OpenAI

Learning Dexterous In-Hand Manipulation

Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafał Józefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, Wojciech Zaremba

Setup & Task

Simulation Environment

Training

Transfer

Real-World Environment

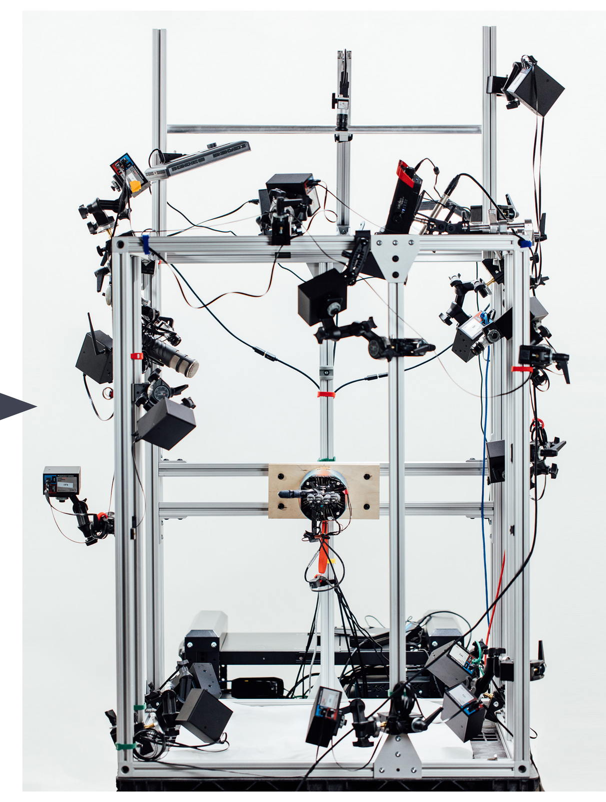

Deployment

Goal

Train entirely in simulation and achieve zero-shot transfer to real robot hand.

System Overview

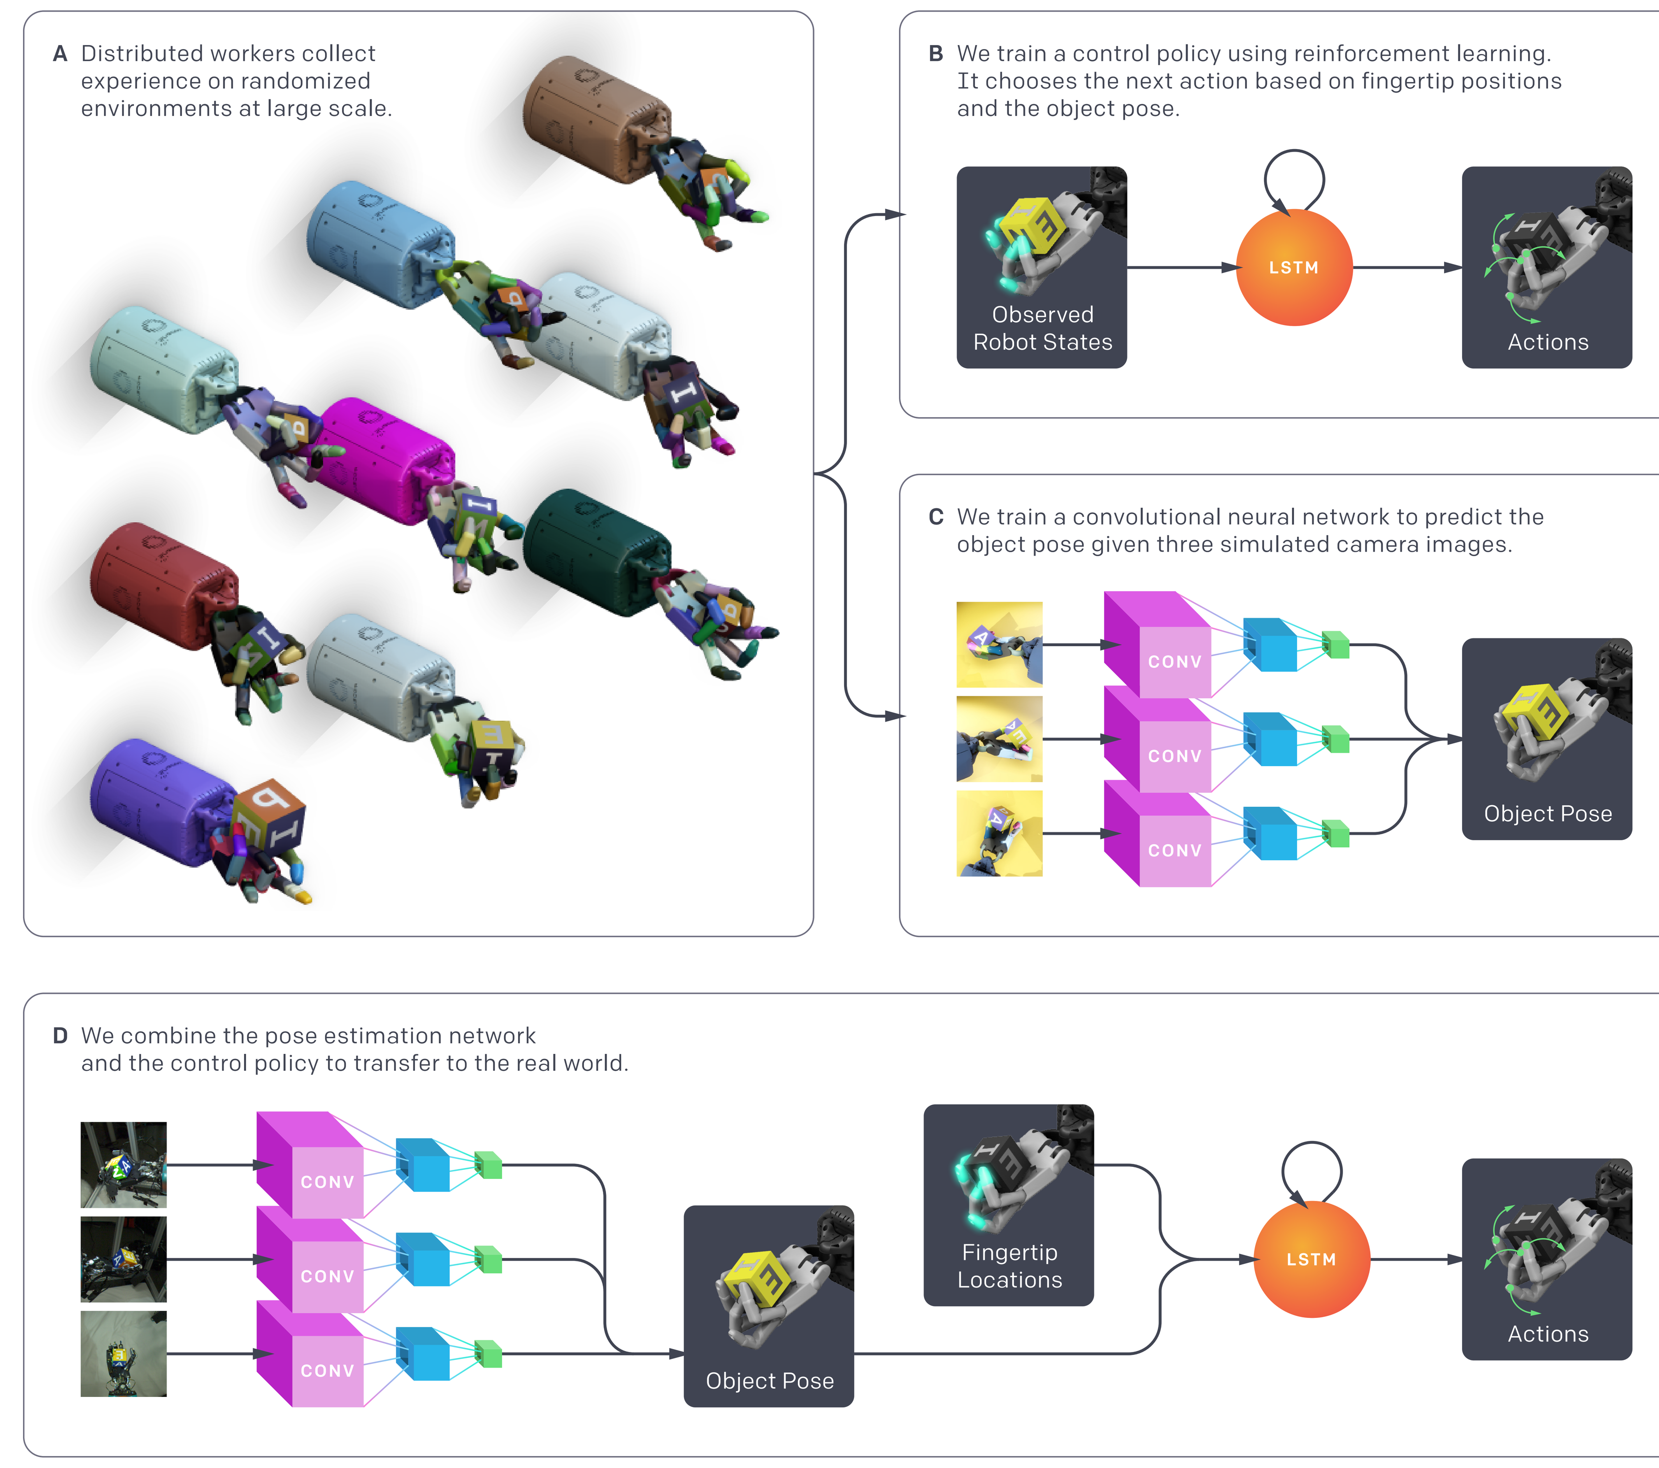

Key Elements

1. Domain Randomization

Physics Randomizations

Object Dimensions	Actuator Force Gains
Object And Robot Link Masses	Joint Limits
Surface Friction	Gravity Vector
Robot Joint Damping	Noisy Observations
Backlash	Noisy Actions

Appearance Randomizations

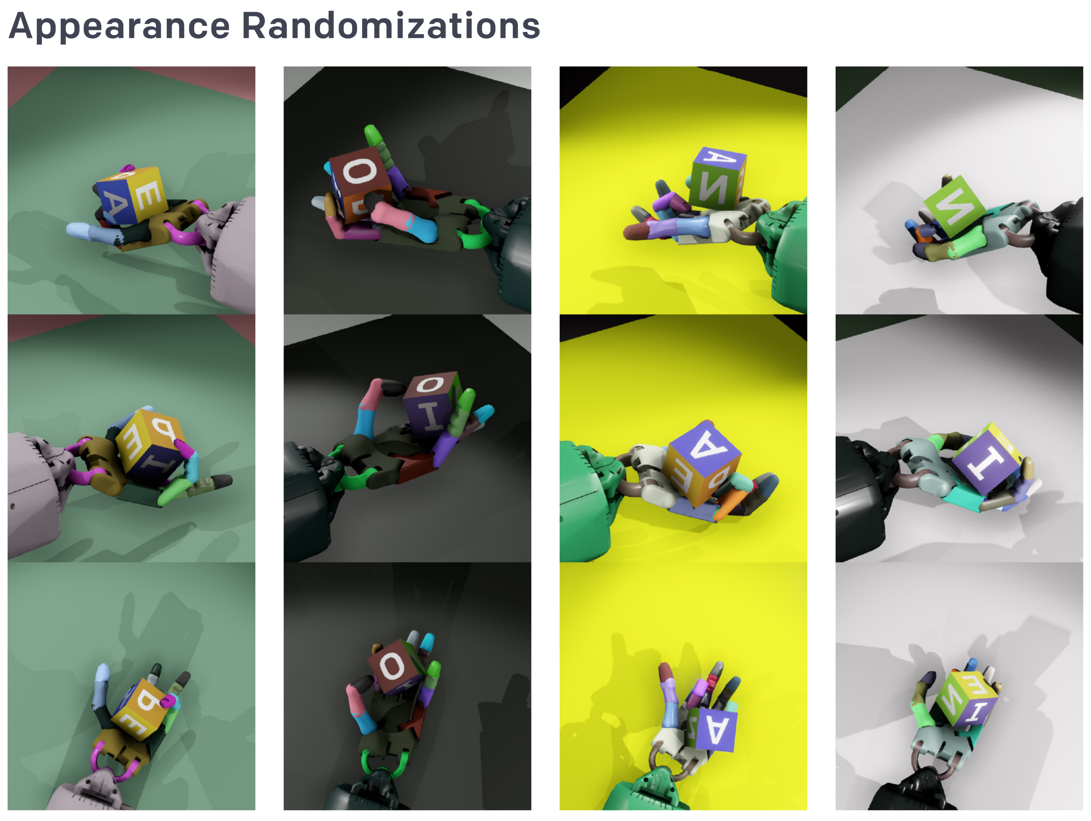

2. Memory Augmented Policy

Policy Architecture

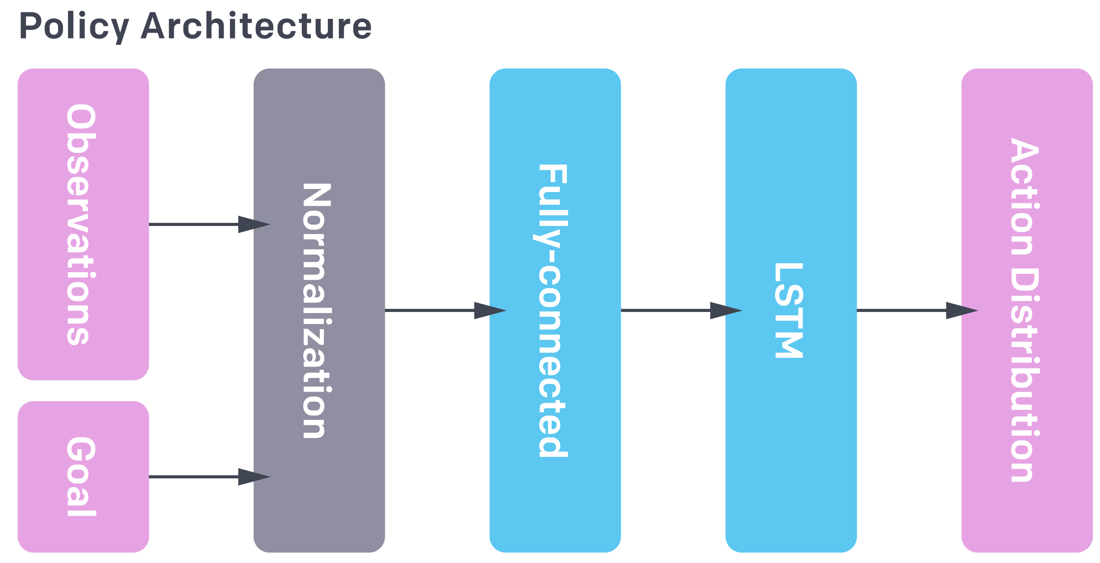

Effect of Memory in Simulation

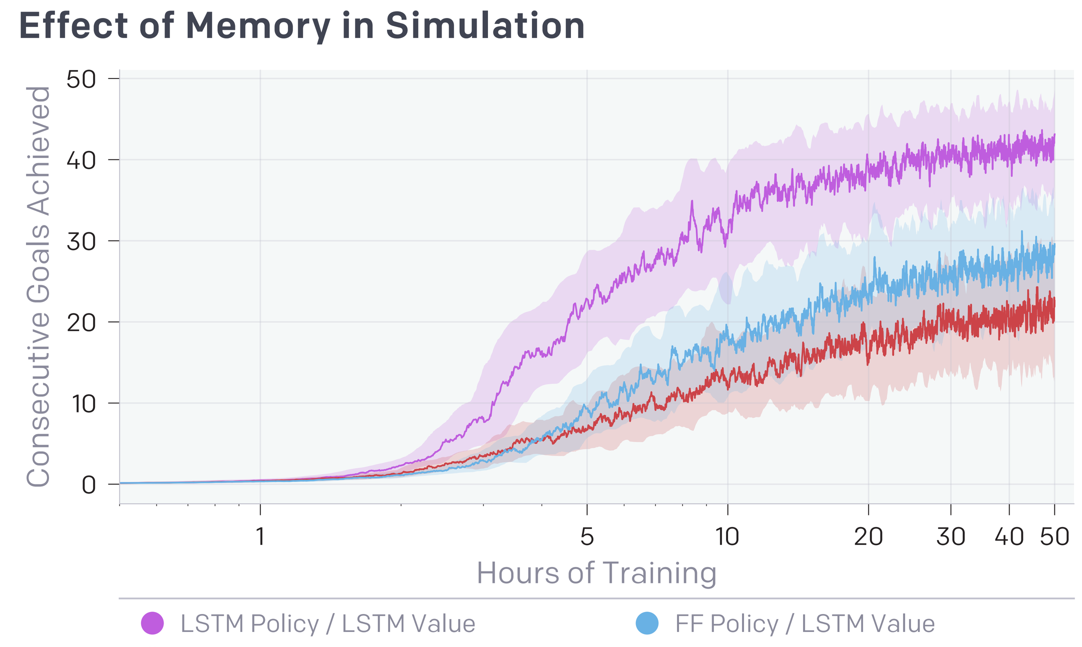

3. Large Scale RL

Distributed Training with Rapid

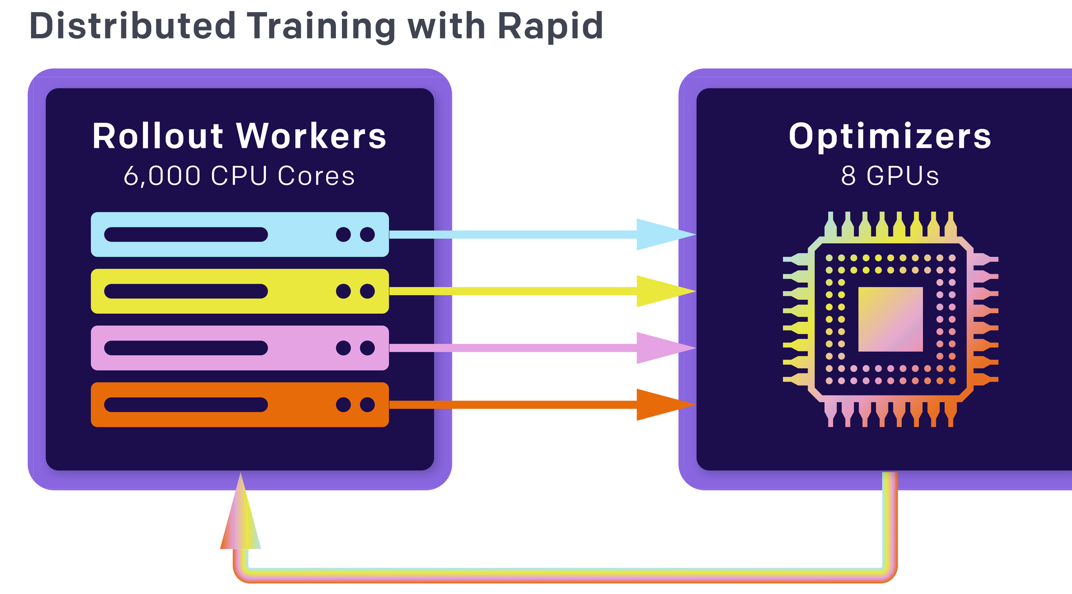

Effect of Scale in Simulation

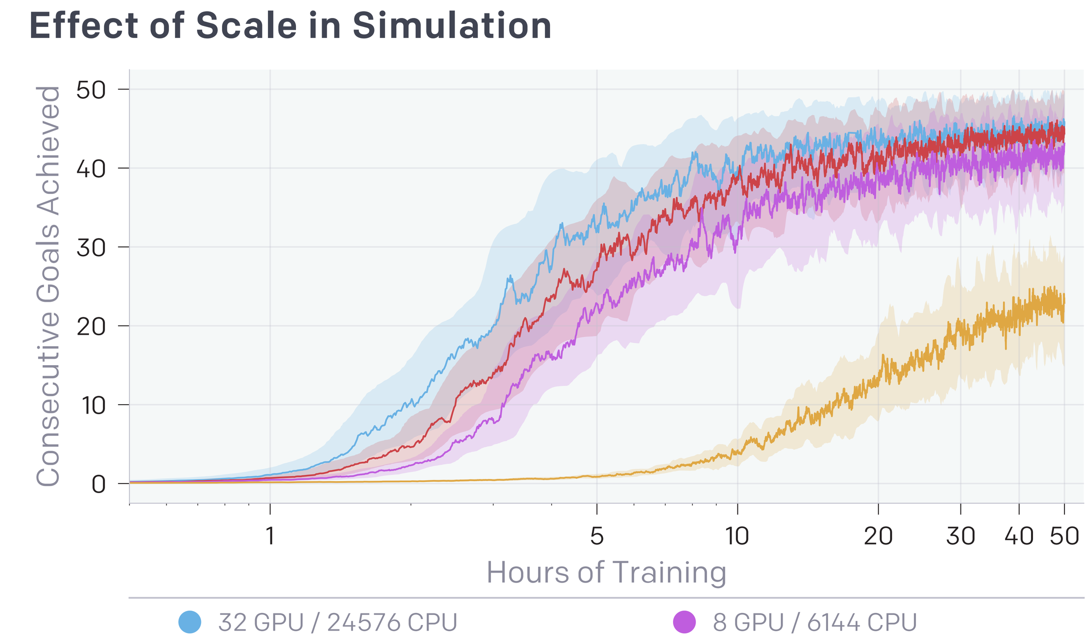

Quantitative Results

Randomizations	Object Tracking	Policy	Number of Successes*		
			Median	Mean	Max
None	Motion Tracking	LSTM	0	1.1 ± 1.9	6
All	Motion Tracking	FF	3.5	4.7 ± 4.1	15
All	Motion Tracking	LSTM	13	18.8 ± 17.1	50
All	Vision	LSTM	11.5	15.2 ± 14.3	46

* Measured across 10 trials. Each trial ends when the block is dropped or if 50 successes are achieved.

Qualitative Results

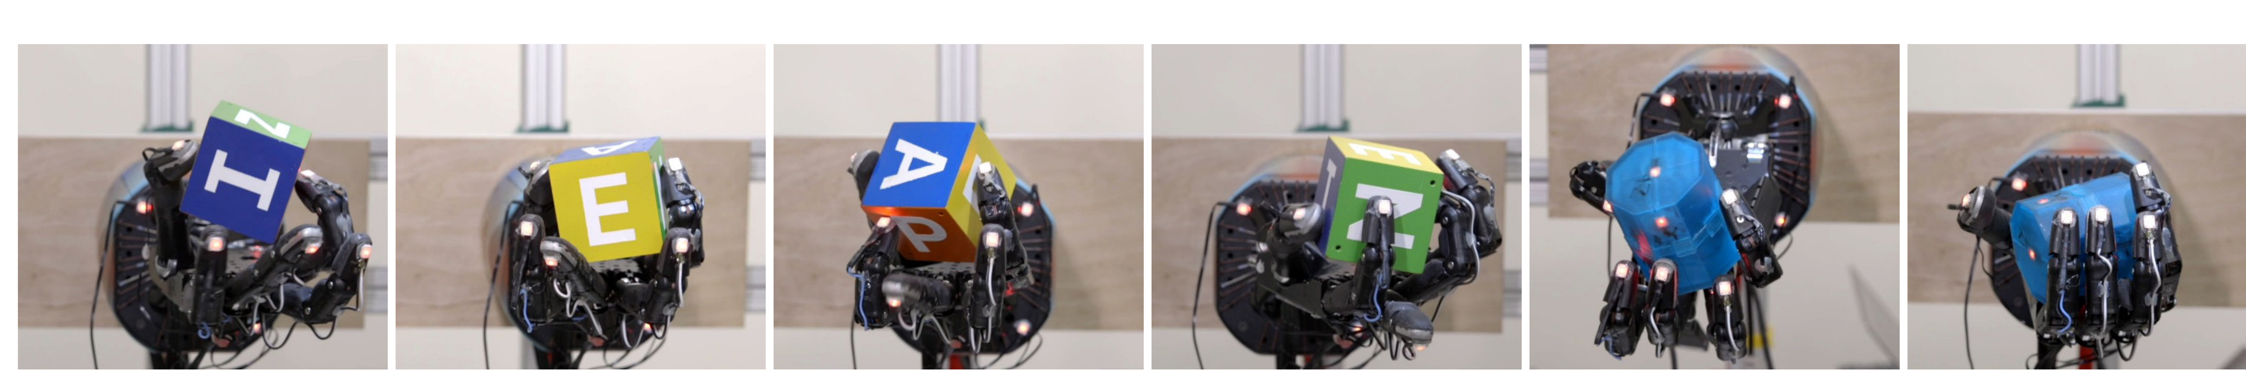

Tip Pinch Grasp | Palmar Pinch Grasp | Tripod Grasp | Quadpod Grasp | 5-finger Precision Grasp | Power Grasp

Classified according to "The GRASP Taxonomy of Human Grasp Types", Feix et al., 2016.

Learn More

https://blog.openai.com/learning-dexterity/

https://arxiv.org/abs/1808.00177

https://youtu.be/jwSbzNHGfIM