Learning from Electromyography Synergies to Grasp Novel Objects by Superquadric Representation

Abhi Gupta, Jingya Bi, Ashwin Jayaraman, Max Xu, David Watkins and Professor Peter Allen
Columbia University: Fu Foundation School of Engineering and Applied Sciences

CTRL-labs

INTRODUCTION

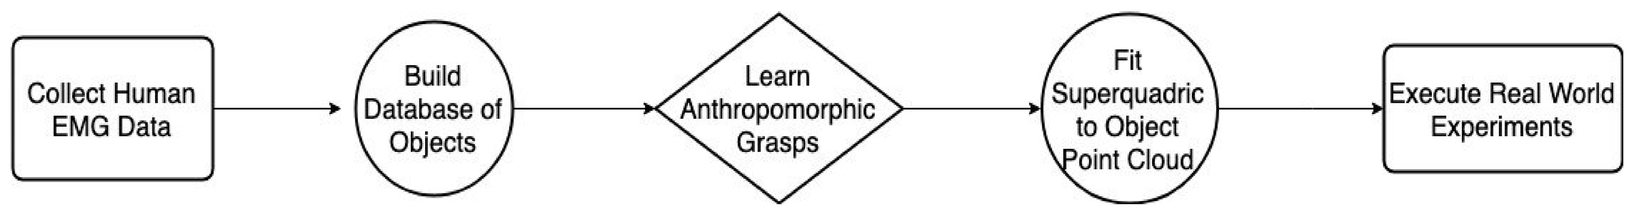

The objective is to learn grasp synergies with high fidelity given object pose and geometry. However, under-actuated, anthropomorphic hands require complex, high dimensional control strategies. Including object pose and geometry further increase the size of the state space. Therefore, grasping in unstructured environments in the same fashion as humans proves to be non-trivial.

DIMENSIONALITY REDUCTION WITH EMG SYNERGIES

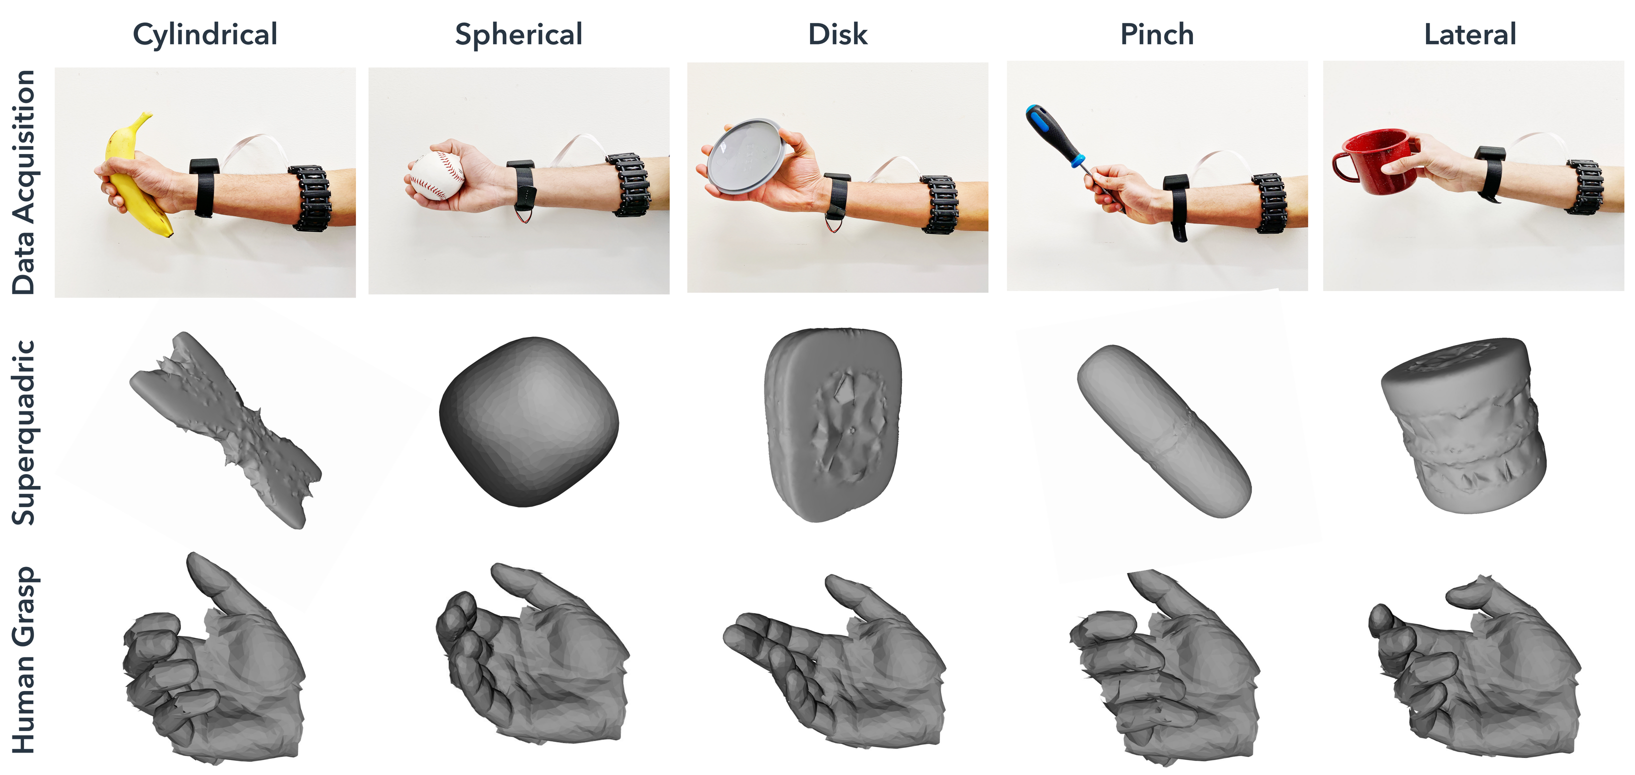

$X = \{x^{jointstates}, x^{EMG}, x^{objectgeometry}, x^{objectsize}\} \subset \mathbb{R}^{42}$

$h = f(x) \quad x \in X \subset \mathbb{R}^{42}$ f is an encoder function that maps x to space of synergies

$y = g(h) \quad y \in Y \subset \mathbb{R}^{42}$ g is a decoder function that maps h back to the original dimension

$x \sim g(f(x))$ g(f(x)) is an approximation of x

Encoding human grasps collected from the Ctrl-Labs' EMG device in a latent space will effectively extract synergies. Grasps of the same type should be clustered in the lower dimensional space to allow for the sampling of a grasp with given type and object geometry. This is not possible when applying principle component analysis (PCA). A deep autoencoder approach has been shown to perform extremely well.

OBJECT GEOMETRY PARAMETRIZATION

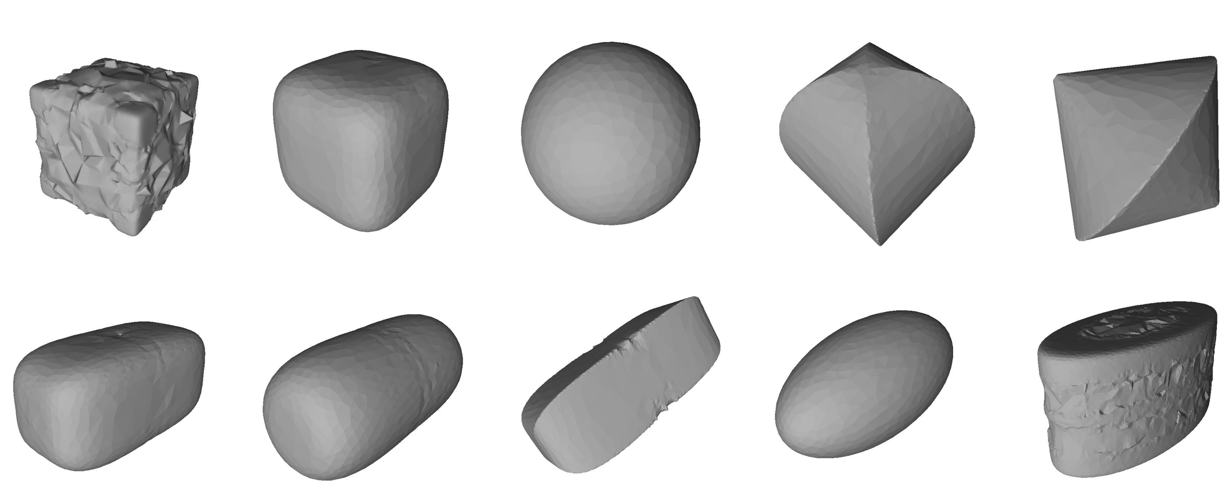

$$f(a,x,y,z): \left(\left|\frac{x}{a_1}\right|^{\frac{2}{\epsilon_2}} + \left|\frac{y}{a_2}\right|^{\frac{2}{\epsilon_2}} \right)^{\frac{\epsilon_2}{\epsilon_1}} + \left|\frac{z}{a_3}\right|^{\frac{2}{\epsilon_1}} = 1.$$

Object geometry was parametrized with the model of a superquadric. Varying these parameters generated a set of points lying on the surface of the respective object. After sampling a sufficient amount of points, Poisson reconstruction was used to generate a databases of unseen object meshes for training in simulation.

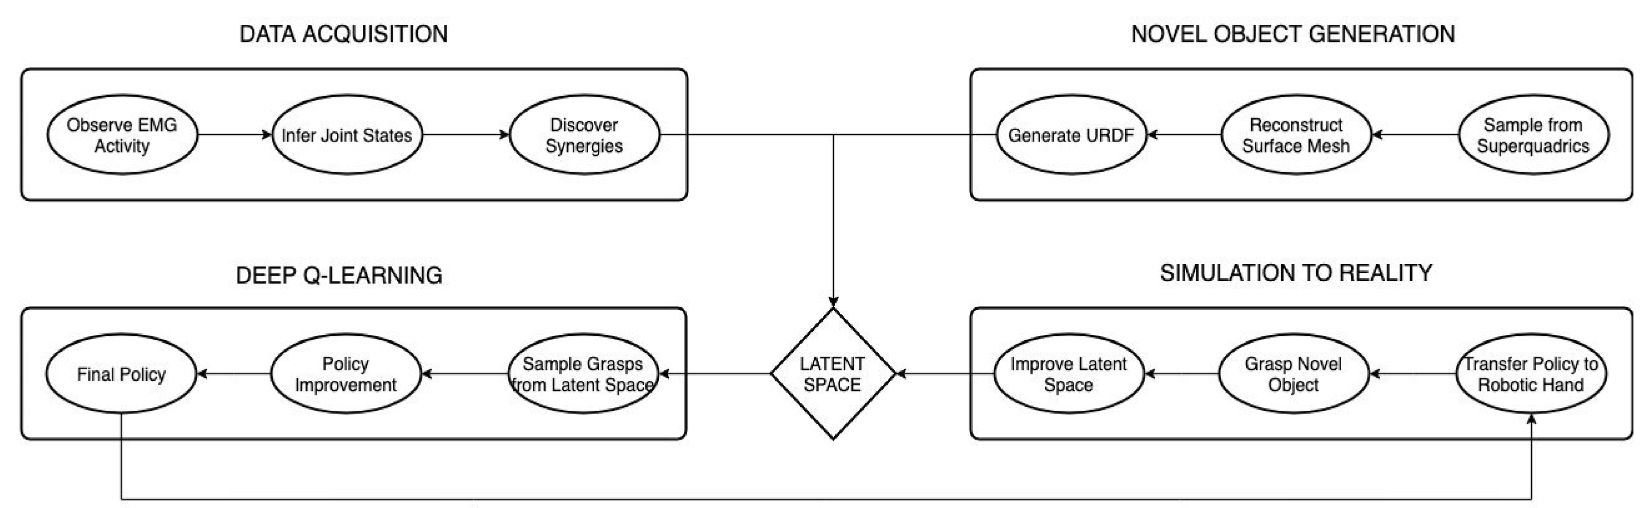

REINFORCEMENT LEARNING ENVIRONMENT

A reinforcement learning environment simulating a robotic arm-hand unit with artificial contact forces was developed using PyBullet. In particular, the robot is introduced to table top objects of varying size and geometry by superquadric representation. The scenes on the right demonstrate the end-effector keeping a constant distance from the sphere.

Although the hand has 18 joints, only 5 are fully-actuated. An empirically estimated linear relationship models all other joints as a proportion of these 5.

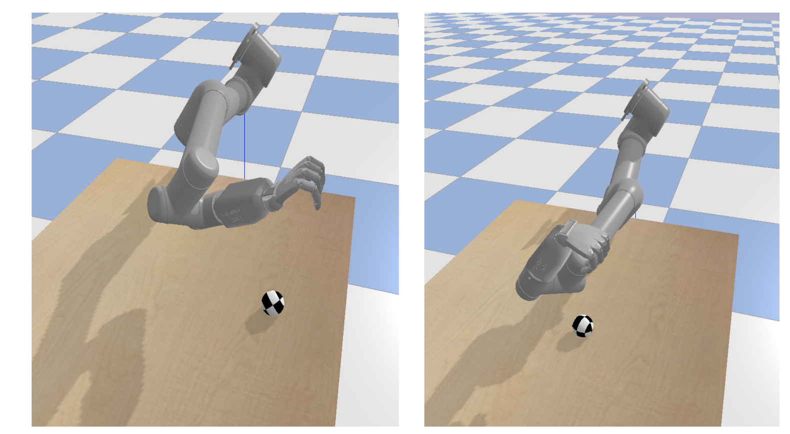

Inverse Kinematics

Moving the end-effector to a given hand pose requires computing the joint angles of the arm.

Forward Kinematics

Computing the amount of contact between the hand and object requires poses of both bodies. Forward kinematics will compute the Cartesian fingertip positions as a function of the joint angles of the arm and hand.

GRASPING WITH DEEP Q-LEARNING

State Space

$$S = \{s^{arm}, s^{hand}, s^{object}\} \subset \mathbb{R}^{26}$$
$$s^{object} = \{position, orientation, geometry, size\}$$
$$geometry = \{a_1, a_2, a_3, \epsilon_1, \epsilon_2\}$$
$$size = \{n\}$$

The motion of a robotic end-effector either operates in *cartesian space* or *configuration space*. Restricting the state space size of the arm becomes difficult under the assumption that it can reach any position and orientation. Controlling the joint angles of the arm however, constraints the possible hand poses because no joint can move beyond its lower and upper limits. Therefore, the hand and arm states are a vector of joint angles.

Algorithms

We implemented Policy Gradient and Deep Q Learning Algorithms. We used Neural Networks for training since they were the state of the art algorithms for Robotics.

Action Space

$$A \leftarrow A + I_a * \delta_\theta$$

It consists of moving the joint angles by an angle of 0.01 radians.

Reward Structure

$$Reward = \begin{cases} -e & e \geq 1 \\ c & c \geq 1 \;\&\; d = False \\ 100 & d = True \end{cases}$$

e: distance of the hand from the object
c: number of contact points
d: success of the grasp

The reward is given to the agent piece-wise. Because the reward at a successful grasp tends to be sparse, intermediate rewards discourage the robot from moving away and incentivize contact between the hand and the object.

RESULTS

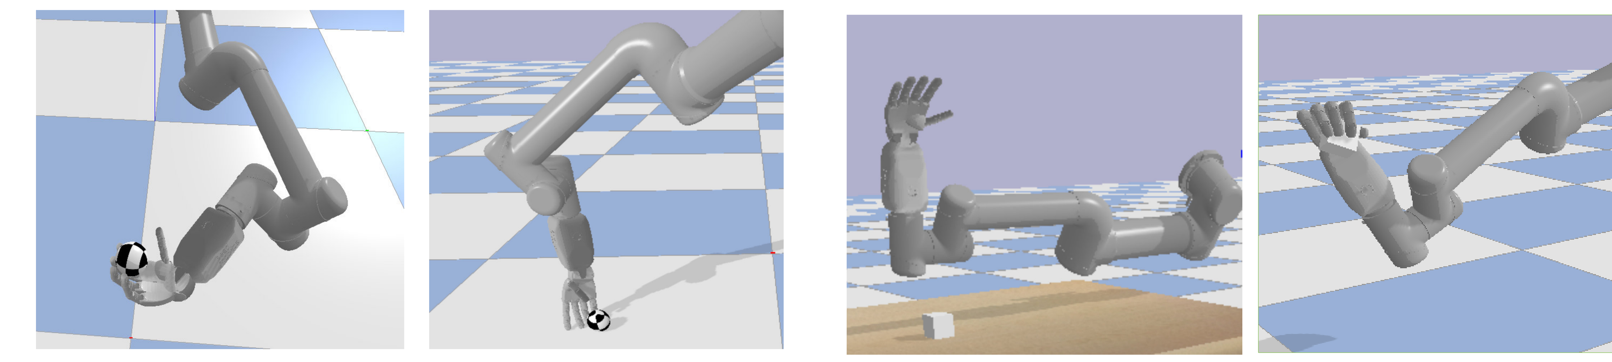

Introducing a prior of a stable grasp to the network pushes the robotic hand to pick up objects far more quickly than otherwise. As the first two scenes exhibit above, the hand will only learn to touch the object rather than lift it up as a result of rewarding high contact points. Nonetheless, the agent eventually succeeds at performing an anthropomorphic grasp of arbitrary objects.

Results below demonstrate the importance of choosing initial parameters to an RL policy. Motivated by this observation, collecting more samples of human grasps will ultimately improve learning and transferring from simulation to reality.

Algorithms	Performance[1]	No of Iterations
Deep Q Learning	52	1000
Deep Q Learning	23	1000
Deep Q Learning	5	100
Policy Gradient	2	300

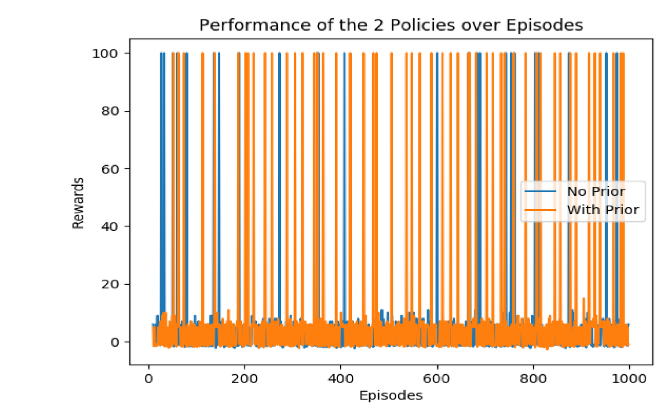

SIMULATION TO REALITY

Simulation uses a database of superquadrics with their preferred grasps. In experiment, superquadrics are fitted to a 3D point cloud of table top objects to sample an appropriate starting grasp for the learned policy.

The robot trained in simulation is capable of executing the learned policy in experiment. The tabletop scene portrays the hand moving toward a cube to perform a grasp.

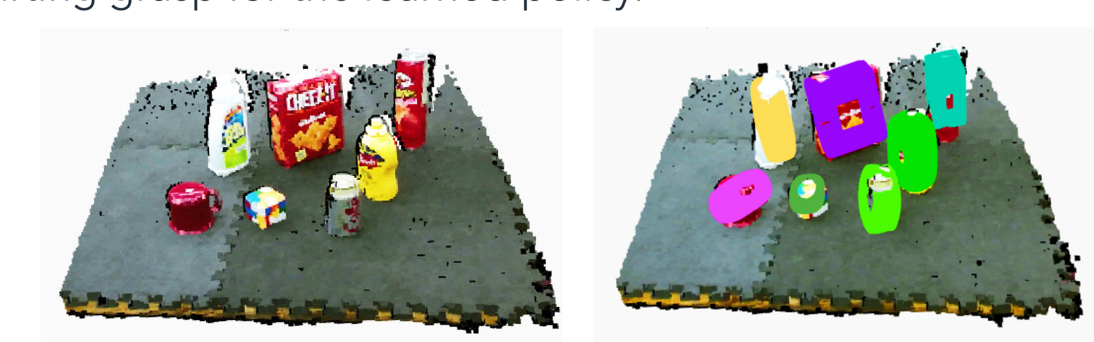

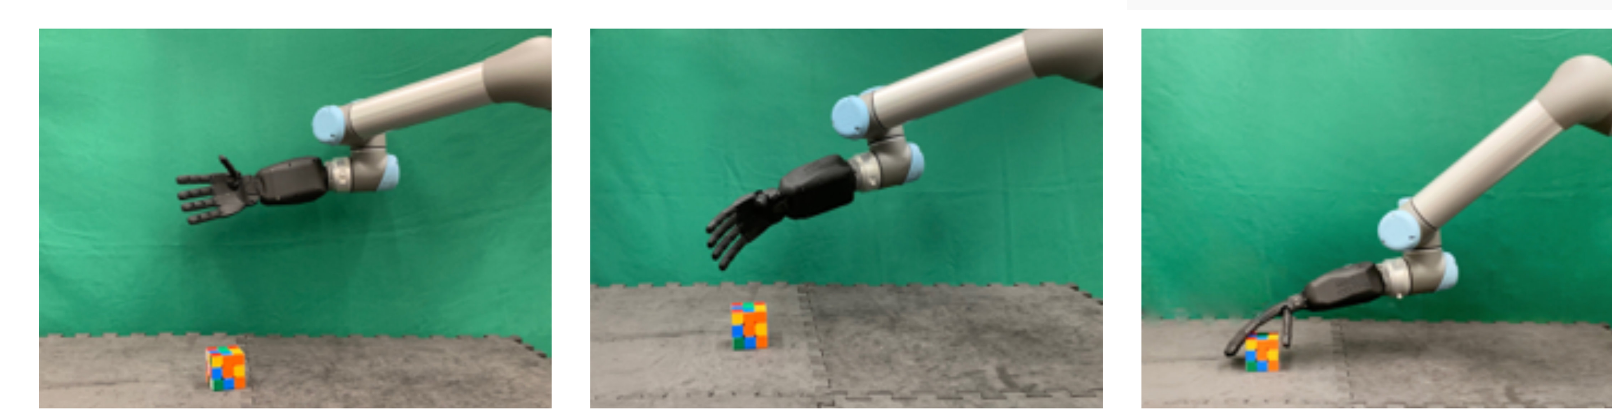

REFERENCES

T. Lampe and M. Riedmiller, "Acquiring visual servoing reaching and grasping skills using neural reinforcement learning," The 2013 International Joint Conference on Neural Networks (IJCNN), Dallas, TX, 2013, pp. 1-8.

Wu, Bohan & Akinola, Iretiayo & K. Allen, Peter. (2019). Pixel-Attentive Policy Gradient for Multi-Fingered Grasping in Cluttered Scenes.

Starke, Julia & Eichmann, Christian & Ottenhaus, Simon & Asfour, Tamim. (2018). Synergy-Based, Data-Driven Generation of Object-Specific Grasps for Anthropomorphic Hands.

Makhal, Abhijit & Thomas, Federico & Perez-Gracia, Alba. (2018). Grasping Unknown Objects in Clutter by Superquadric Representation. 292-299. 10.1109/IRC.2018.00062.

Ariel Sharone) Anders, Ariel. (2014). Learning a strategy for whole-arm grasping.

Korkmaz, Semih. (2018). Training a Robotic Hand to Grasp Using Reinforcement Learning.