


FLEX: Full-Body Grasping Without Full-Body Grasps

Purva Tendulkar Dídac Surís Carl Vondrick

Task and Challenges


Full-body grasping of objects in presence of obstacles.



MoCap Challenges No Generalization No Obstacle Consideration

Key Idea

- Full-Body Grasp \rightarrow Full-Body + Hand-Grasp
- No explicit full-body grasps required for training.



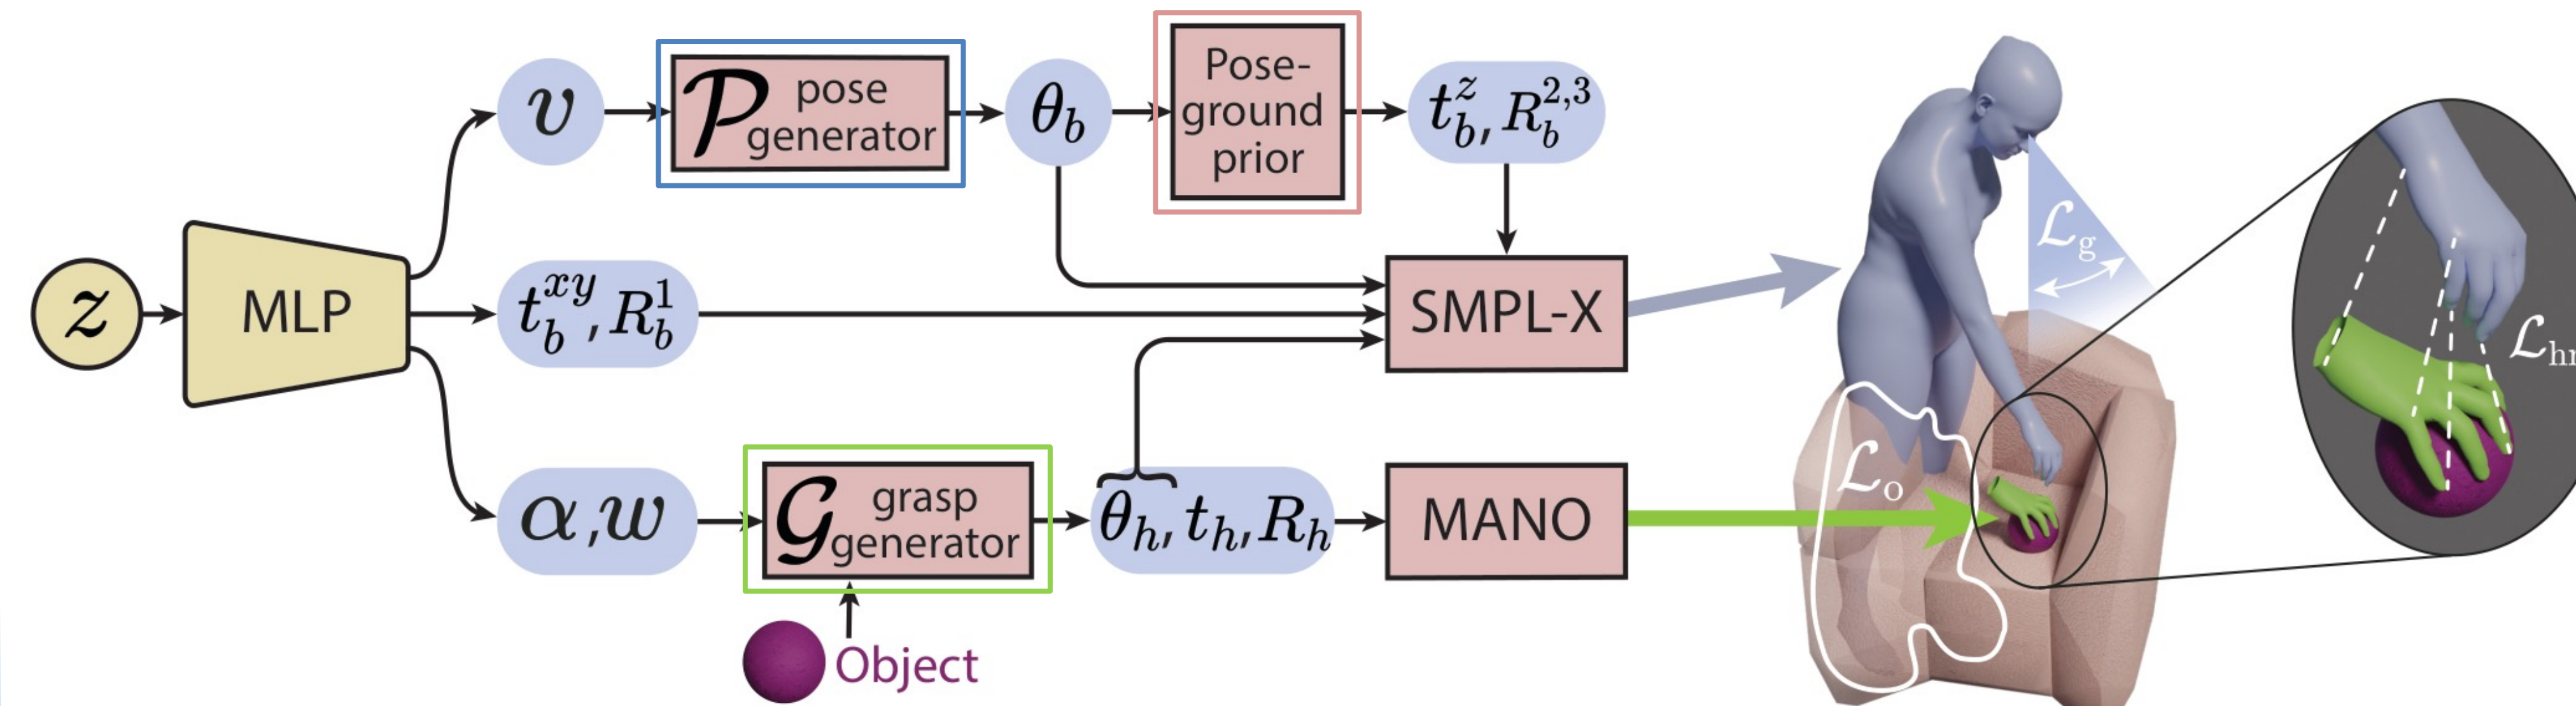
ReplicaGrasp Dataset



Approach

FLEX (Full-Body Latent Exploration) generates a 3D human grasping the desired object, given

1. Pre-trained right-hand grasping model \mathcal{G} that can predict global MANO parameters $\{\theta_h, t_h, R_h\}$
2. Pre-trained full-body pose prior \mathcal{P} that can generate feasible full-body poses θ_b
3. Learnt pose-ground prior which predicts the floor given the human pose θ_b




FLEX performs a gradient-based search over hand & body priors to minimize hand-matching and obstacle losses.

Key Insights

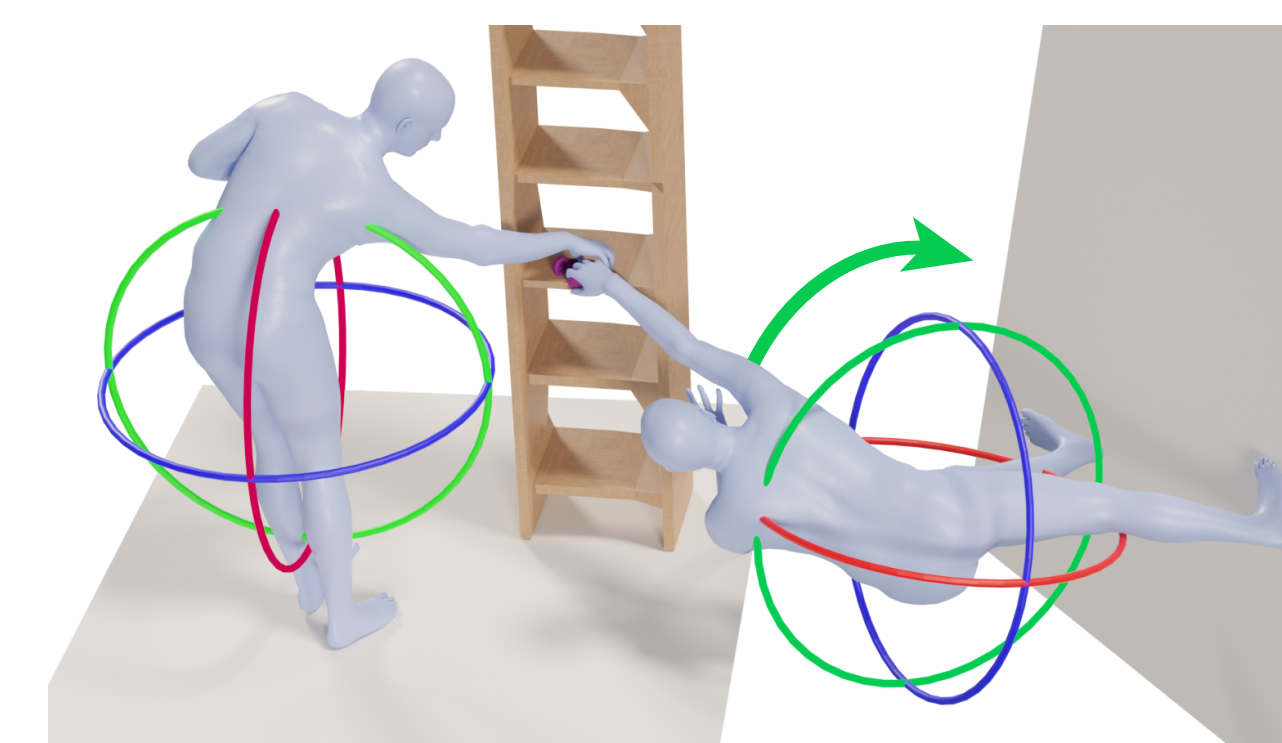
- Searching in the latent space of the hand-grasping model \mathcal{G} allows full-body obstacle consideration.
- We penalize *all* vertices in the connected components of the resulting body graph other than the largest one.
- The ground position can be predicted from the body pose. This removes 2 DoFs from the optimization.



Hand-grasp search




Obstacle avoidance loss




Pose-ground Prior

Comparison Results



Diversity Results



All humans are samples from FLEX