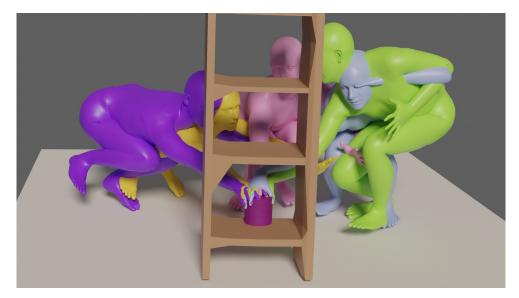




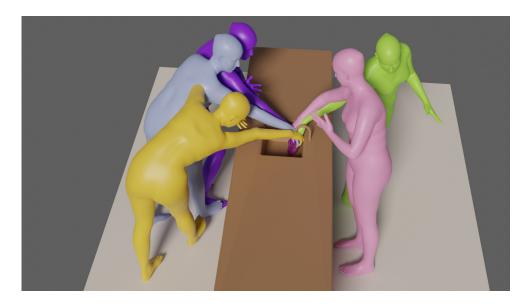












Results: Quantitative

			Repli	icaGrasp		GRAB					
Method	Obj Cont (%)↑		Obst Penet (%) ↓	-	Div _{all} (cm) ↑	Ground (cm)↓		Obj Penet (%)↓		Div _{all} (cm) ↑	Ground (cm)↓
GOAL [28]	1.14	2.14	6.87	0.20	37.89	5.31	1.62	3.50	0.11	7.87	3.56
SAGA [29] FLEX (ours)	1.21 2.20	0.29 2.49	7.27 0.53	1.10 10.37	43.84 69.78	9.40 0.00	2.19 1.63	3.43 2.73	1.06 24.94	17.35 46.46	2.39 0.00

Results: Quantitative

	ReplicaGrasp							GRAB					
Method	Obj Cont (%)↑	Obj Penet (%)↓	Obst Penet (%) ↓	Div _{samp} (cm) ↑	Div _{all} (cm) ↑	Ground (cm)↓	Obj Cont (%)↑	Obj Penet (%)↓	Div _{samp} (cm) ↑	Div _{all} (cm) ↑	Ground (cm)↓		
GOAL [28]	1.14	2.14	6.87	0.20	37.89	5.31	1.62	3.50	0.11	7.87	3.56		
SAGA [29]	1.21	0.29	7.27	1.10	43.84	9.40	2.19	3.43	1.06	17.35	2.39		
FLEX (ours)	2.20	2.49	0.53	10.37	69.78	0.00	1.63	2.73	24.94	46.46	0.00		

FLEX outperforms GOAL and SAGA on the dataset they were trained on, without even using it!

Results: Quantitative

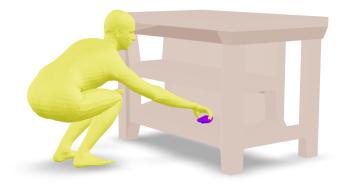
			Repli	caGrasp		GRAB					
Method	Obj Cont (%)↑	Obj Penet (%)↓	Obst Penet (%) ↓	Div _{samp} (cm) ↑	Div _{all} (cm) ↑	Ground (cm)↓	Obj Cont (%)↑	Obj Penet (%)↓	Div _{samp} (cm) ↑	Div _{all} (cm) ↑	Ground (cm)↓
GOAL [28]	1.14	2.14	6.87	0.20	37.89	5.31	1.62	3.50	0.11	7.87	3.56
SAGA [29]	1.21	0.29	7.27	1.10	43.84	9.40	2.19	3.43	1.06	17.35	2.39
FLEX (ours)	2.20	2.49	0.53	10.37	69.78	0.00	1.63	2.73	24.94	46.46	0.00

FLEX also outperforms both baselines on ReplicaGrasp.

Results: Qualitative

Claim: "This pose is natural and realistic for grasping the object.".

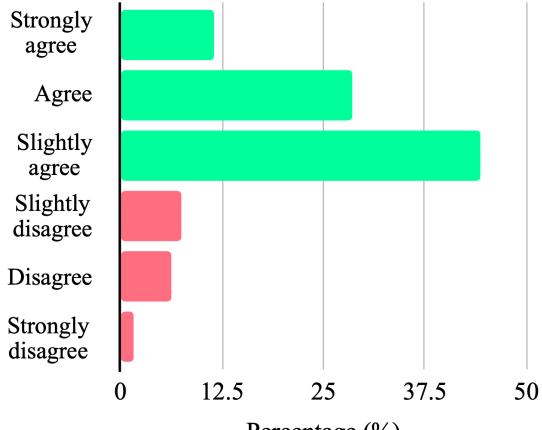
Note: the scene may not show up in the preview, but will be visible once you accept. If it is not visible upon acceptance, please select option "7. The scene is not showing".



0	1. Strongly disagree	O 2. Disagree	0	3. Slightly disagree	0	4. Slightly agree	O 5. Agree	0	6. Strongly agree	0	7. The scene is not showing
---	-------------------------	---------------	---	-------------------------	---	-------------------	------------	---	-------------------	---	--------------------------------

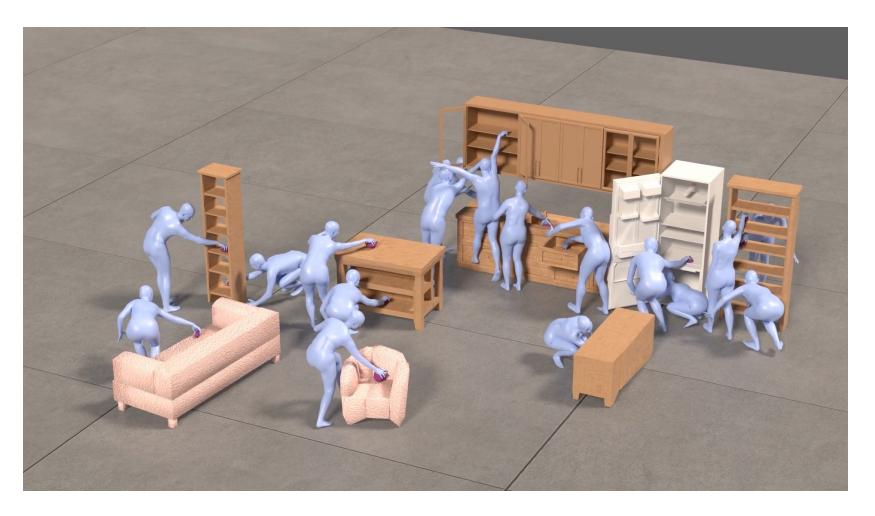
You must select one of the options above. Do not leave unselected.

(Optional) Provide a reason for your selection



Percentage (%)

Conclusion



- FLEX a method to generate full-body grasps for objects in everyday scenes while avoiding obstacles.
- ReplicaGrasp a dataset of 50 objects spawned in everyday scenes in a variety of configurations
- Demonstrate full-body grasping results that outperform baselines quantitatively and qualitatively.

FLEX: Full-Body Grasping Without Full-Body Grasps



Purva Tendulkar



Dídac Surís



Carl Vondrick

CVPR 2023

flex.cs.columbia.edu

Poster: THU-PM-055

COLUMBIA UNIVERSITY