Self-Adaptive Manipulation

Rethinking Embodiments in Embodied Intelligence

Shuran Song

Columbia University
Artificial Intelligence & Robotics Lab

Self-Adaptive Manipulation

Rethinking Embodiments in Embodied Intelligence

Shuran Song

Columbia University
Artificial Intelligence & Robotics Lab

Embodied Intelligence

Develop embodied agents that is able to acquire intelligent behaviors through interacting with their environments.

Navigation

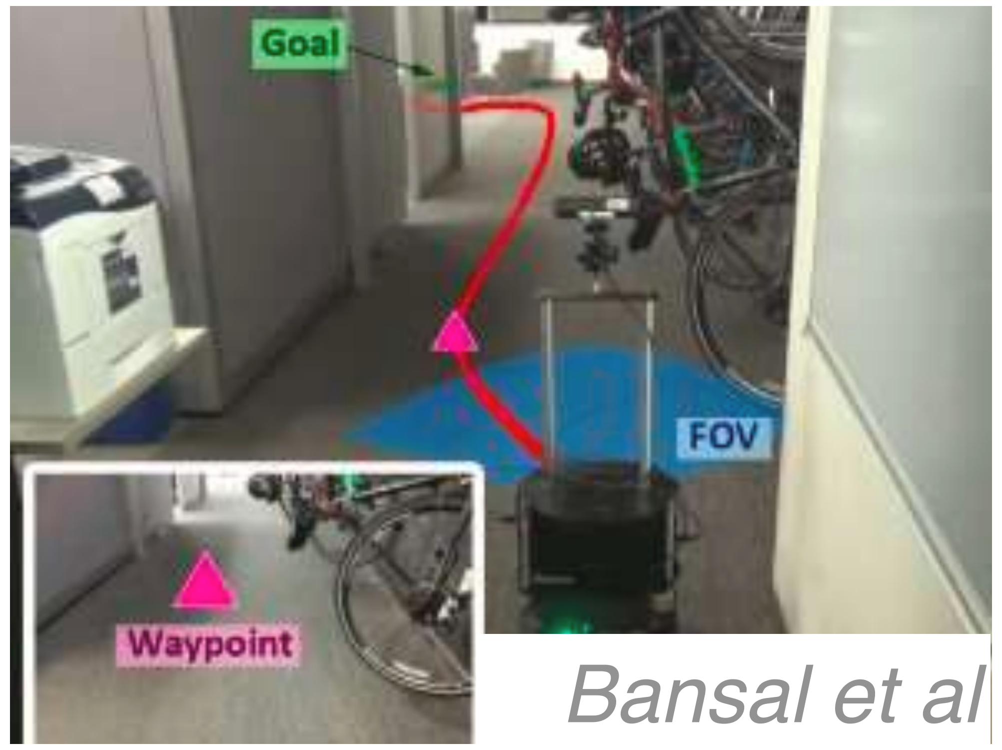

Bansal et al

Object Searching

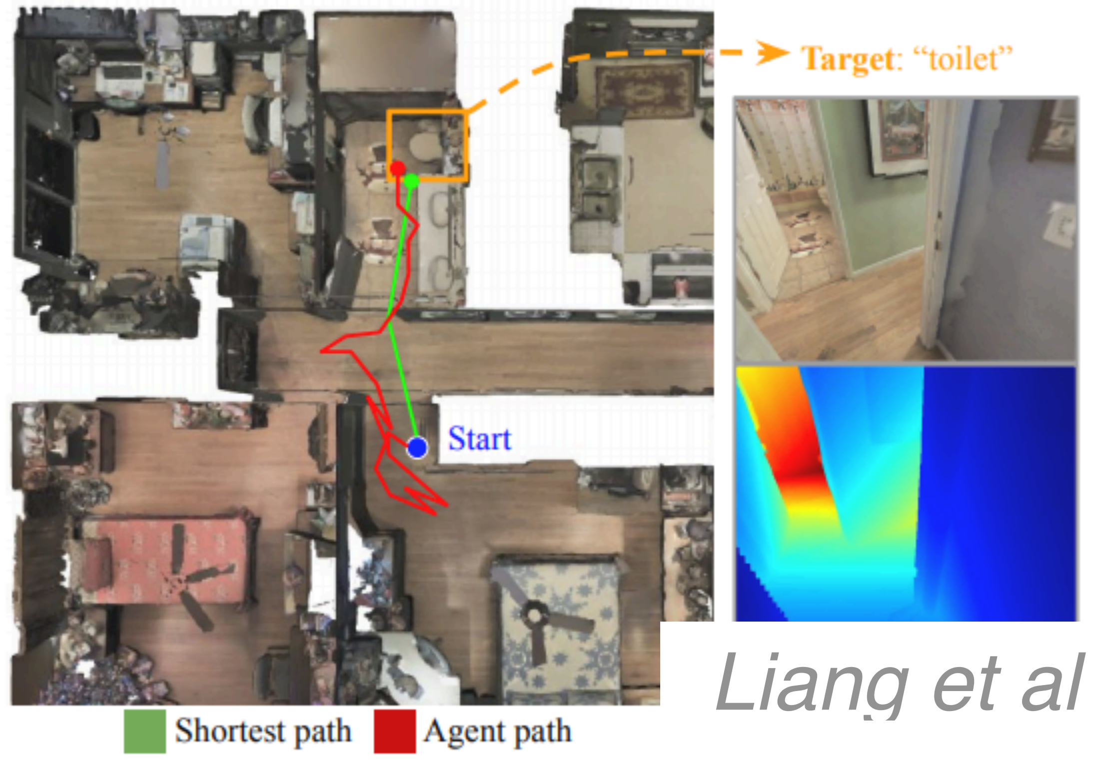

Liang et al

Question Answering

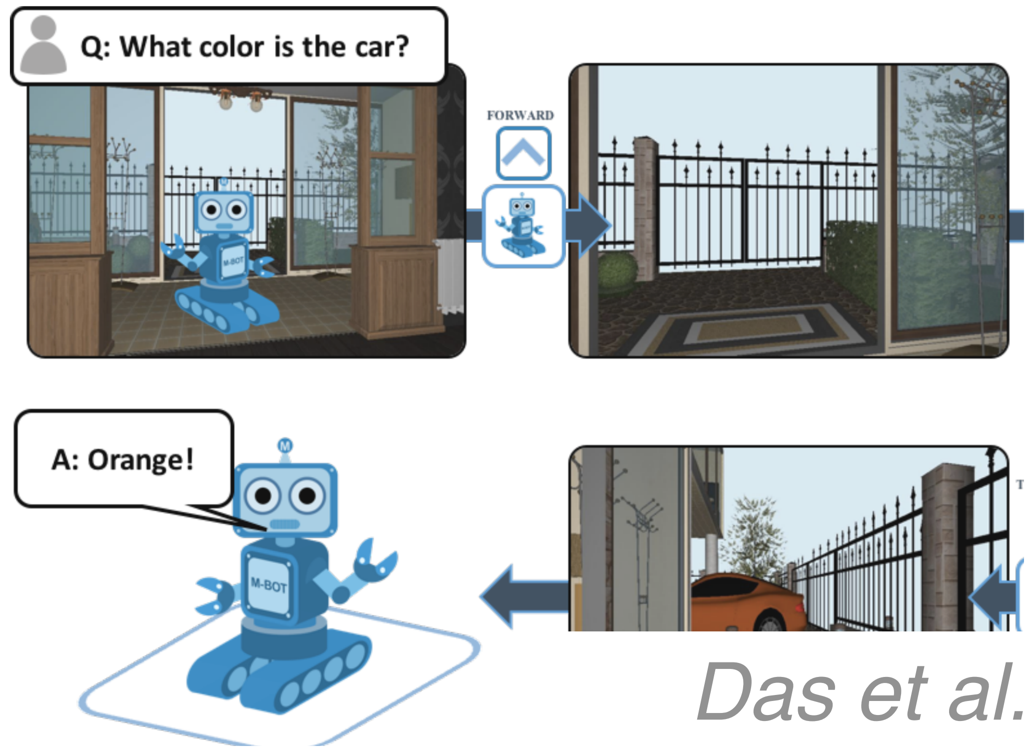

Das et al.

Manipulation

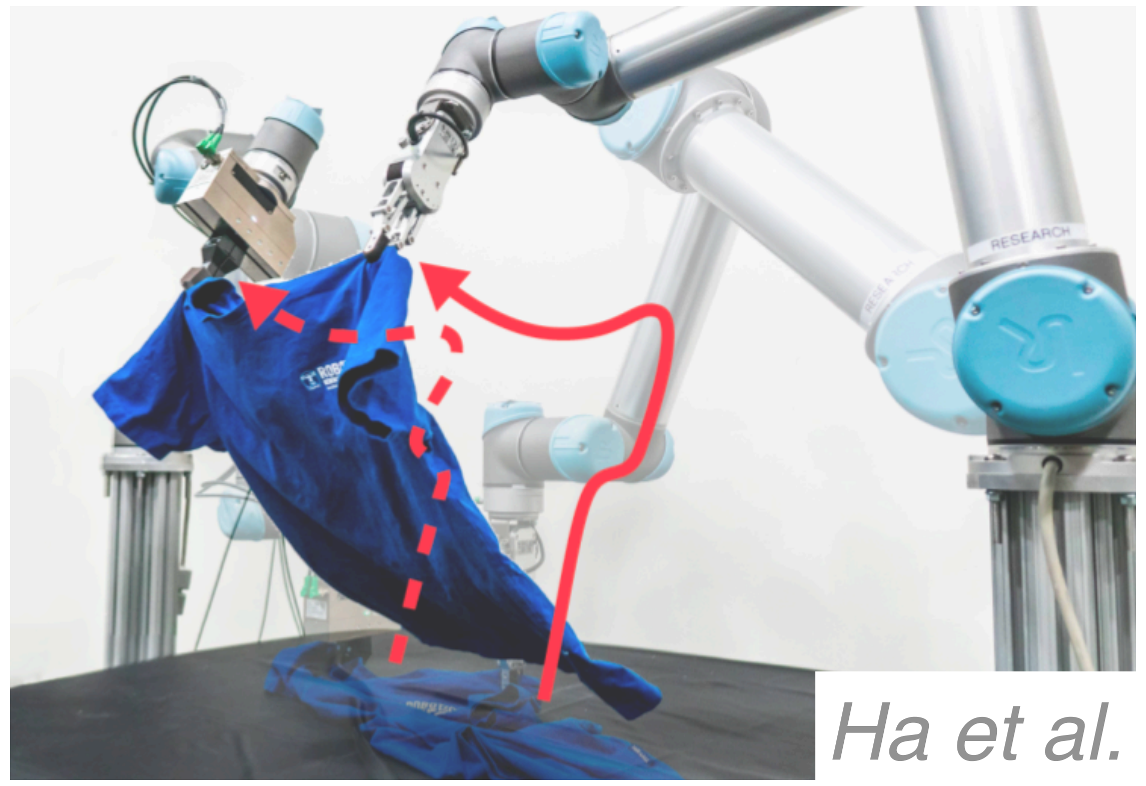

Ha et al.

Multi-agent Collaboration

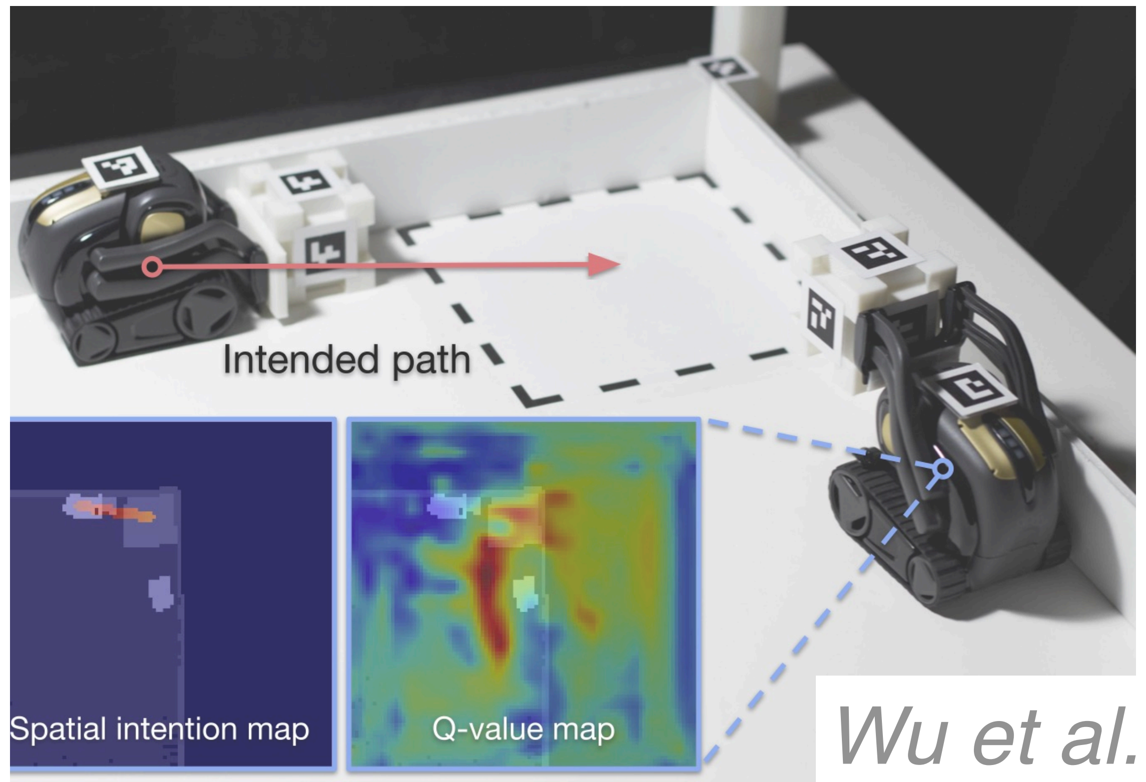

Wu et al.

Motion Planning

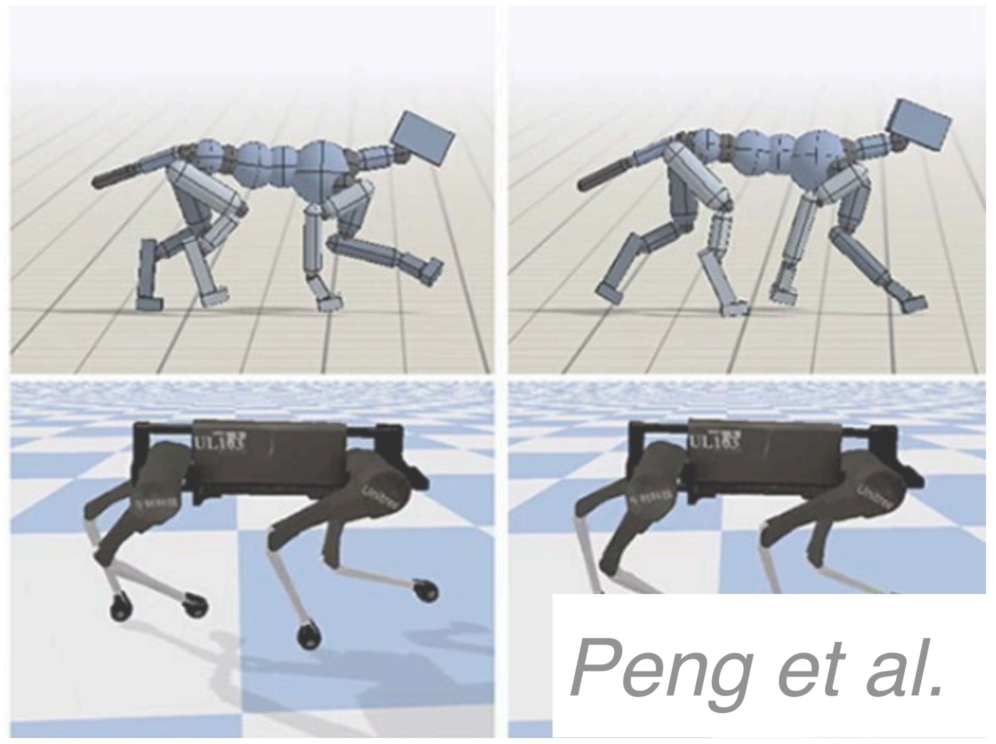

Peng et al.

Embodied Intelligence

Problem formulation

Input: Observation

$$f(\color{red}{x}) \rightarrow y$$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Embodied Intelligence

Problem formulation

X: External environment, objects

Input: Observation

$f(x) \rightarrow y$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Embodied Intelligence

What's missing?

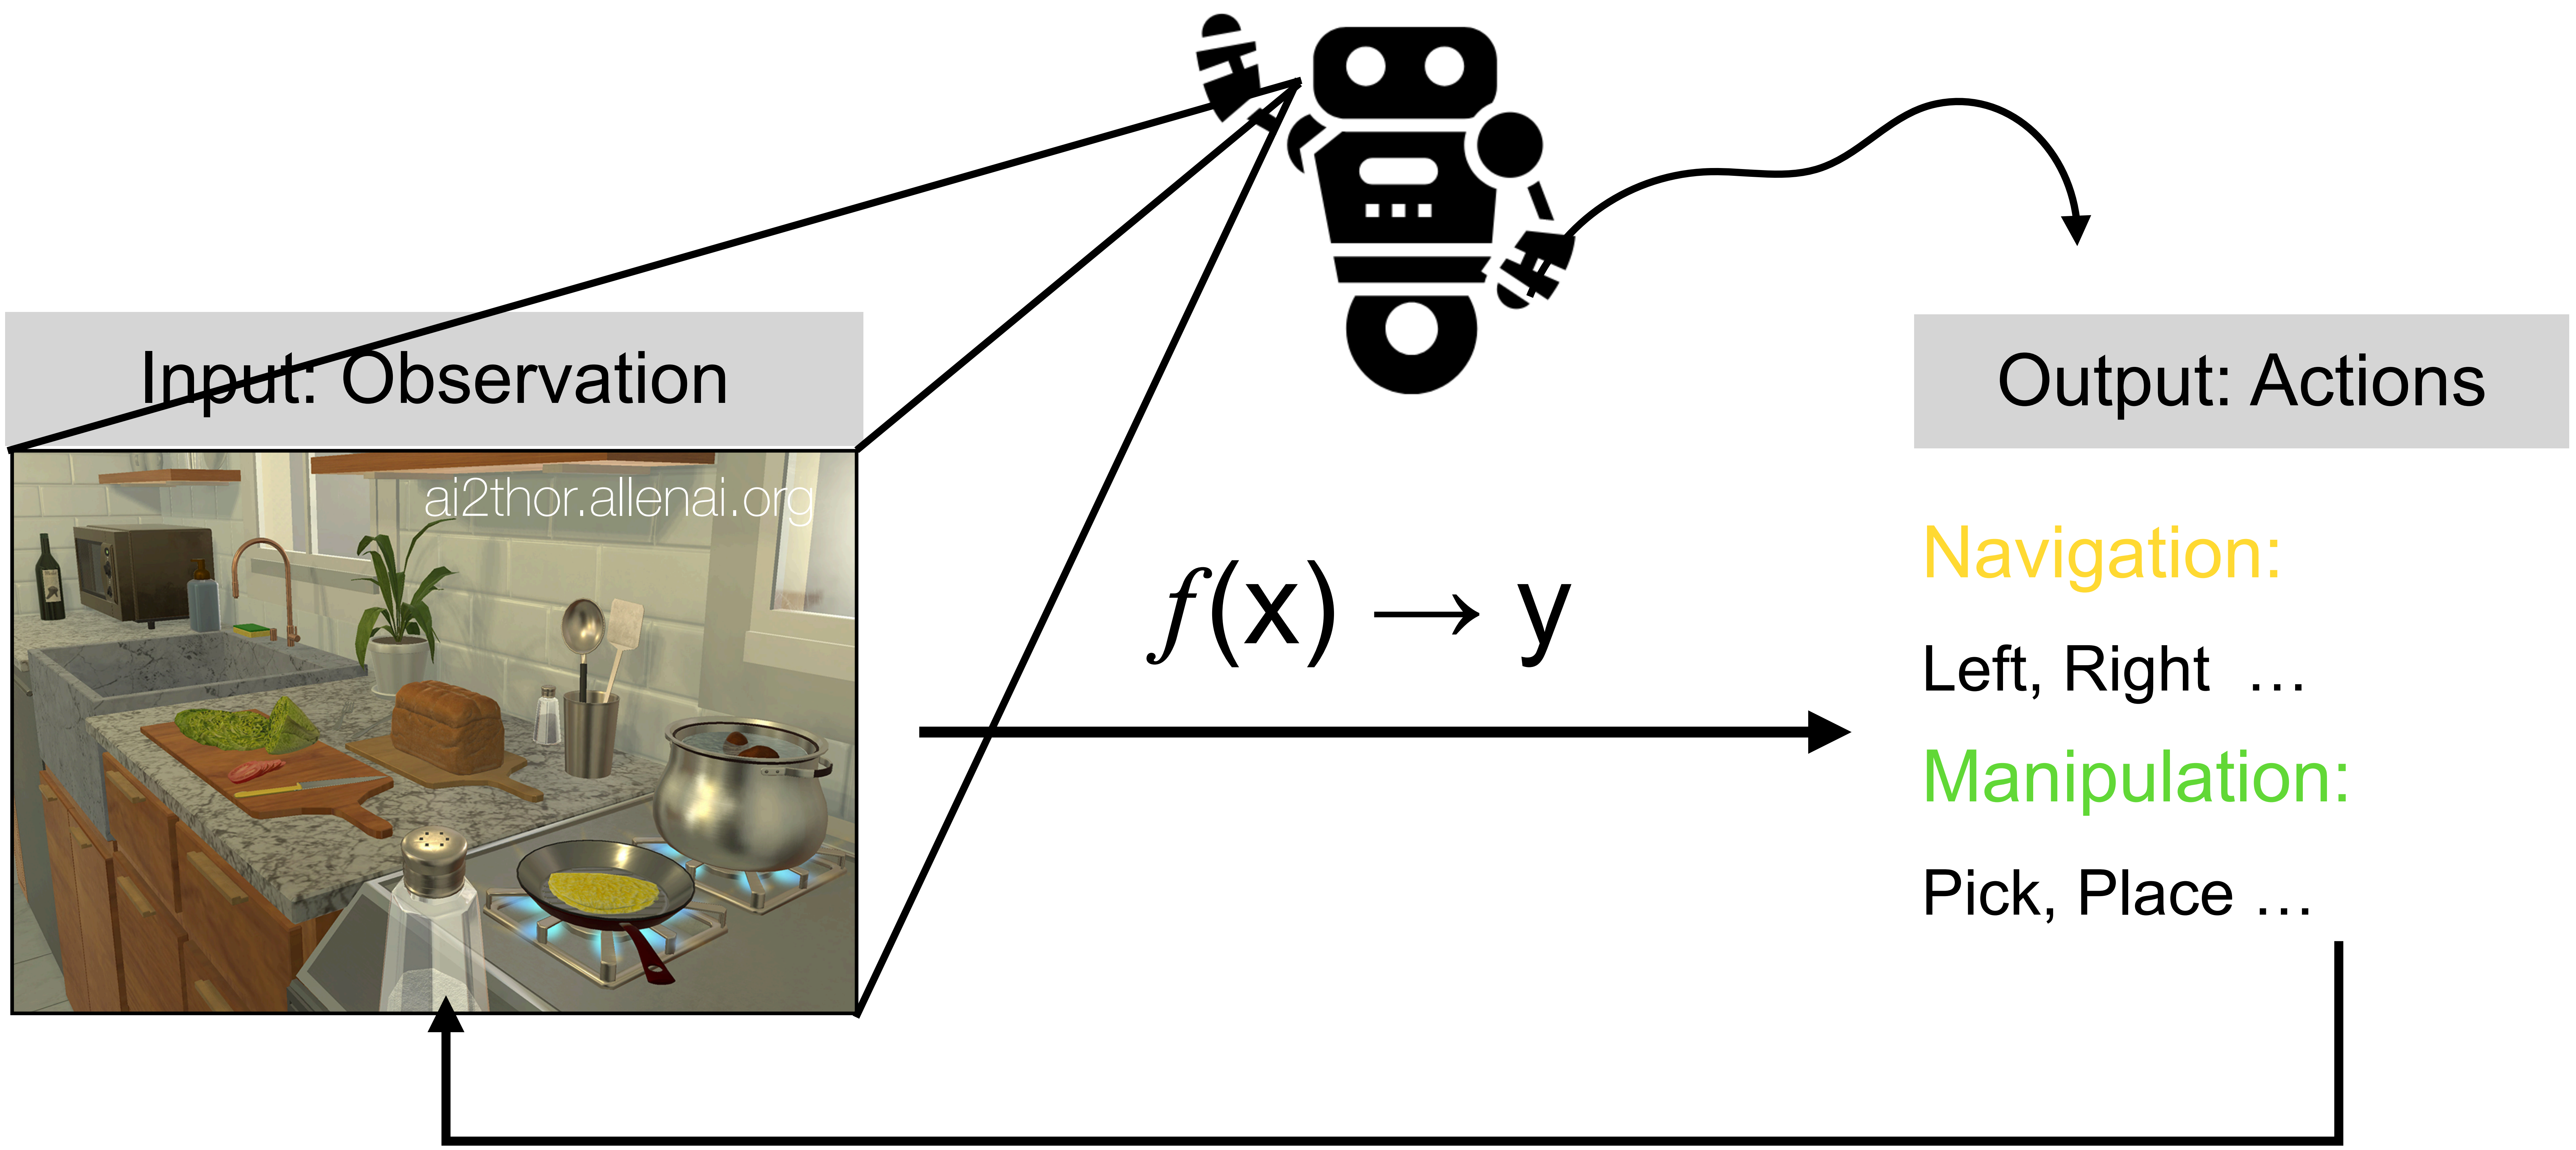

Physical Embodiment of the Agent

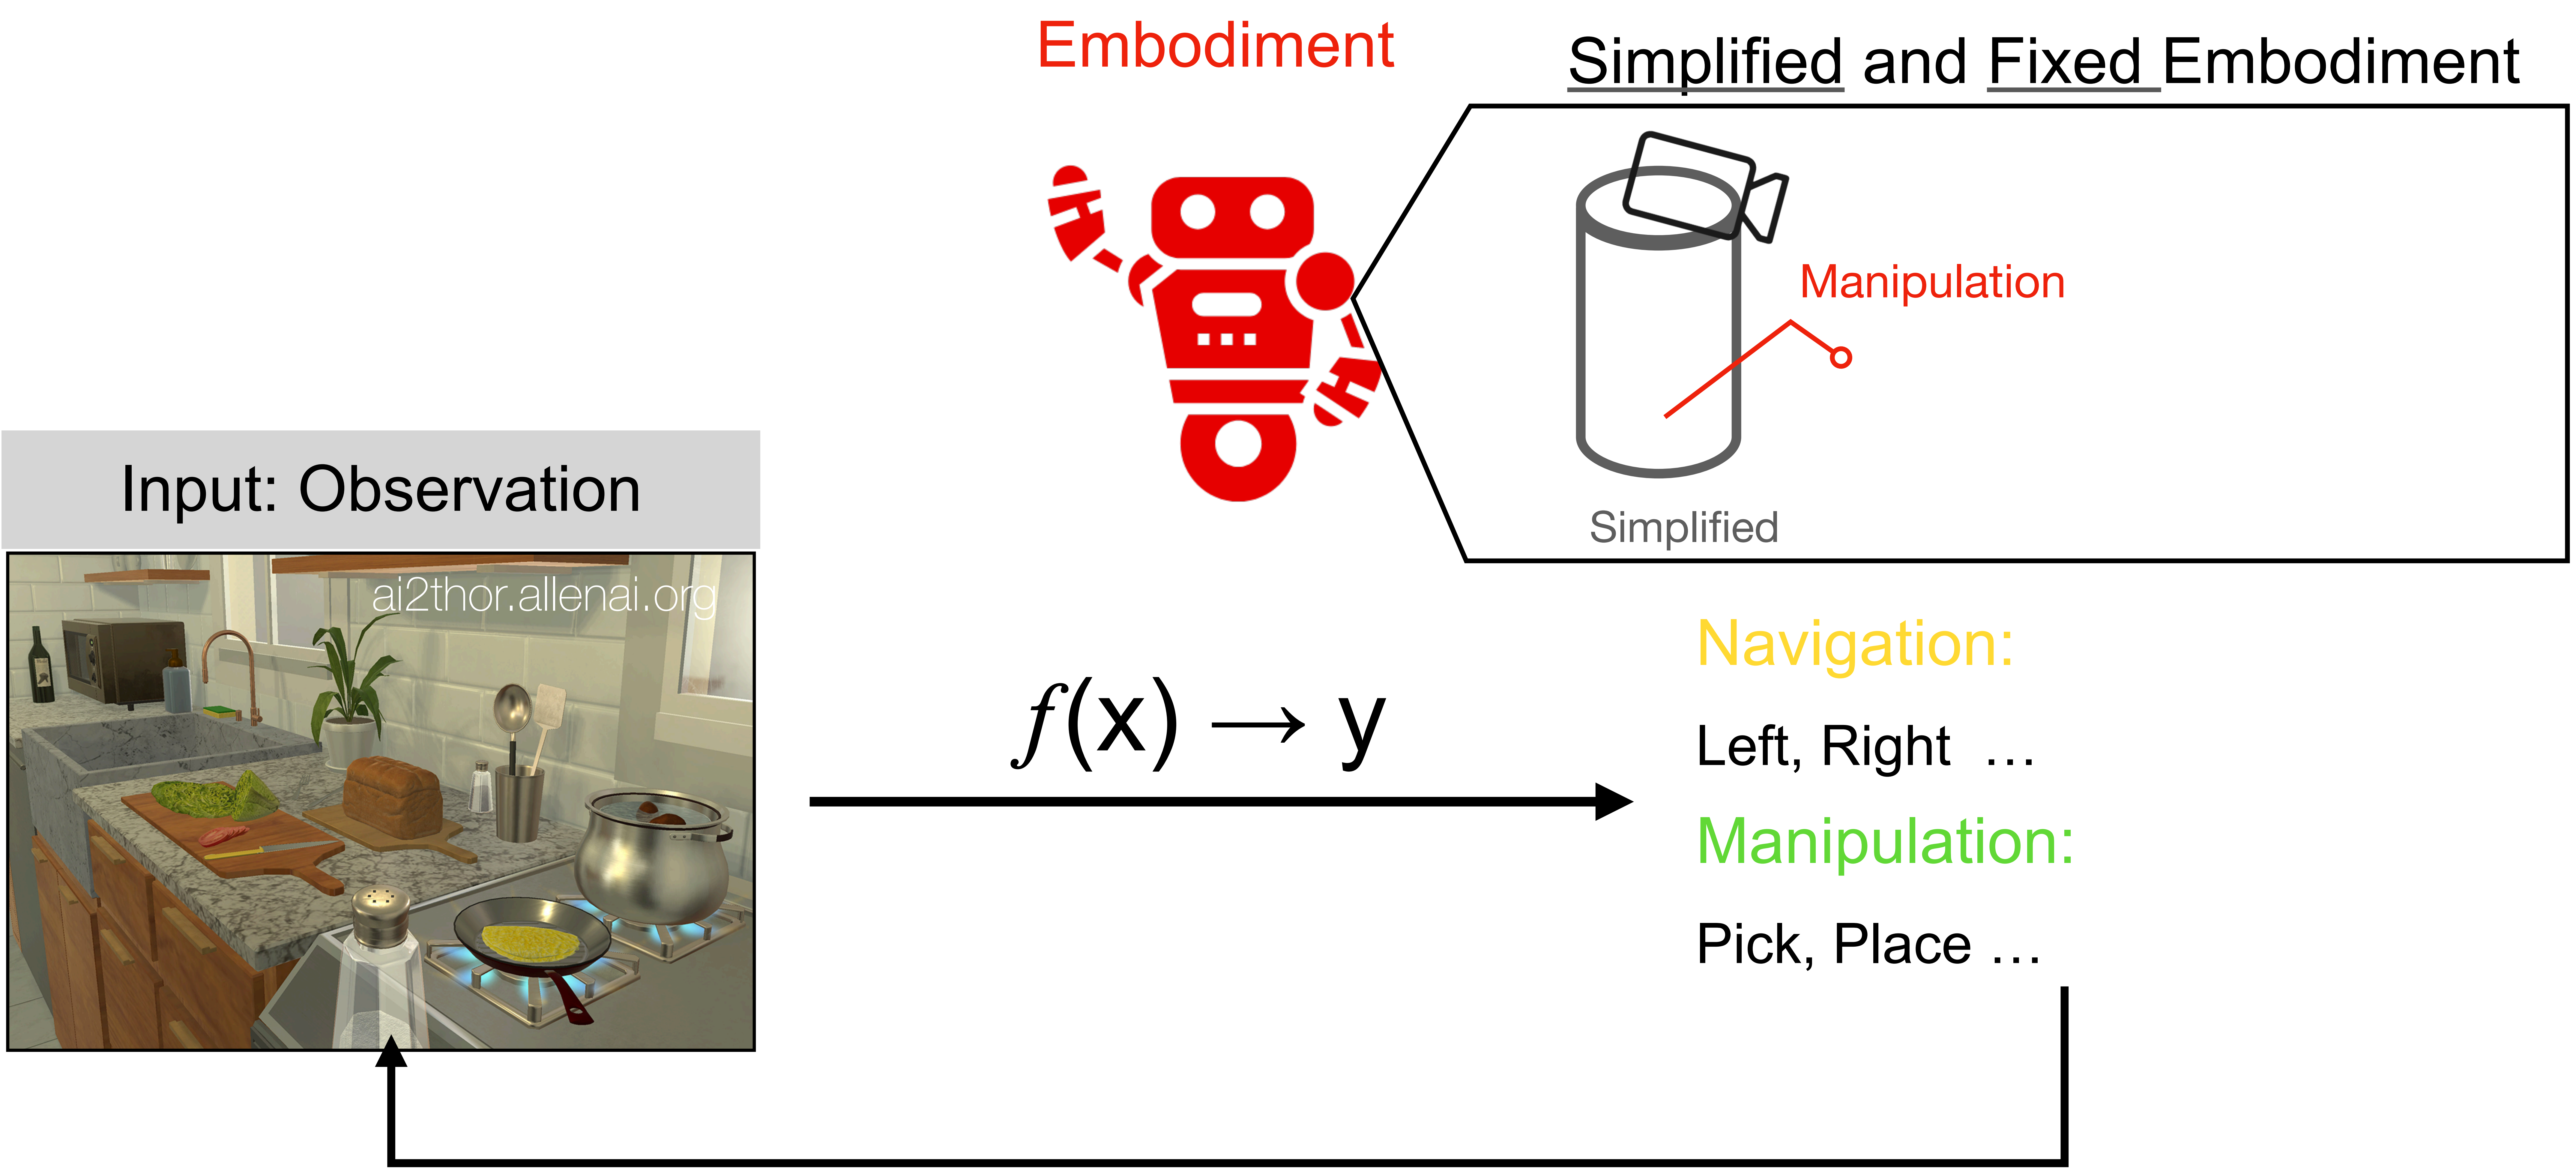

Physical Embodiment of the Agent

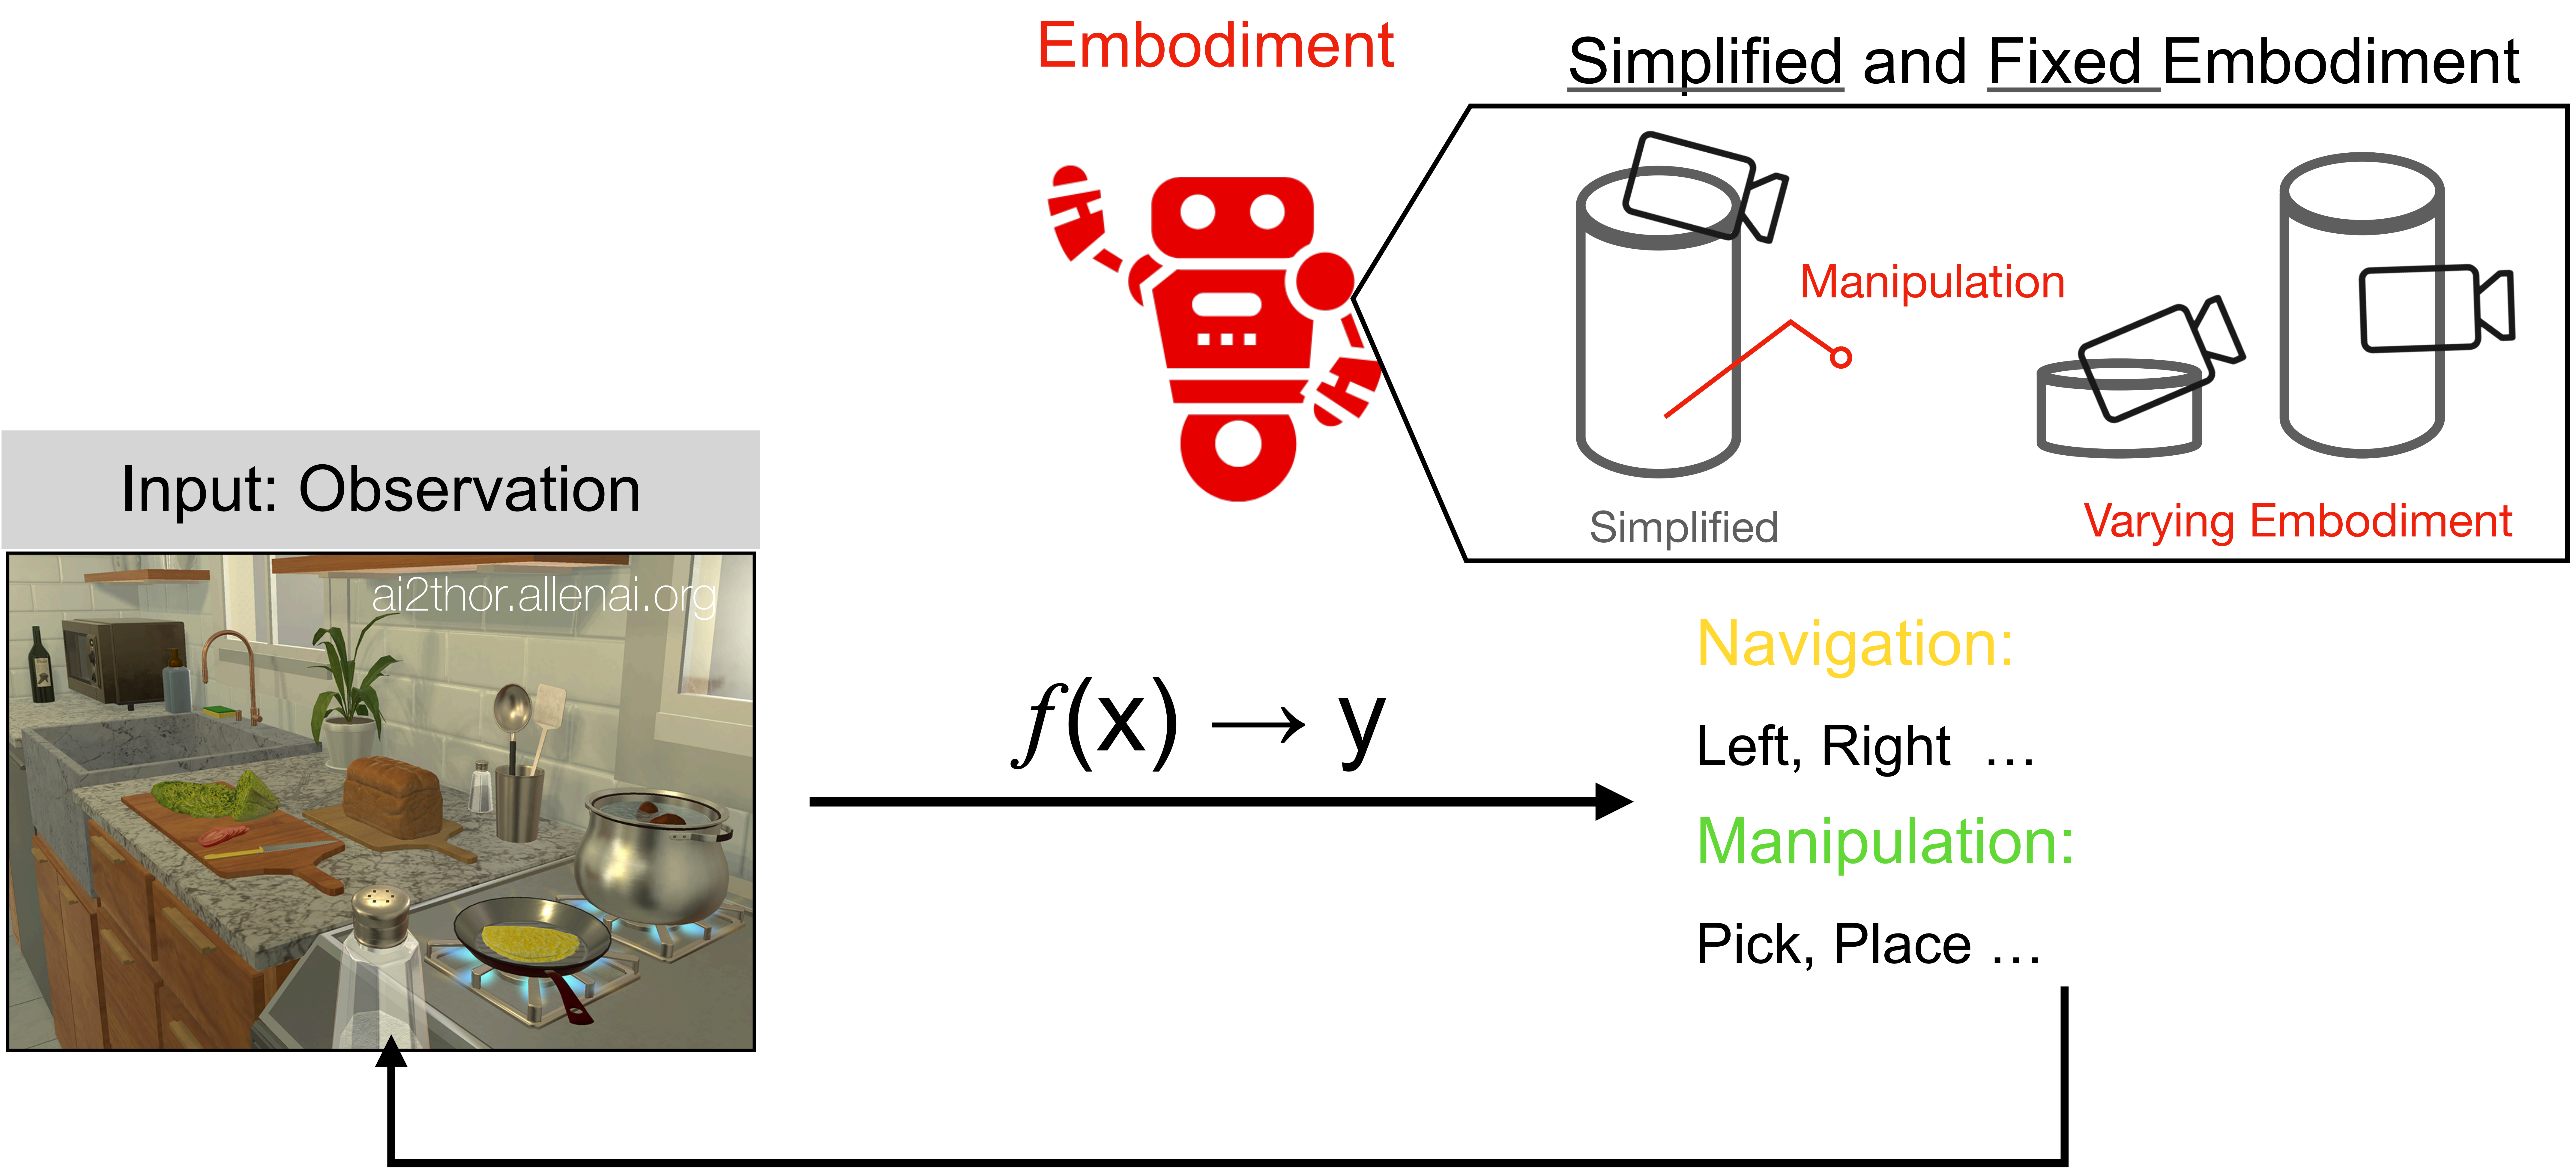

Physical Embodiment of the Agent

Why it matters?

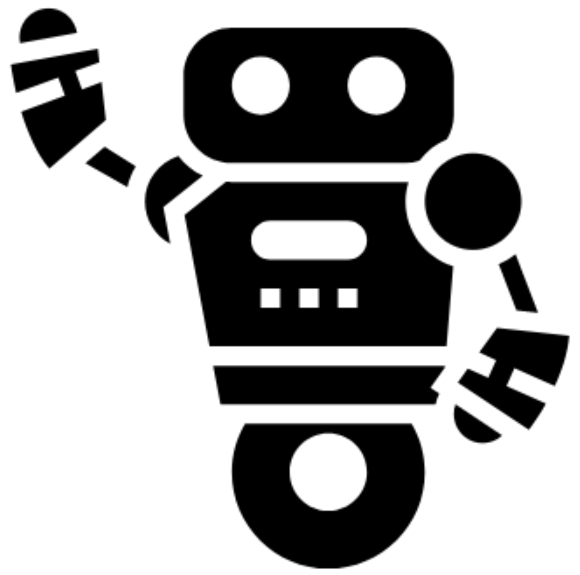

Input: Observation

$f(x) \rightarrow y$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Physical Embodiment of the Agent

Why it matters ?

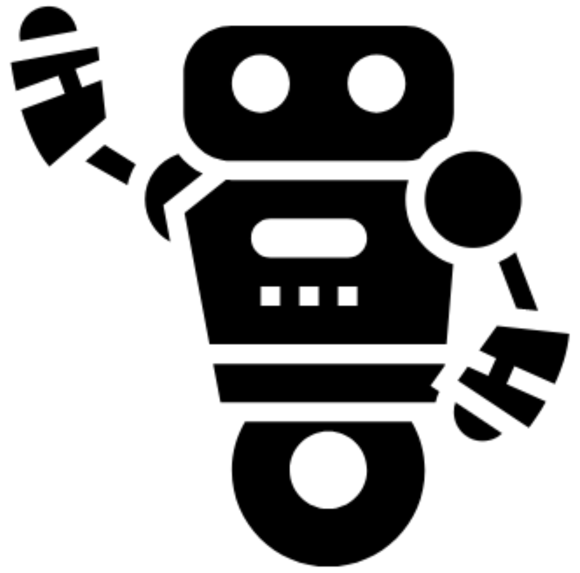

The agent's embodiment can directly influence the sensory input and action space

Input: Observation

$f(x) \rightarrow y$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Physical Embodiment of the Agent

Why it matters ?

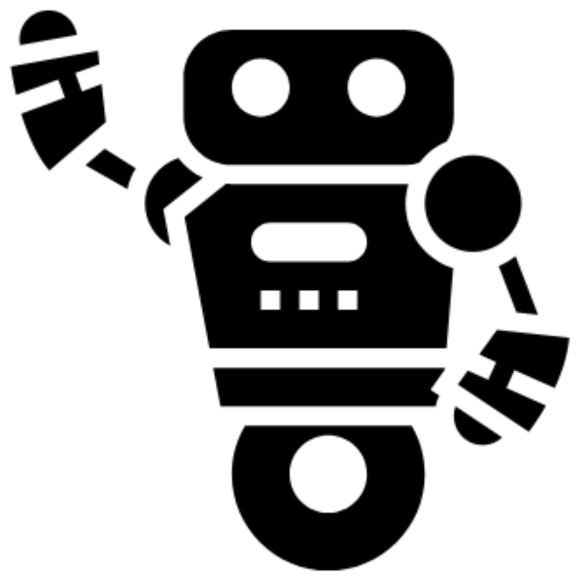

The agent's embodiment can directly influence the sensory input and action space

Input: Observation

Output: Actions

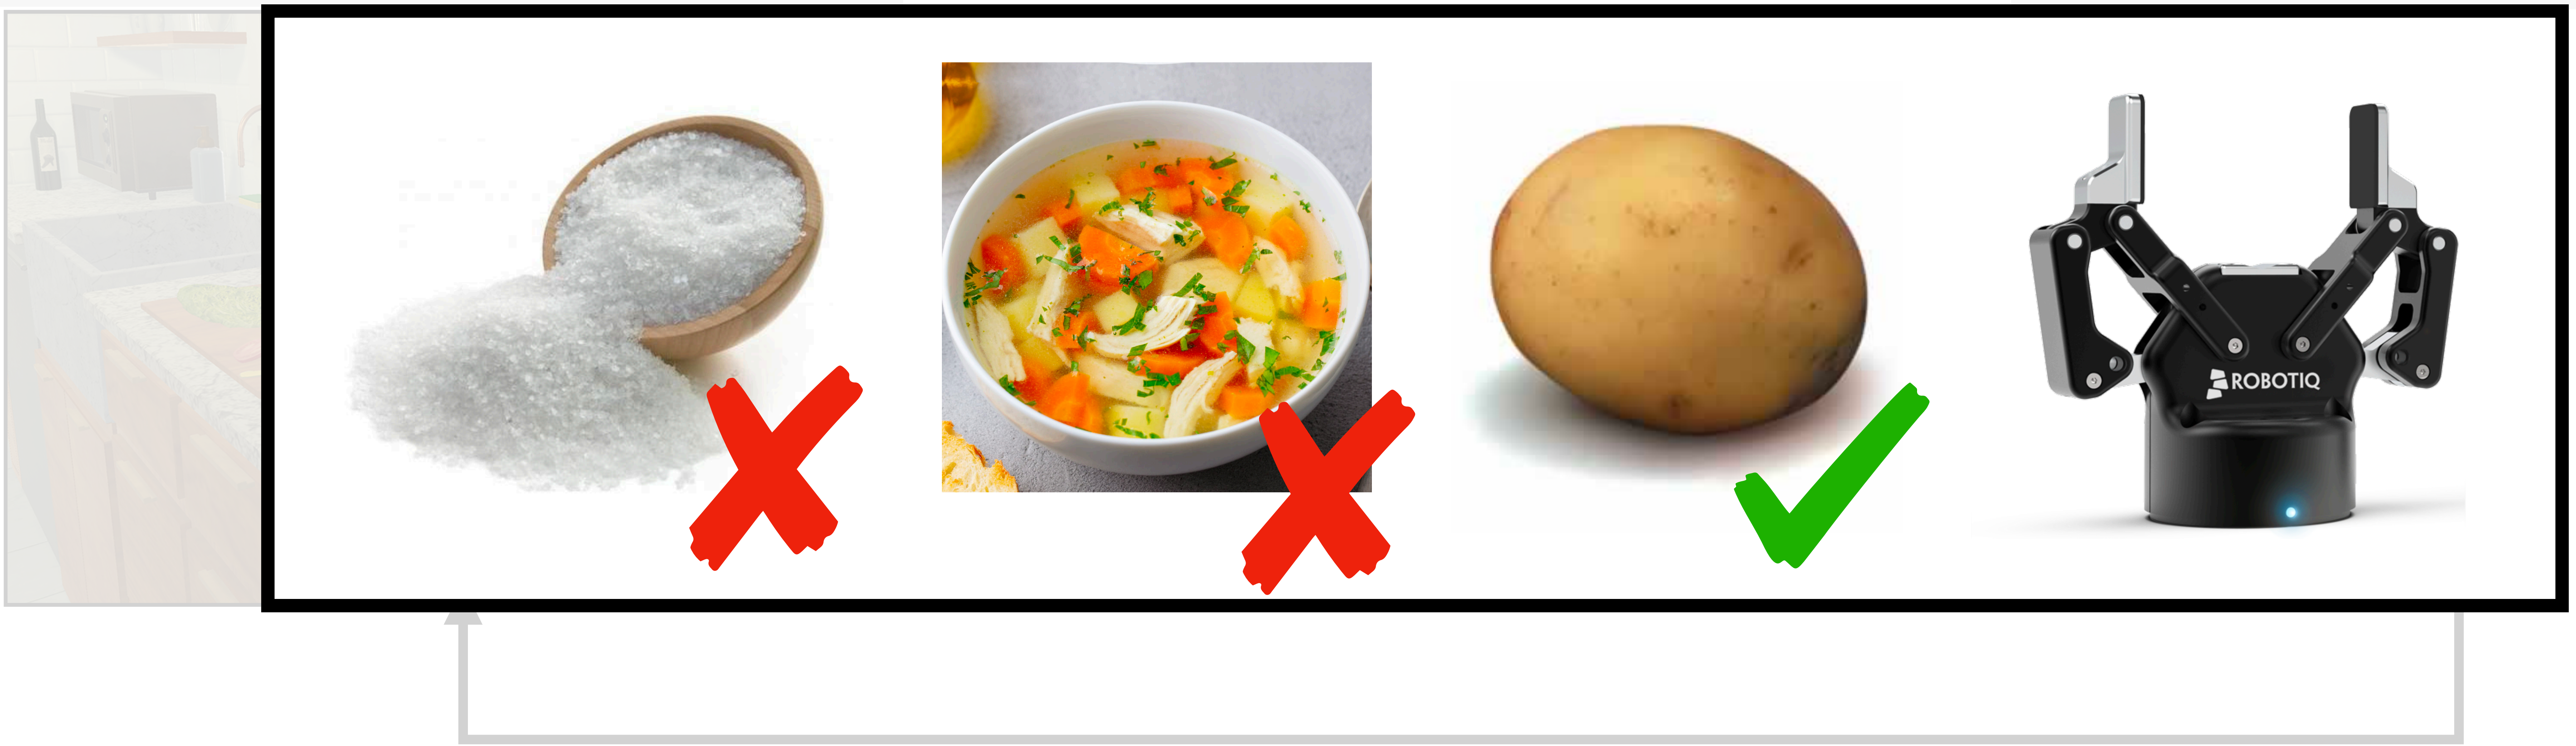

Physical Embodiment of the Agent

Why it matters ?

System's functionality should also influence the agent's embodiment

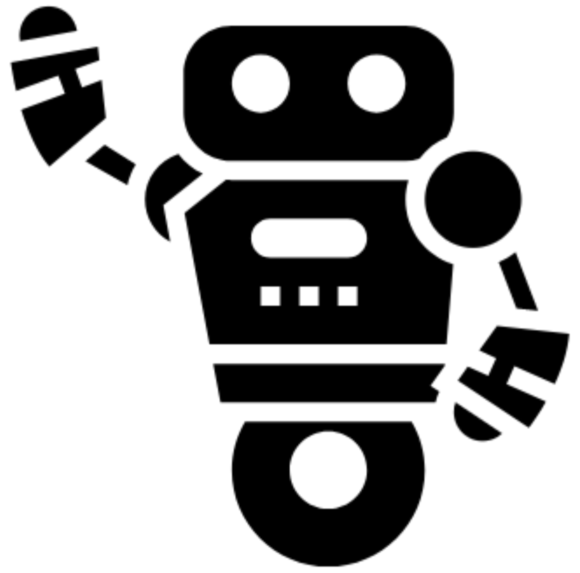

Input: Observation

$$f(x) \rightarrow y$$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Physical Embodiment of the Agent

Why it matters ?

System's functionality should also influence the agent's embodiment

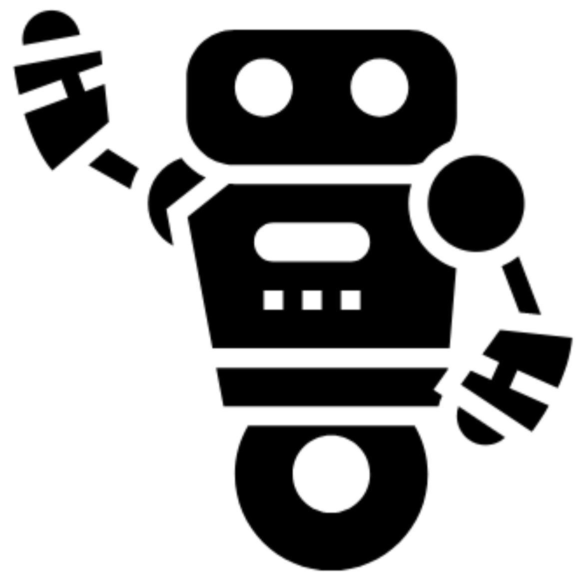

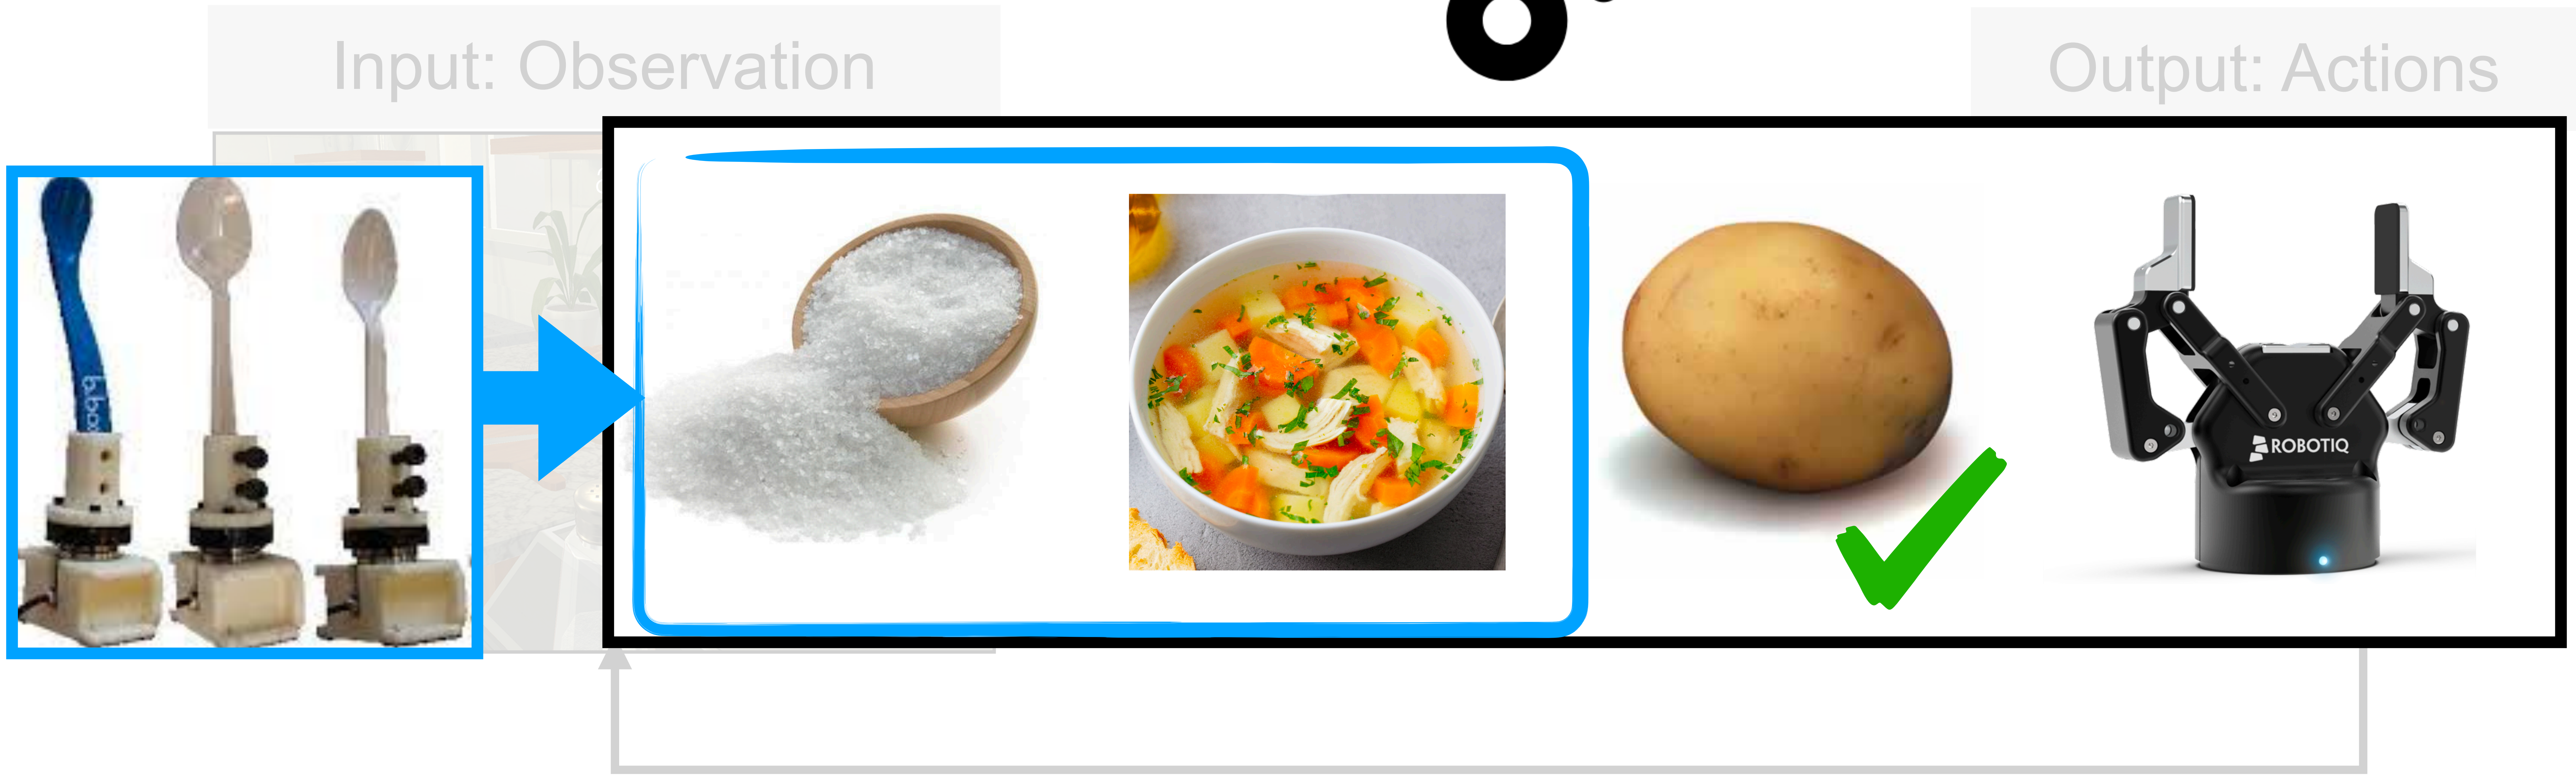

Physical Embodiment of the Agent

Embodiment

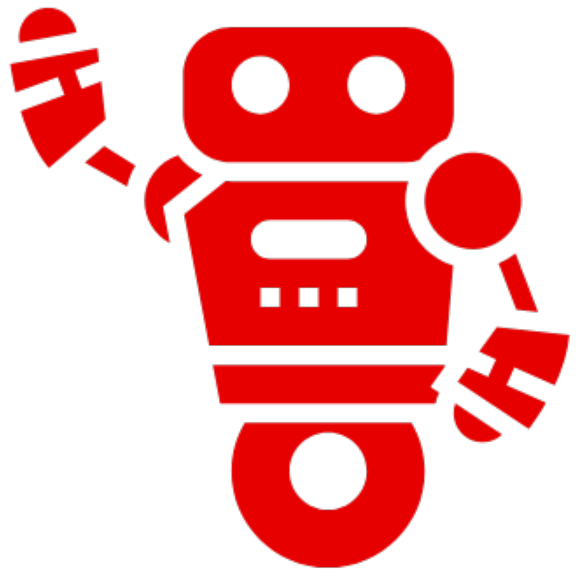

Input: Observation

$f(x) \rightarrow y$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Physical Embodiment of the Agent

Embodiment

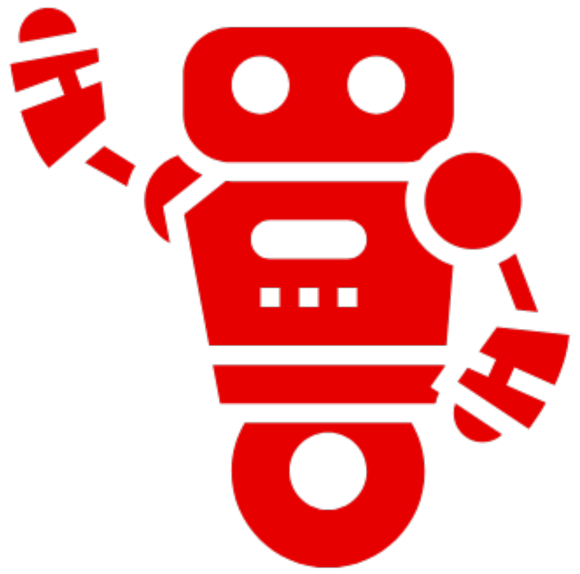

Input: Observation

$$f(x\ ,\ b) \rightarrow y$$

Output: Actions

Navigation:

Left, Right …

Manipulation:

Pick, Place …

Self-Adaptive Manipulation

Self-Adaptive Manipulation

Treating a robot's hardware as an integral part of its behavior —

the learned manipulation policies should be conditioned on the agent's hardware

also inform how hardware should be improved.

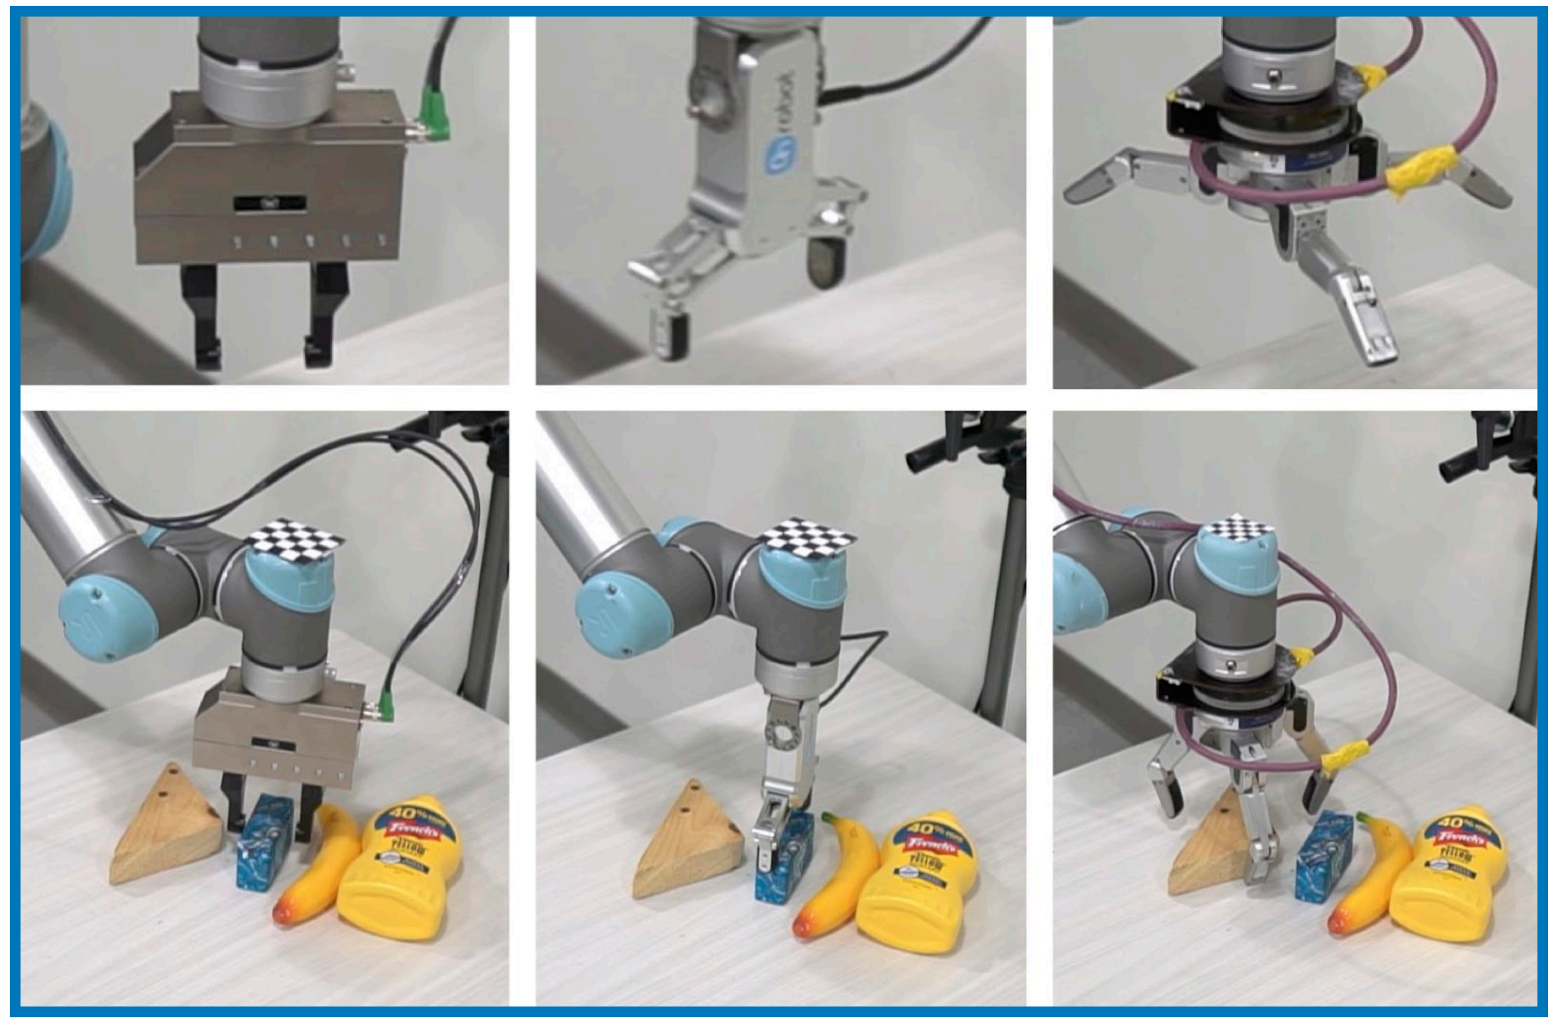

Adagrasp *(ICRA'21)*

https://adagrasp.cs.columbia.edu/

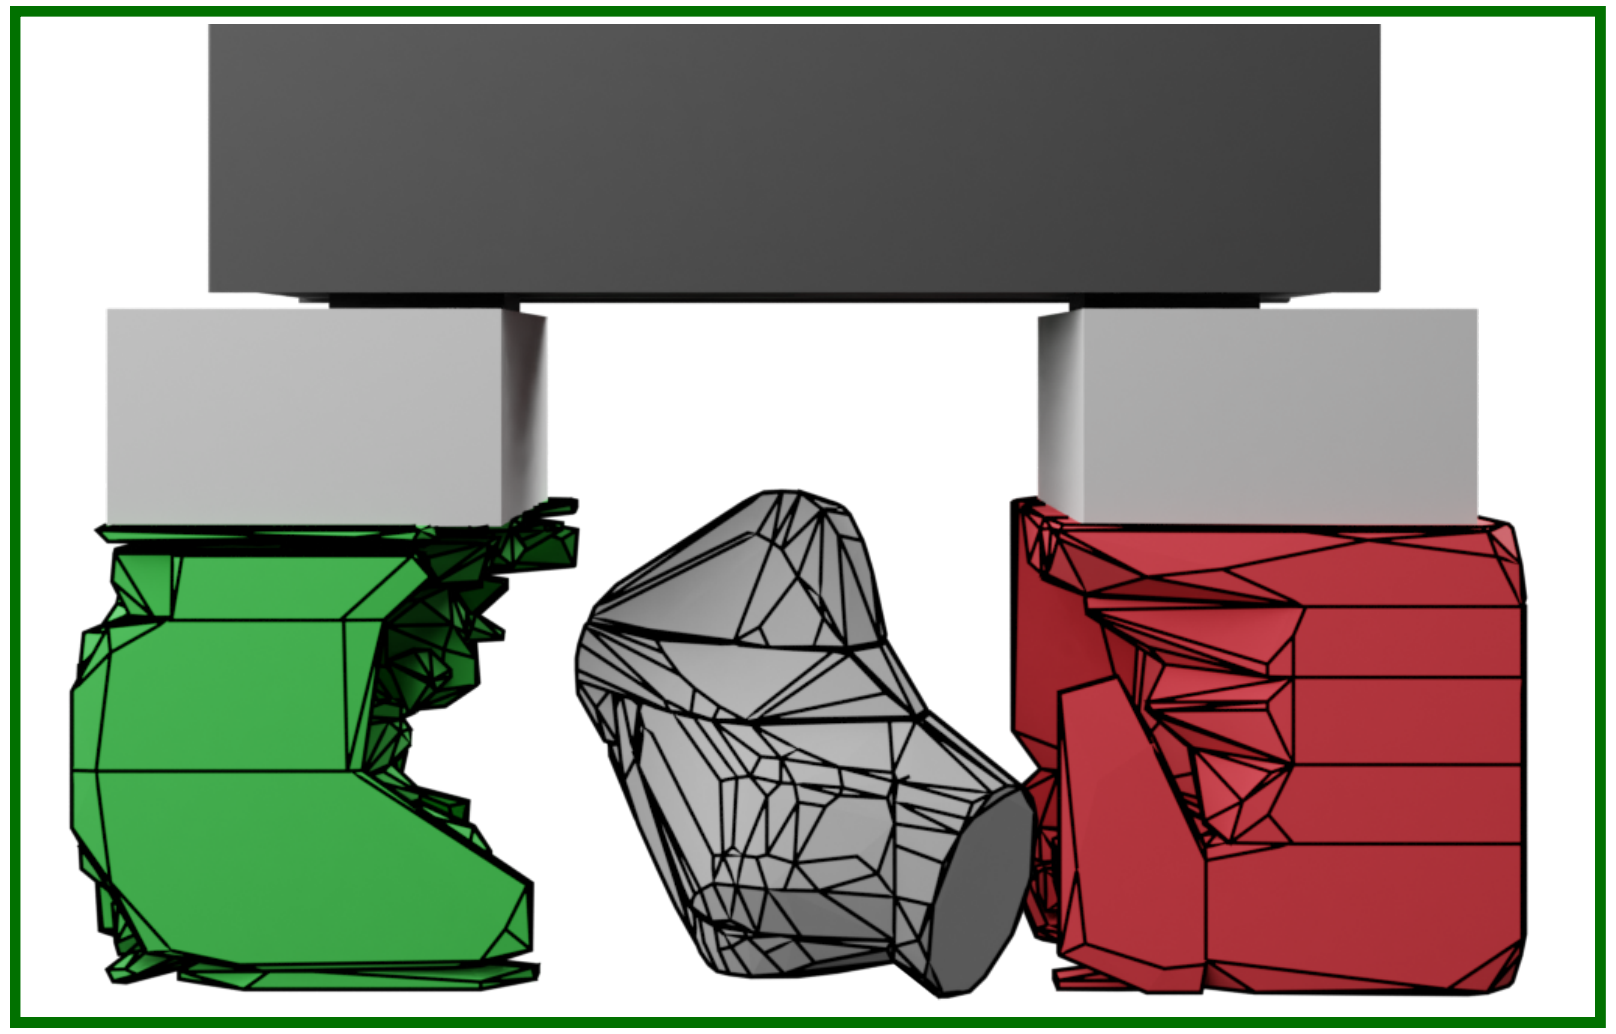

Fit2Form *(CoRL'20)*

https://fit2form.cs.columbia.edu/

Self-Adaptive Manipulation

Treating a robot’s hardware as an integral part of its behavior —

the learned manipulation policies should be conditioned on the agent’s hardware

also inform how hardware should be improved.

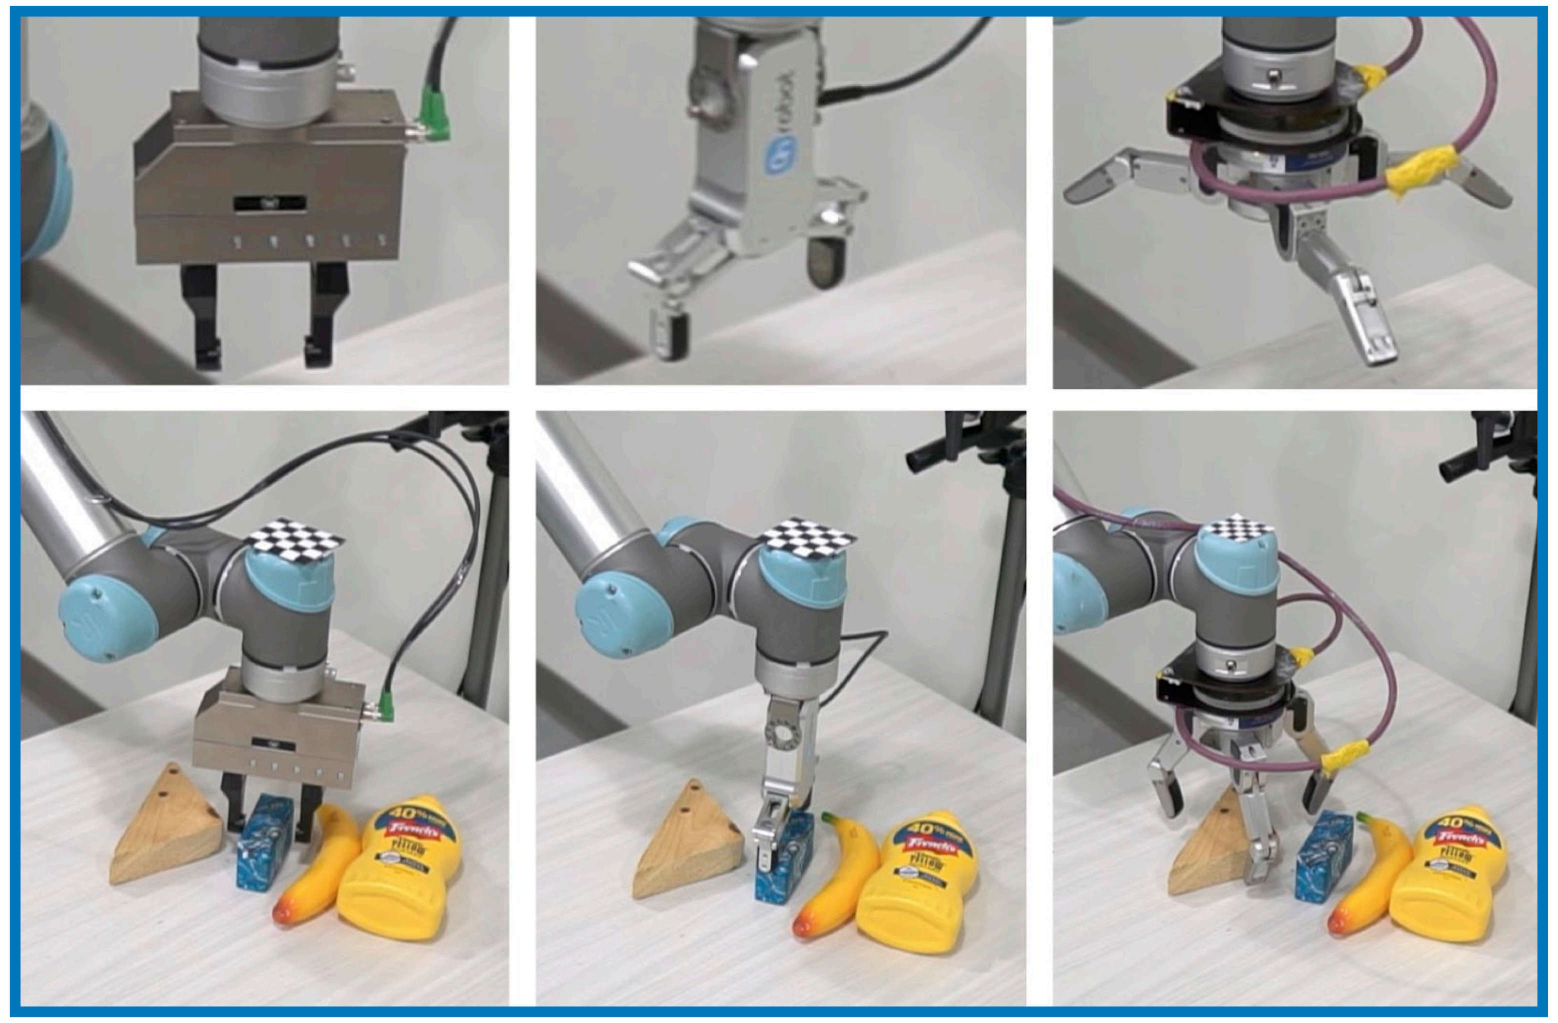

Adagrasp *(ICRA’21)*

AdaGrasp: Learning an Adaptive Gripper-Aware Grasping Policy
ICRA 2021

Zhenjia Xu

Beichun Qi

Shubham Agrawal

Grasping Algorithms Today …

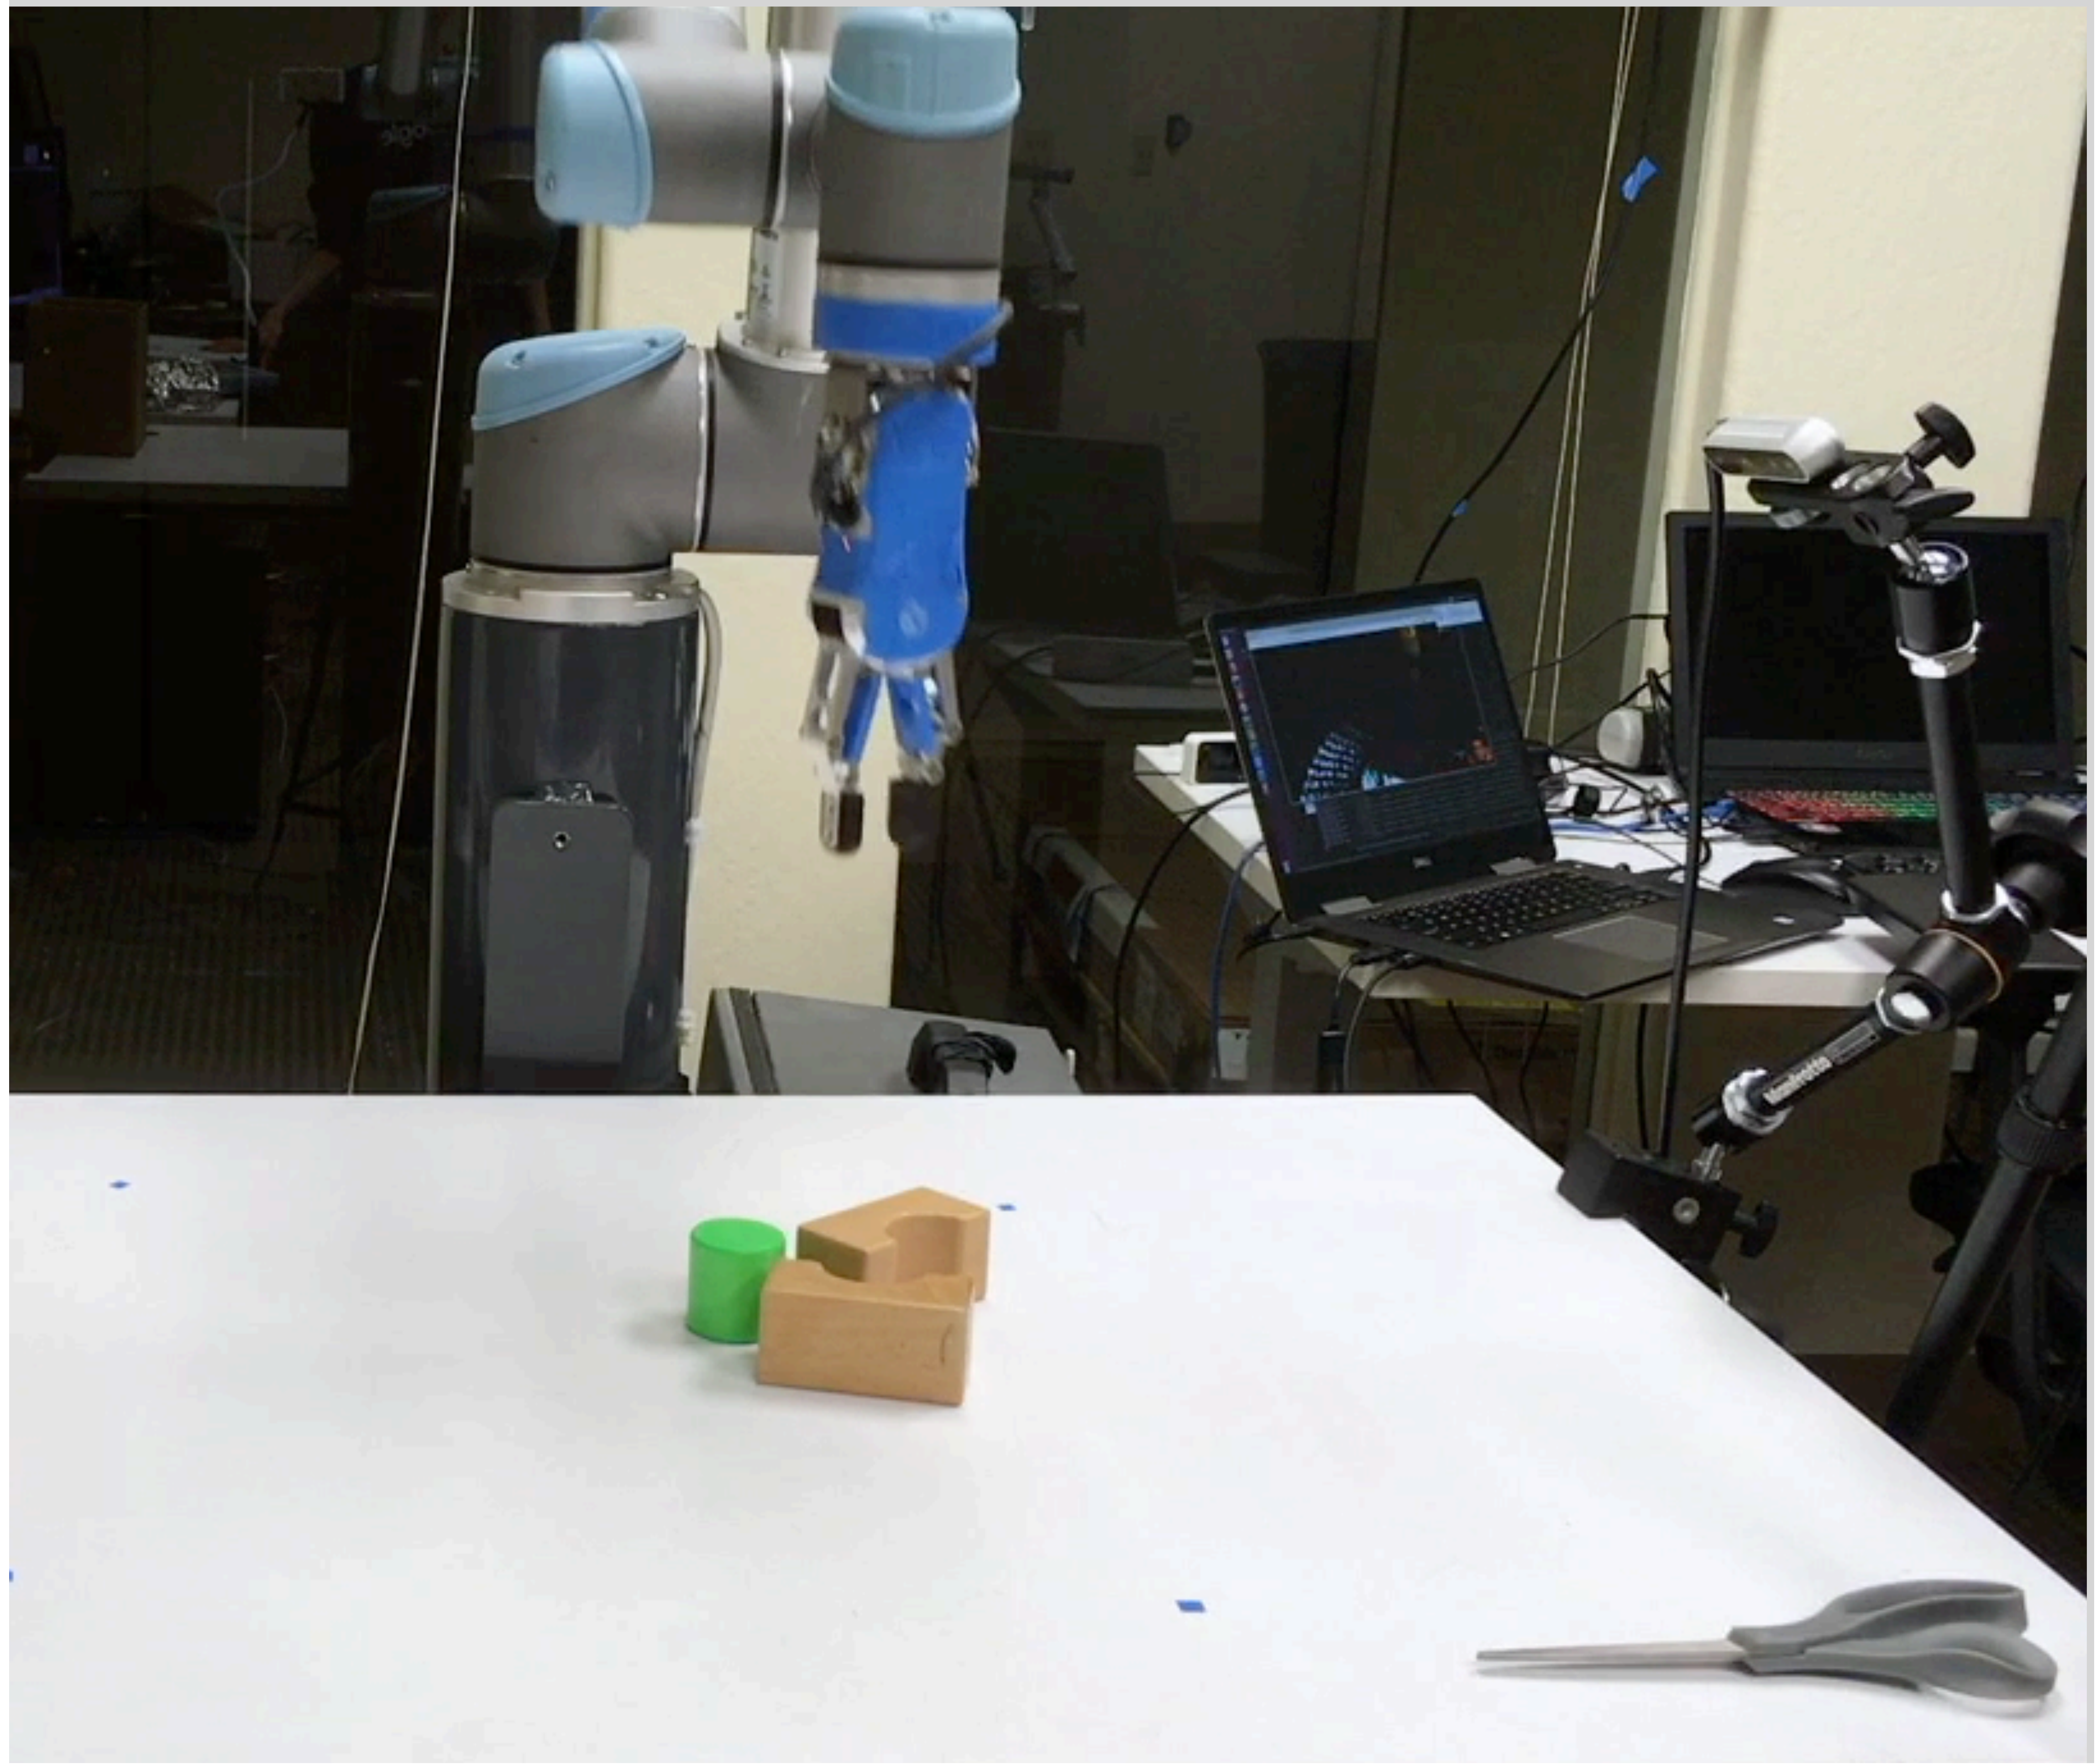

$f(\mathrm{x}) \rightarrow$ action

Prior works in grasping

Mahler et al. ICRA. 2016 DexNet 1.0
Mahler et al. RSS,2017 DexNet 2.0
Morrison et al. RSS, 2018.
Pinto and Gupta, ICRA 2016
Prattichizzo and Trinkle, Springer Handbook of Robotics. 2008
Redmon and Angelova ICRA, 2015
Kalashnikov et al, CORL 2018
Goldfeder et al 2009
Gualtieri and Platt ICRA 2018
Gualtieri et al IROS 2016
Levine et al IJRR 2018
Viereck et al CORL 2017
Weisz and Allen ICRA 2012

Grasping Algorithms Today ...

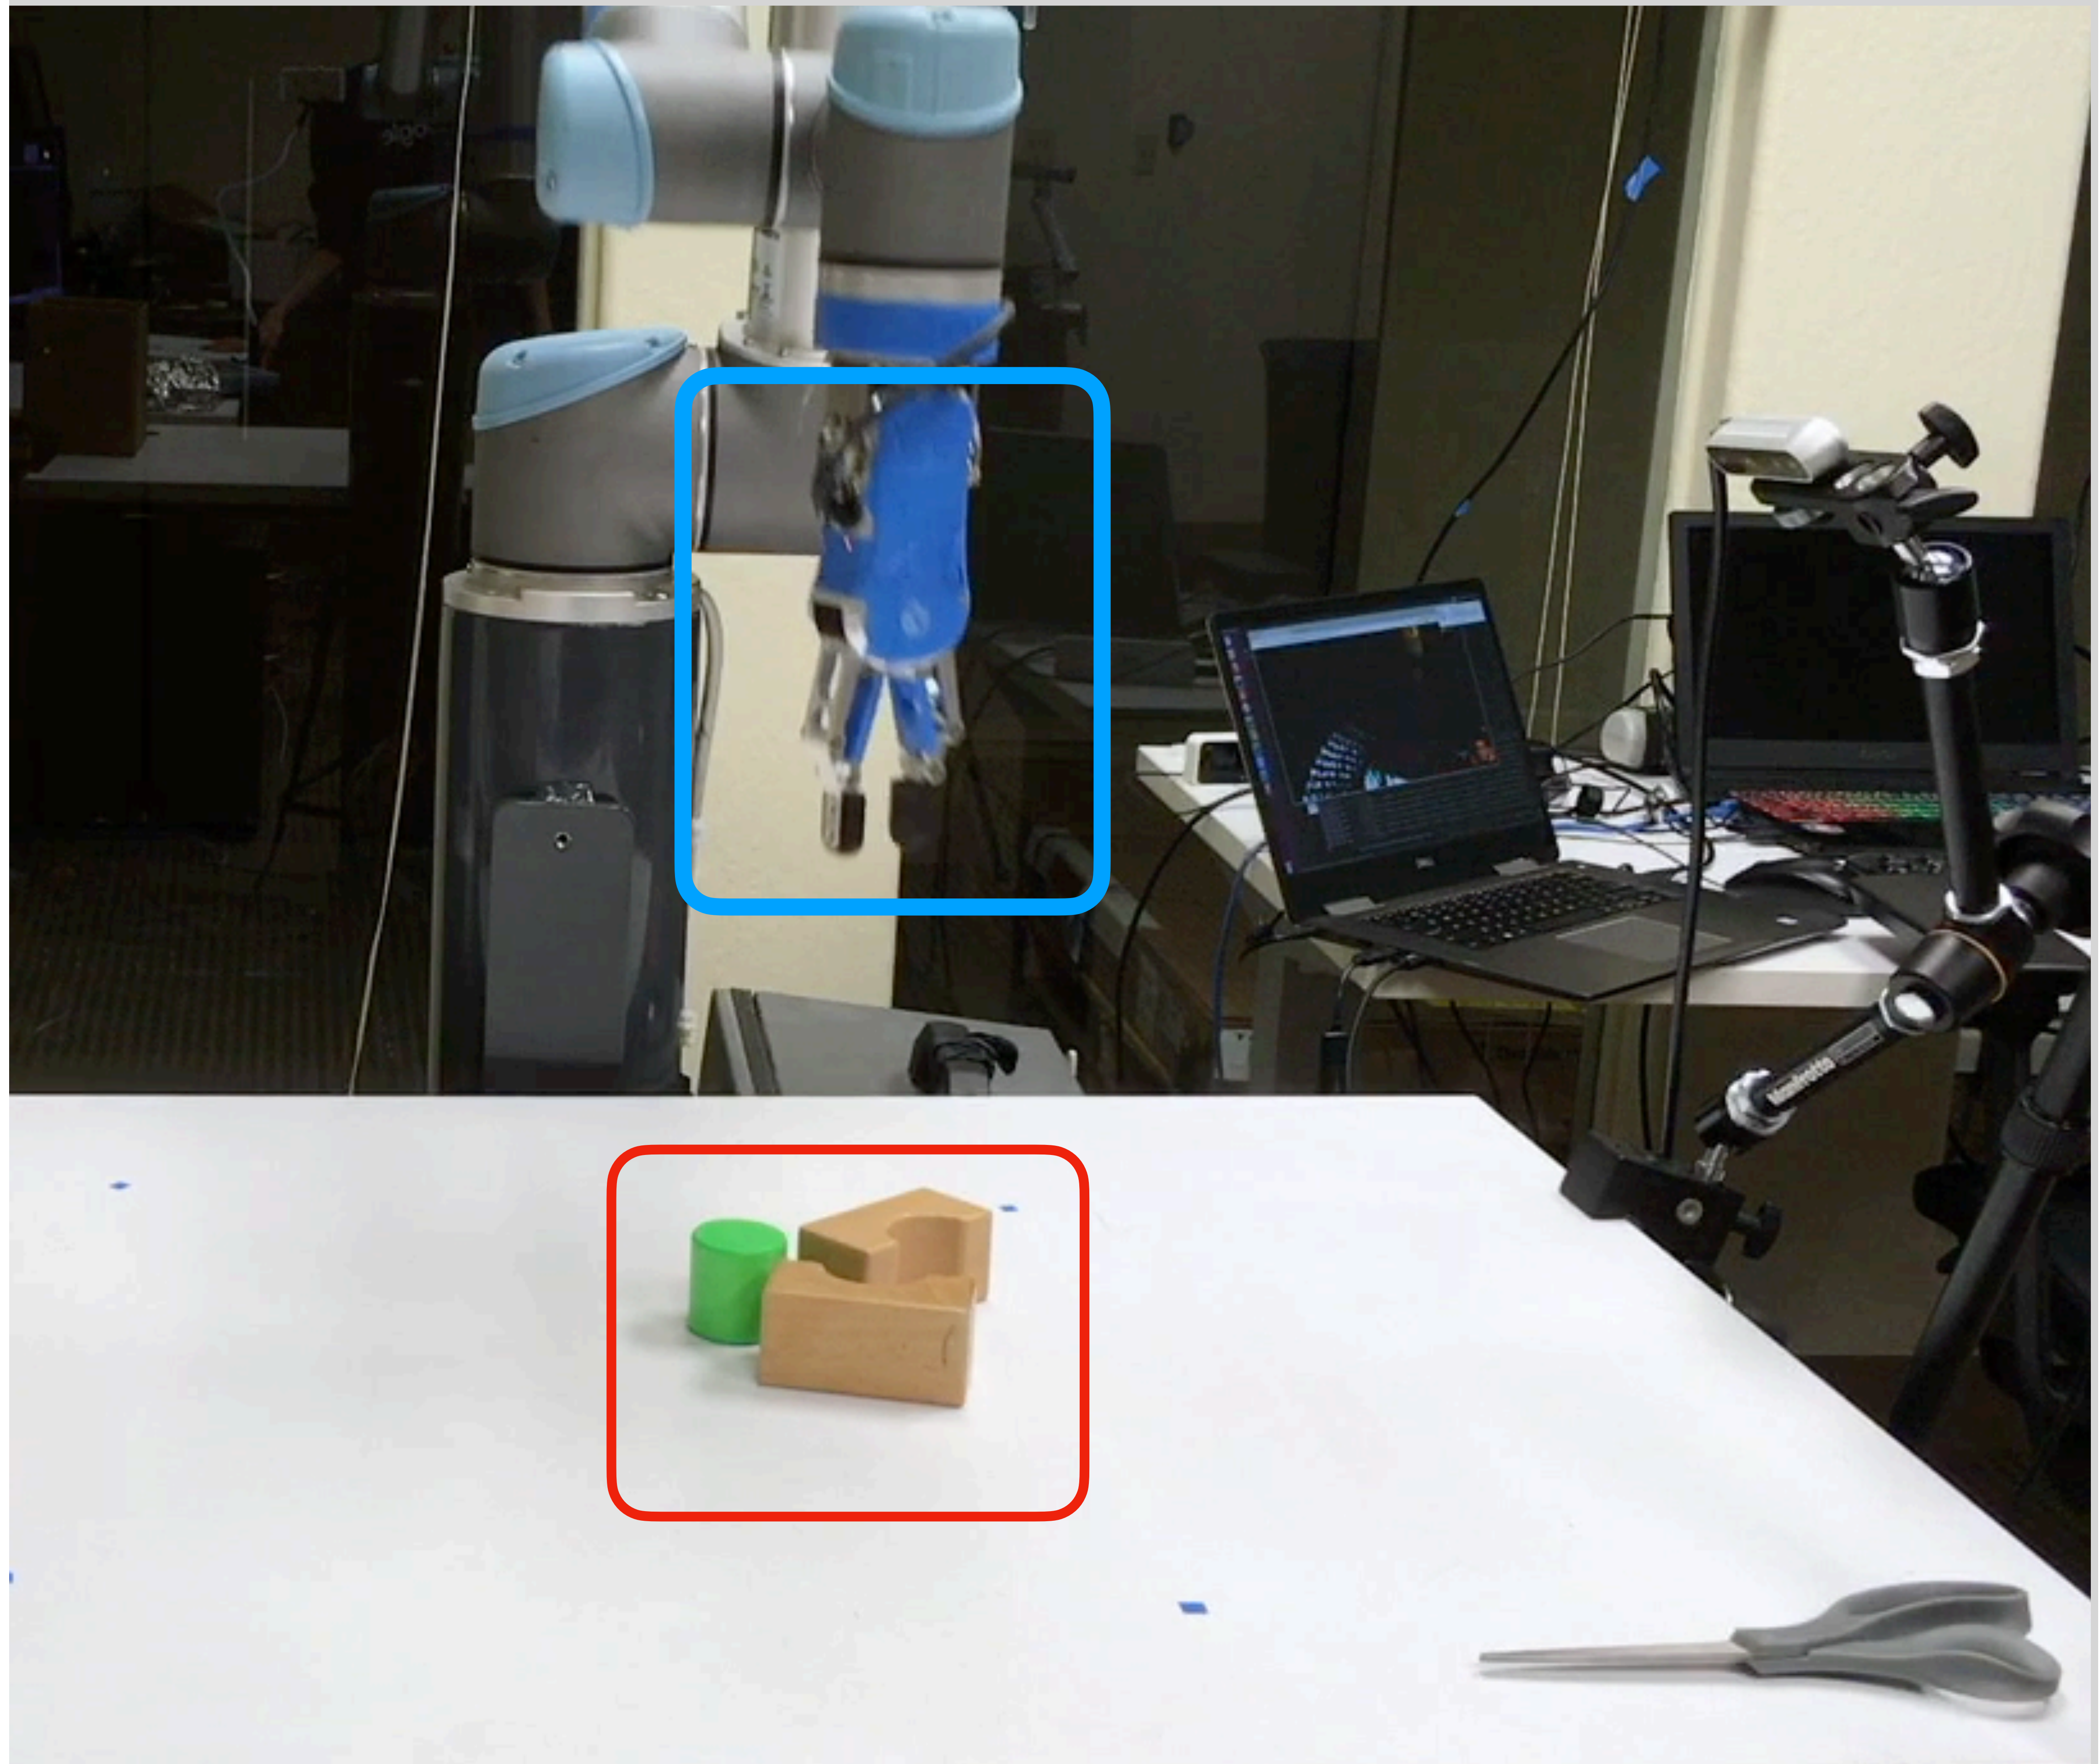

$f(\mathrm{x}) \rightarrow \text{action}$

X: External environment, objects

Where + How to grasp?

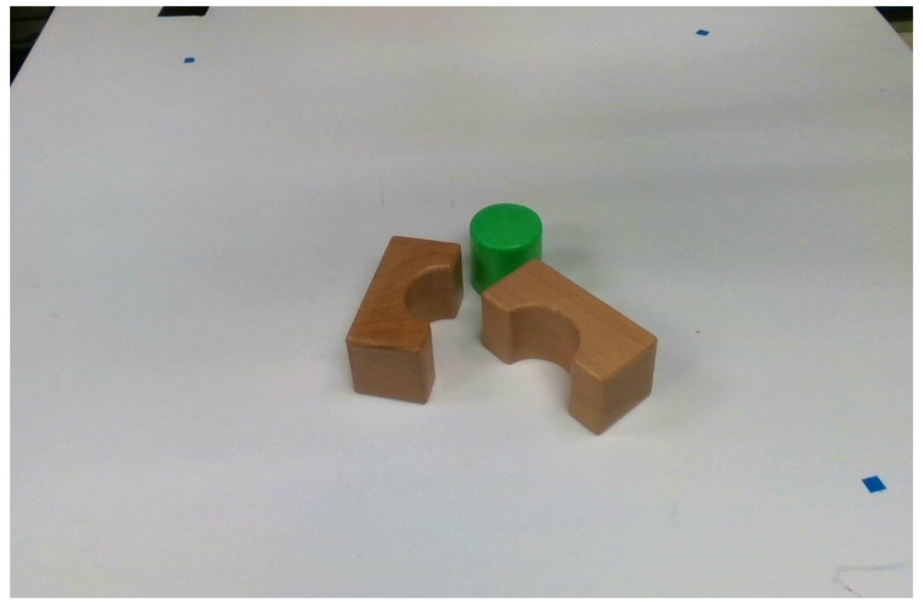

Grasping Algorithms Today ...

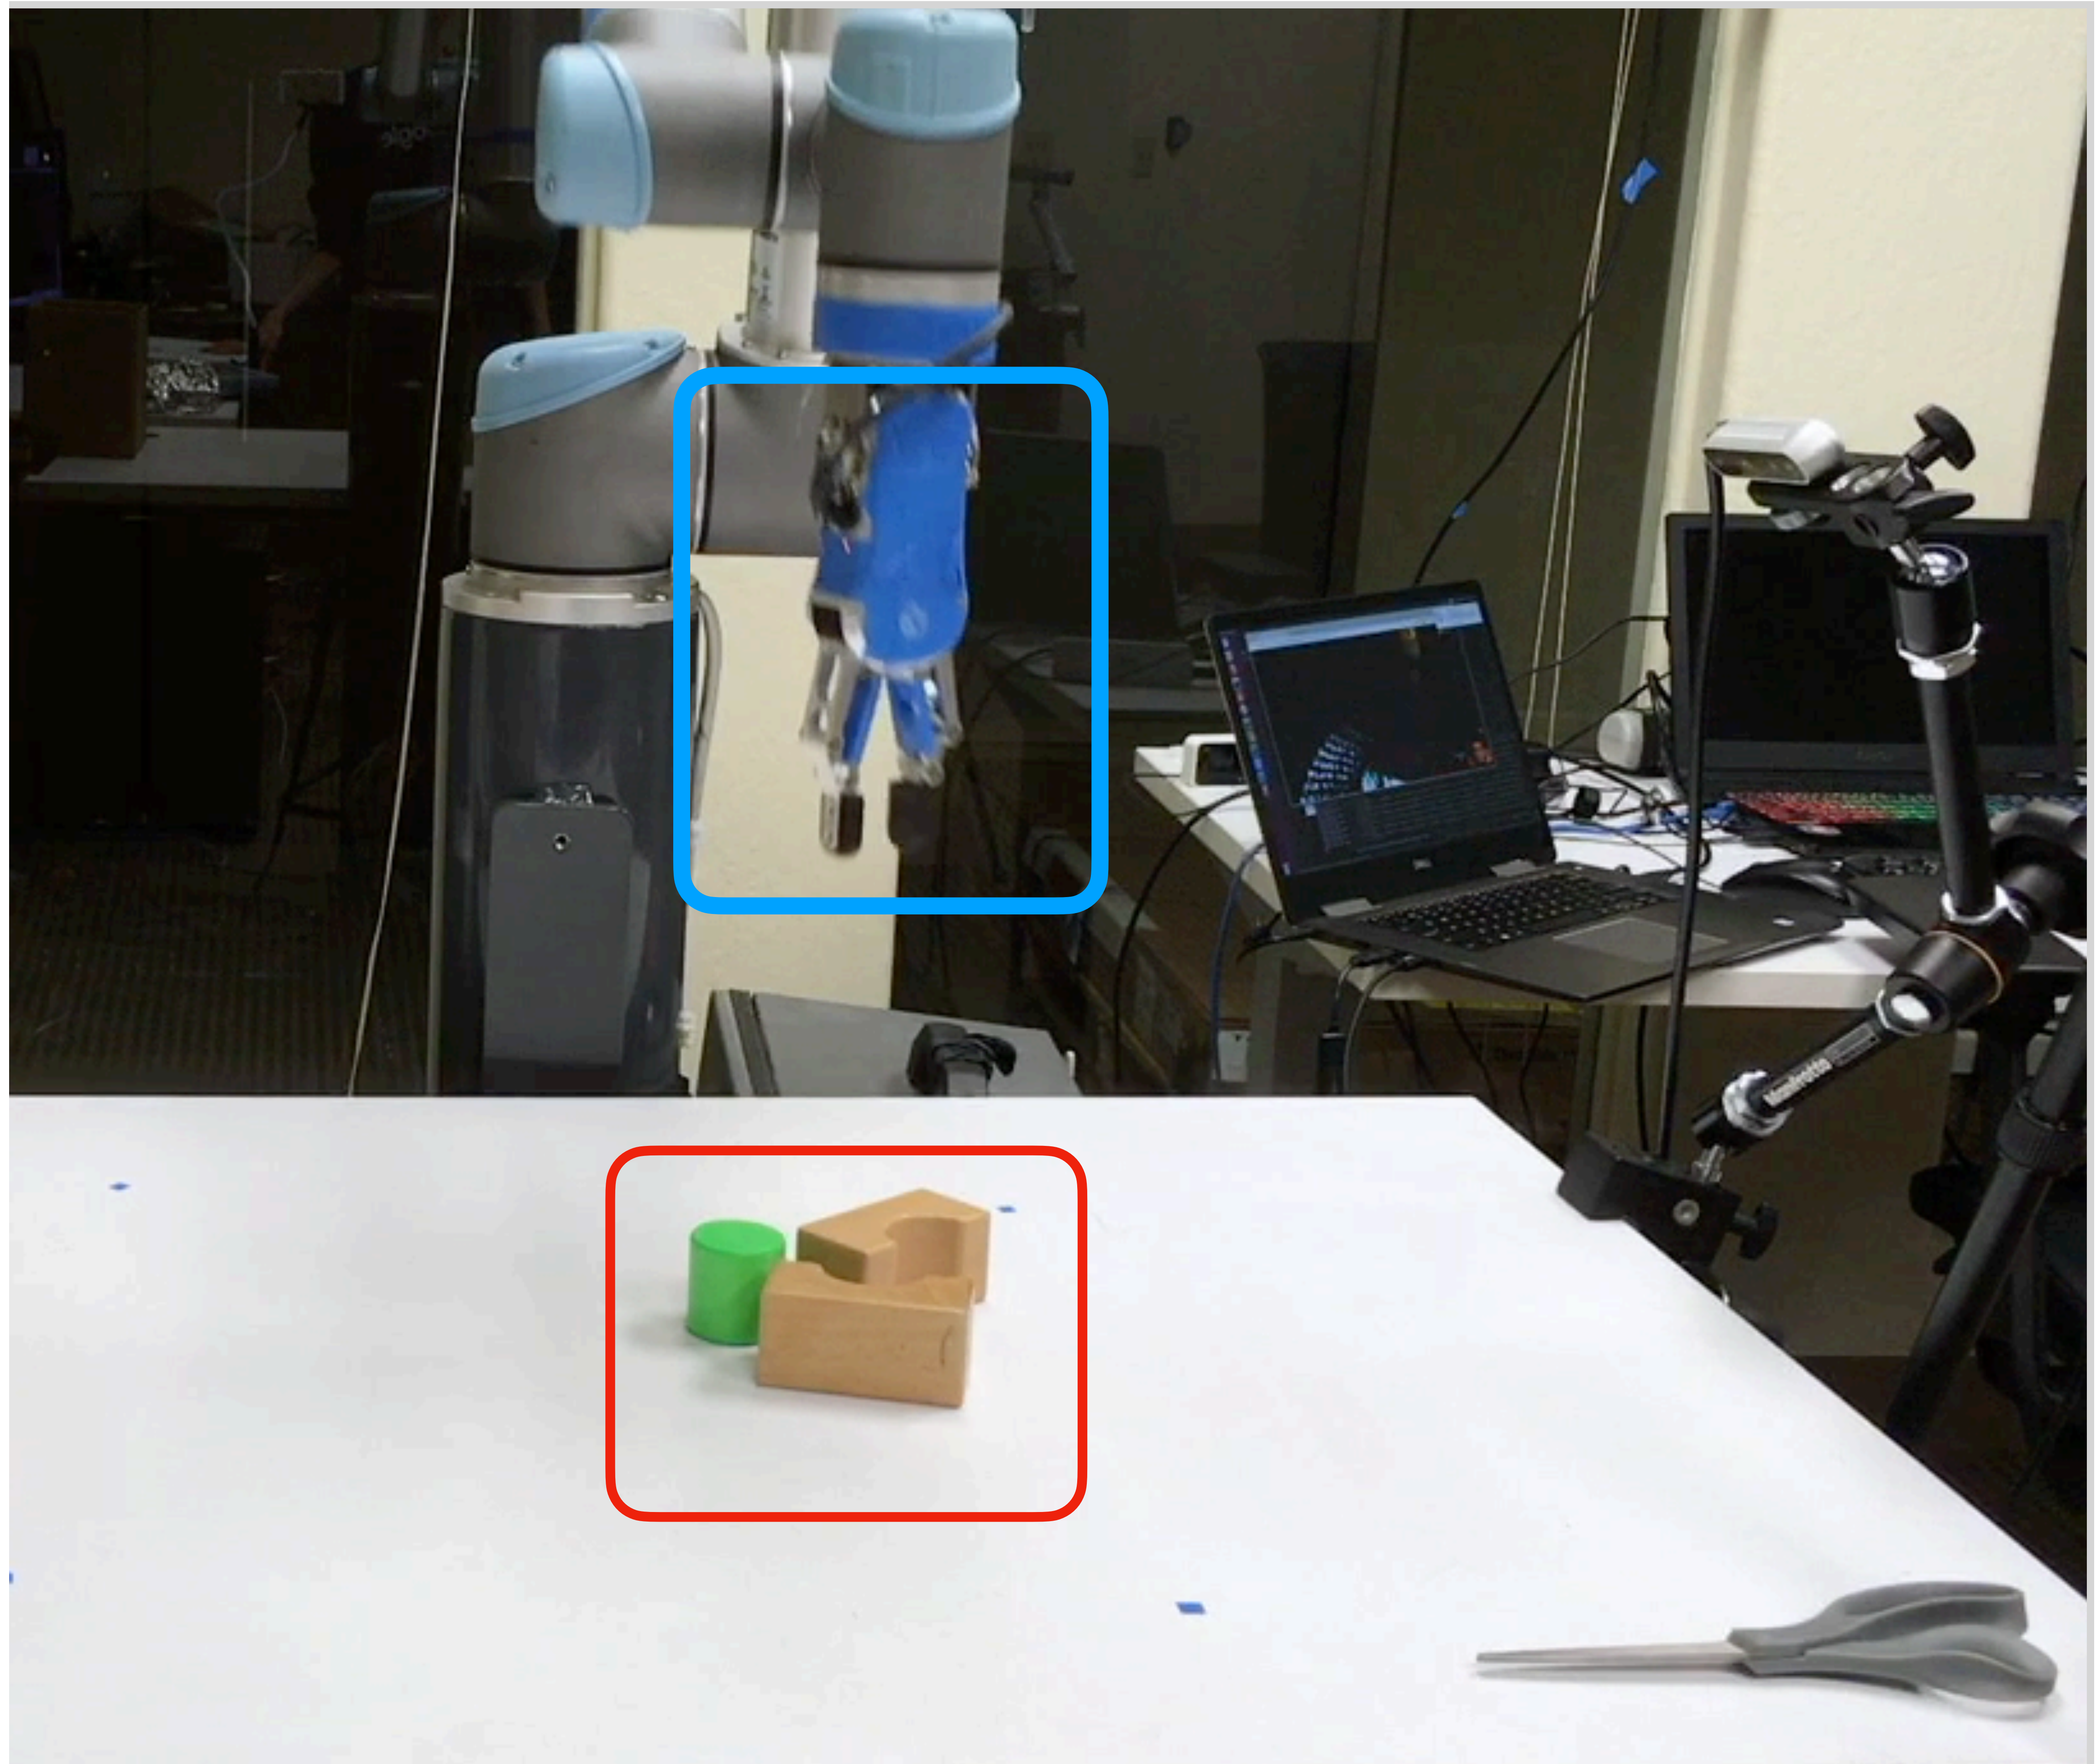

$f(\mathrm{x}, \mathrm{g}) \rightarrow \text{action}$

g: Robot gripper hardware

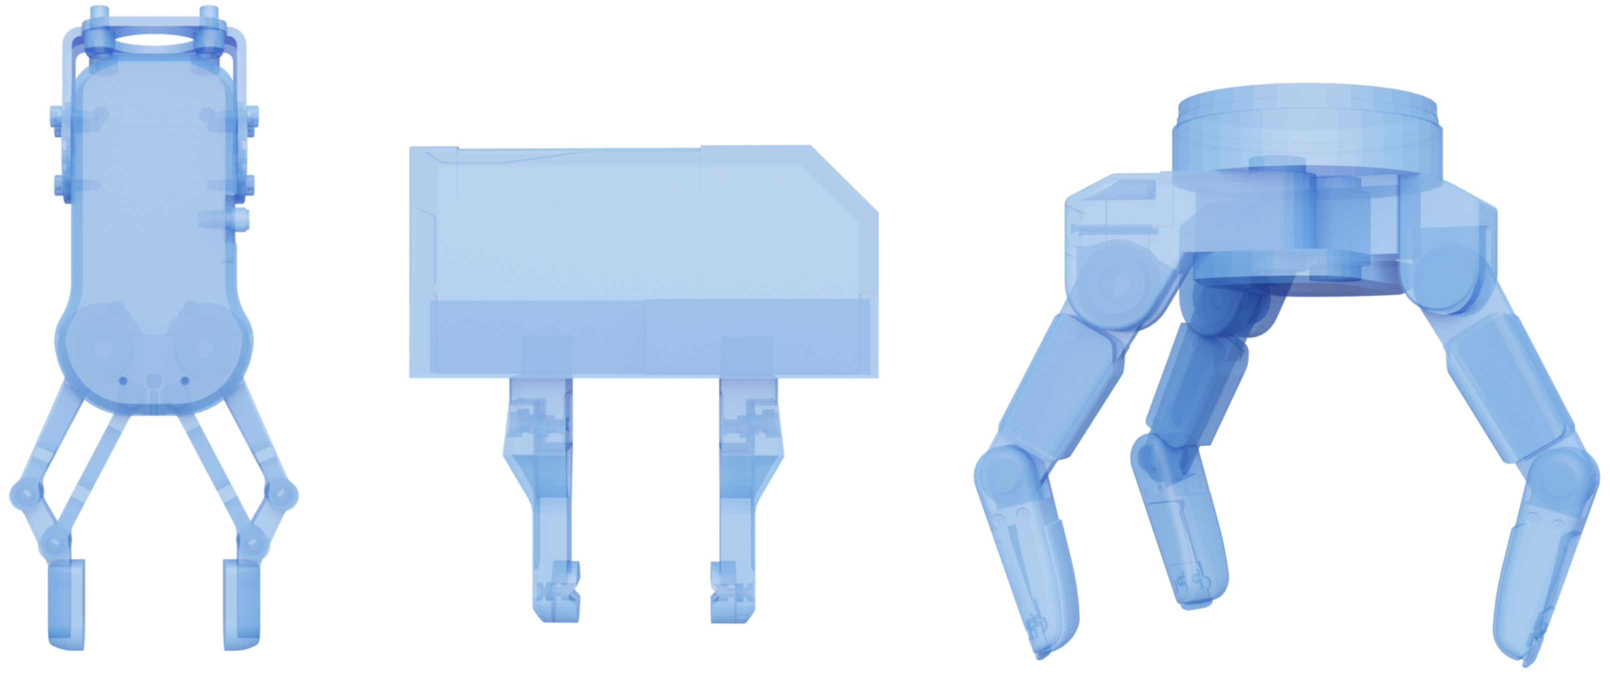

Motivation

Scene

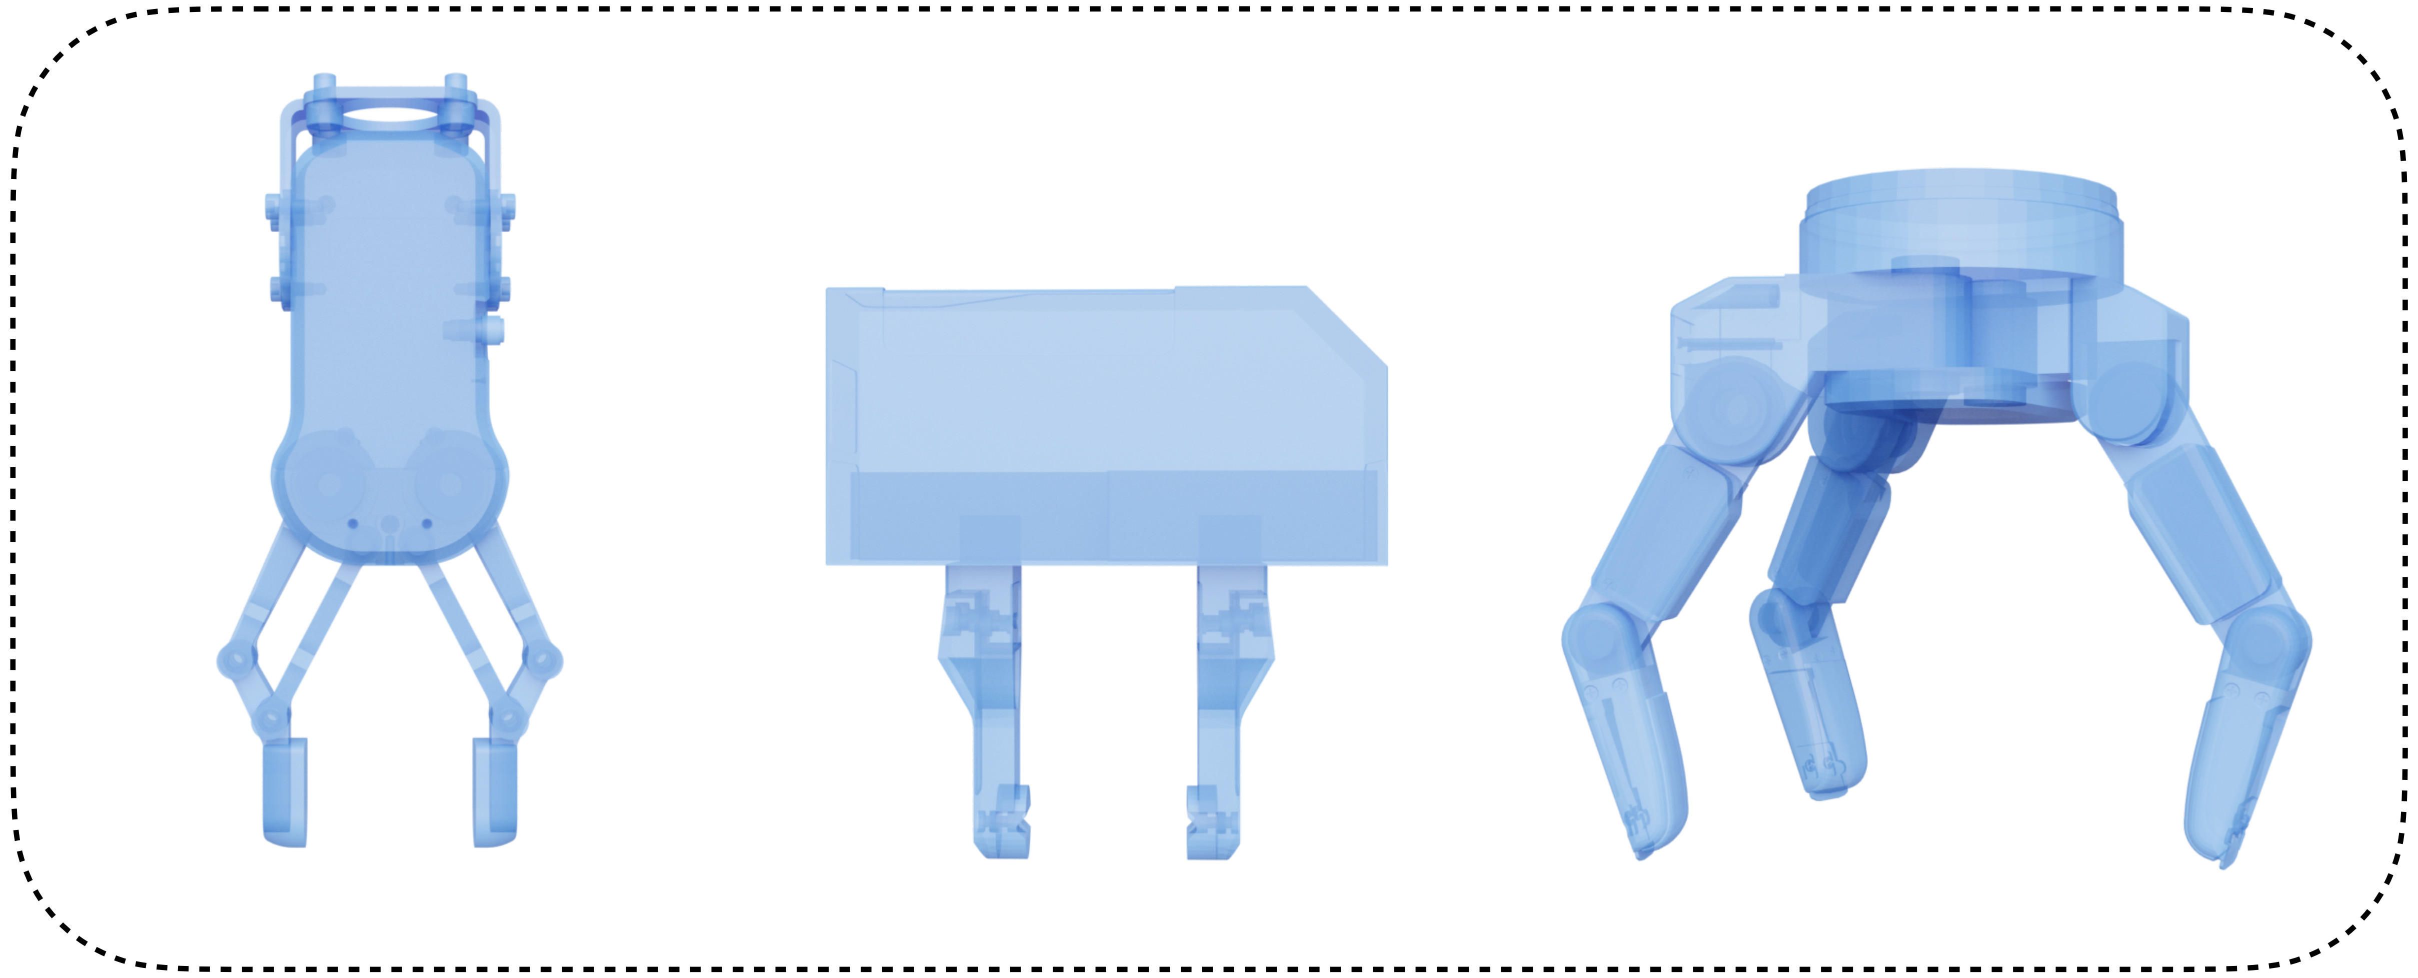

Grippers

Motivation

RG2

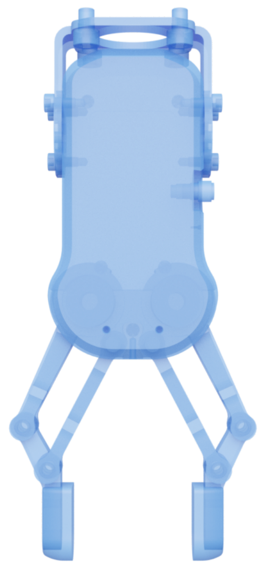

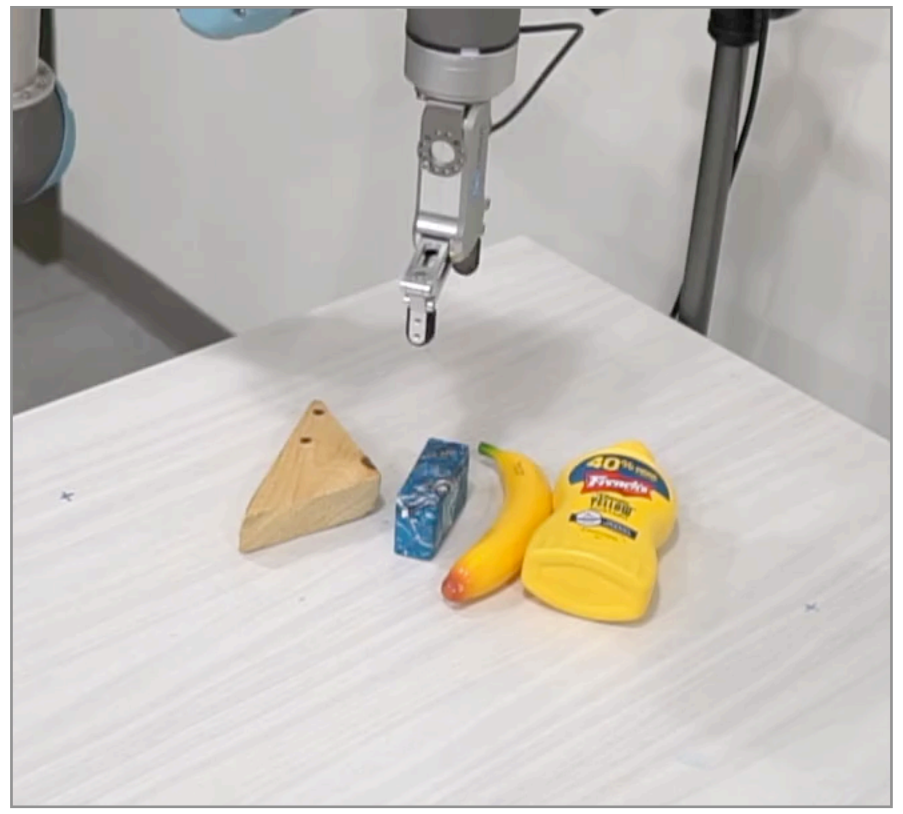

WSG 50

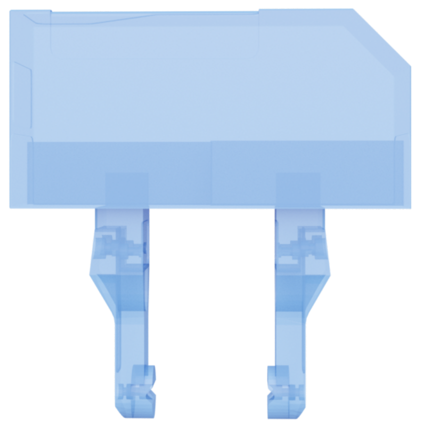

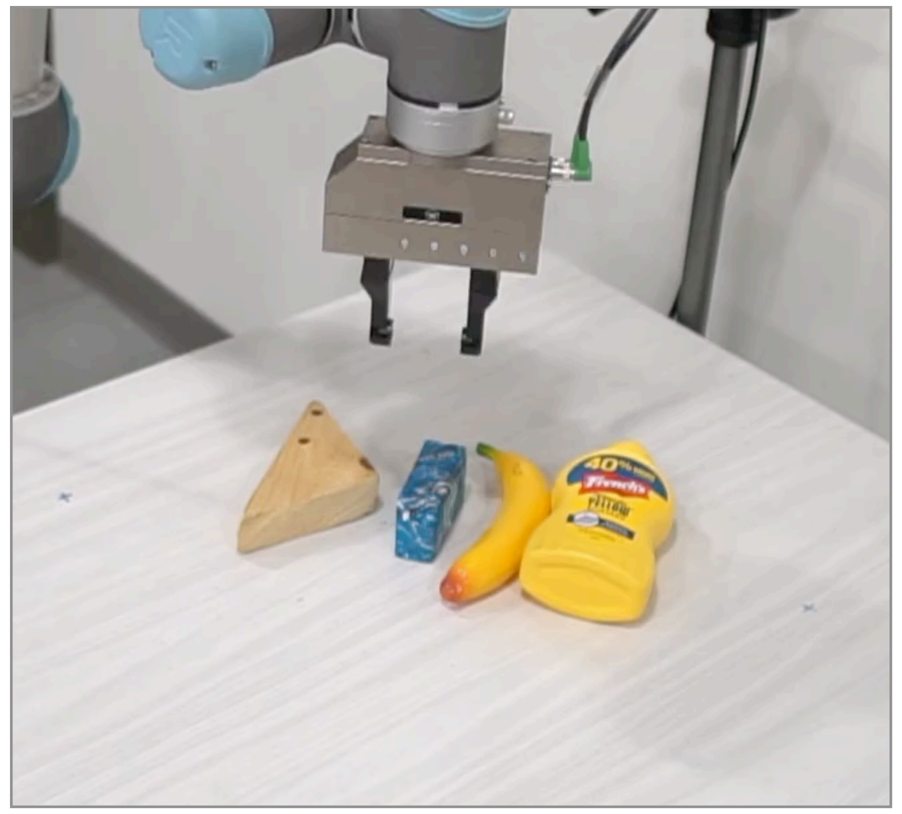

Barrent Hand

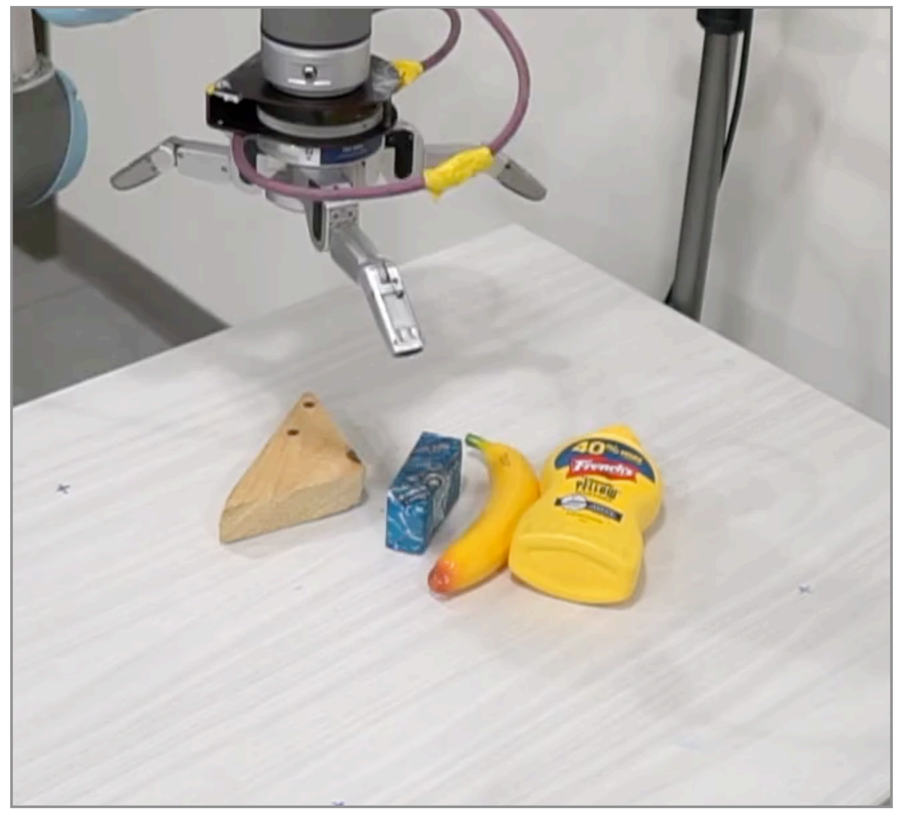

Motivation

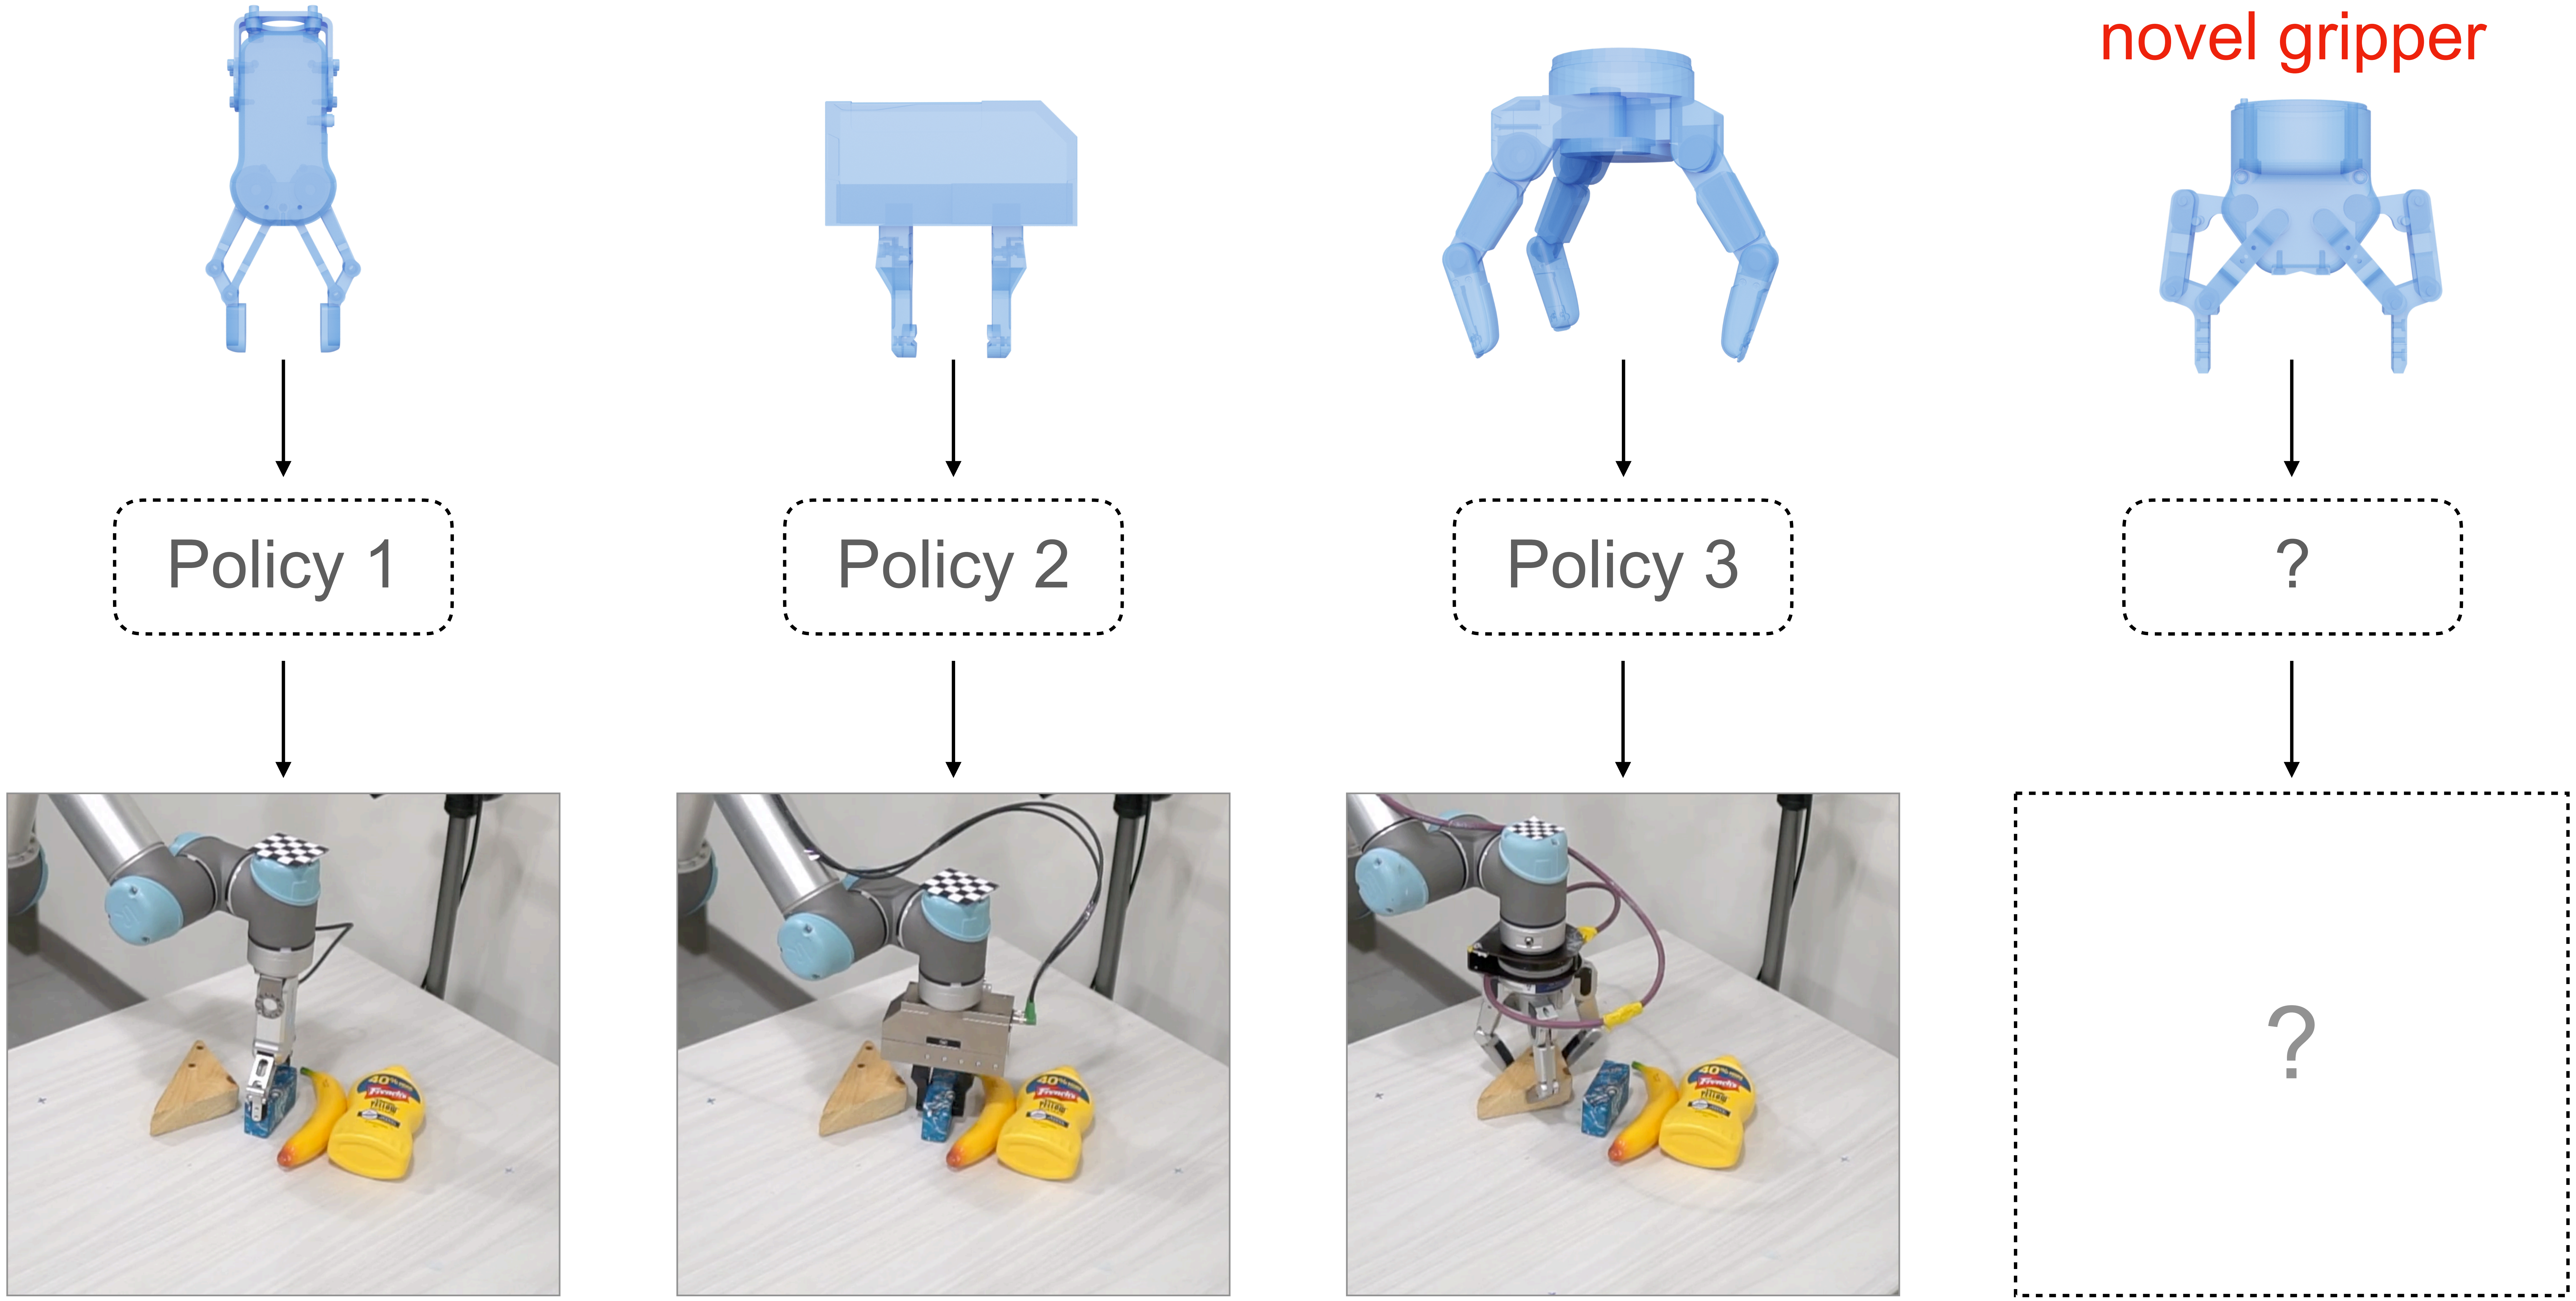

Motivation

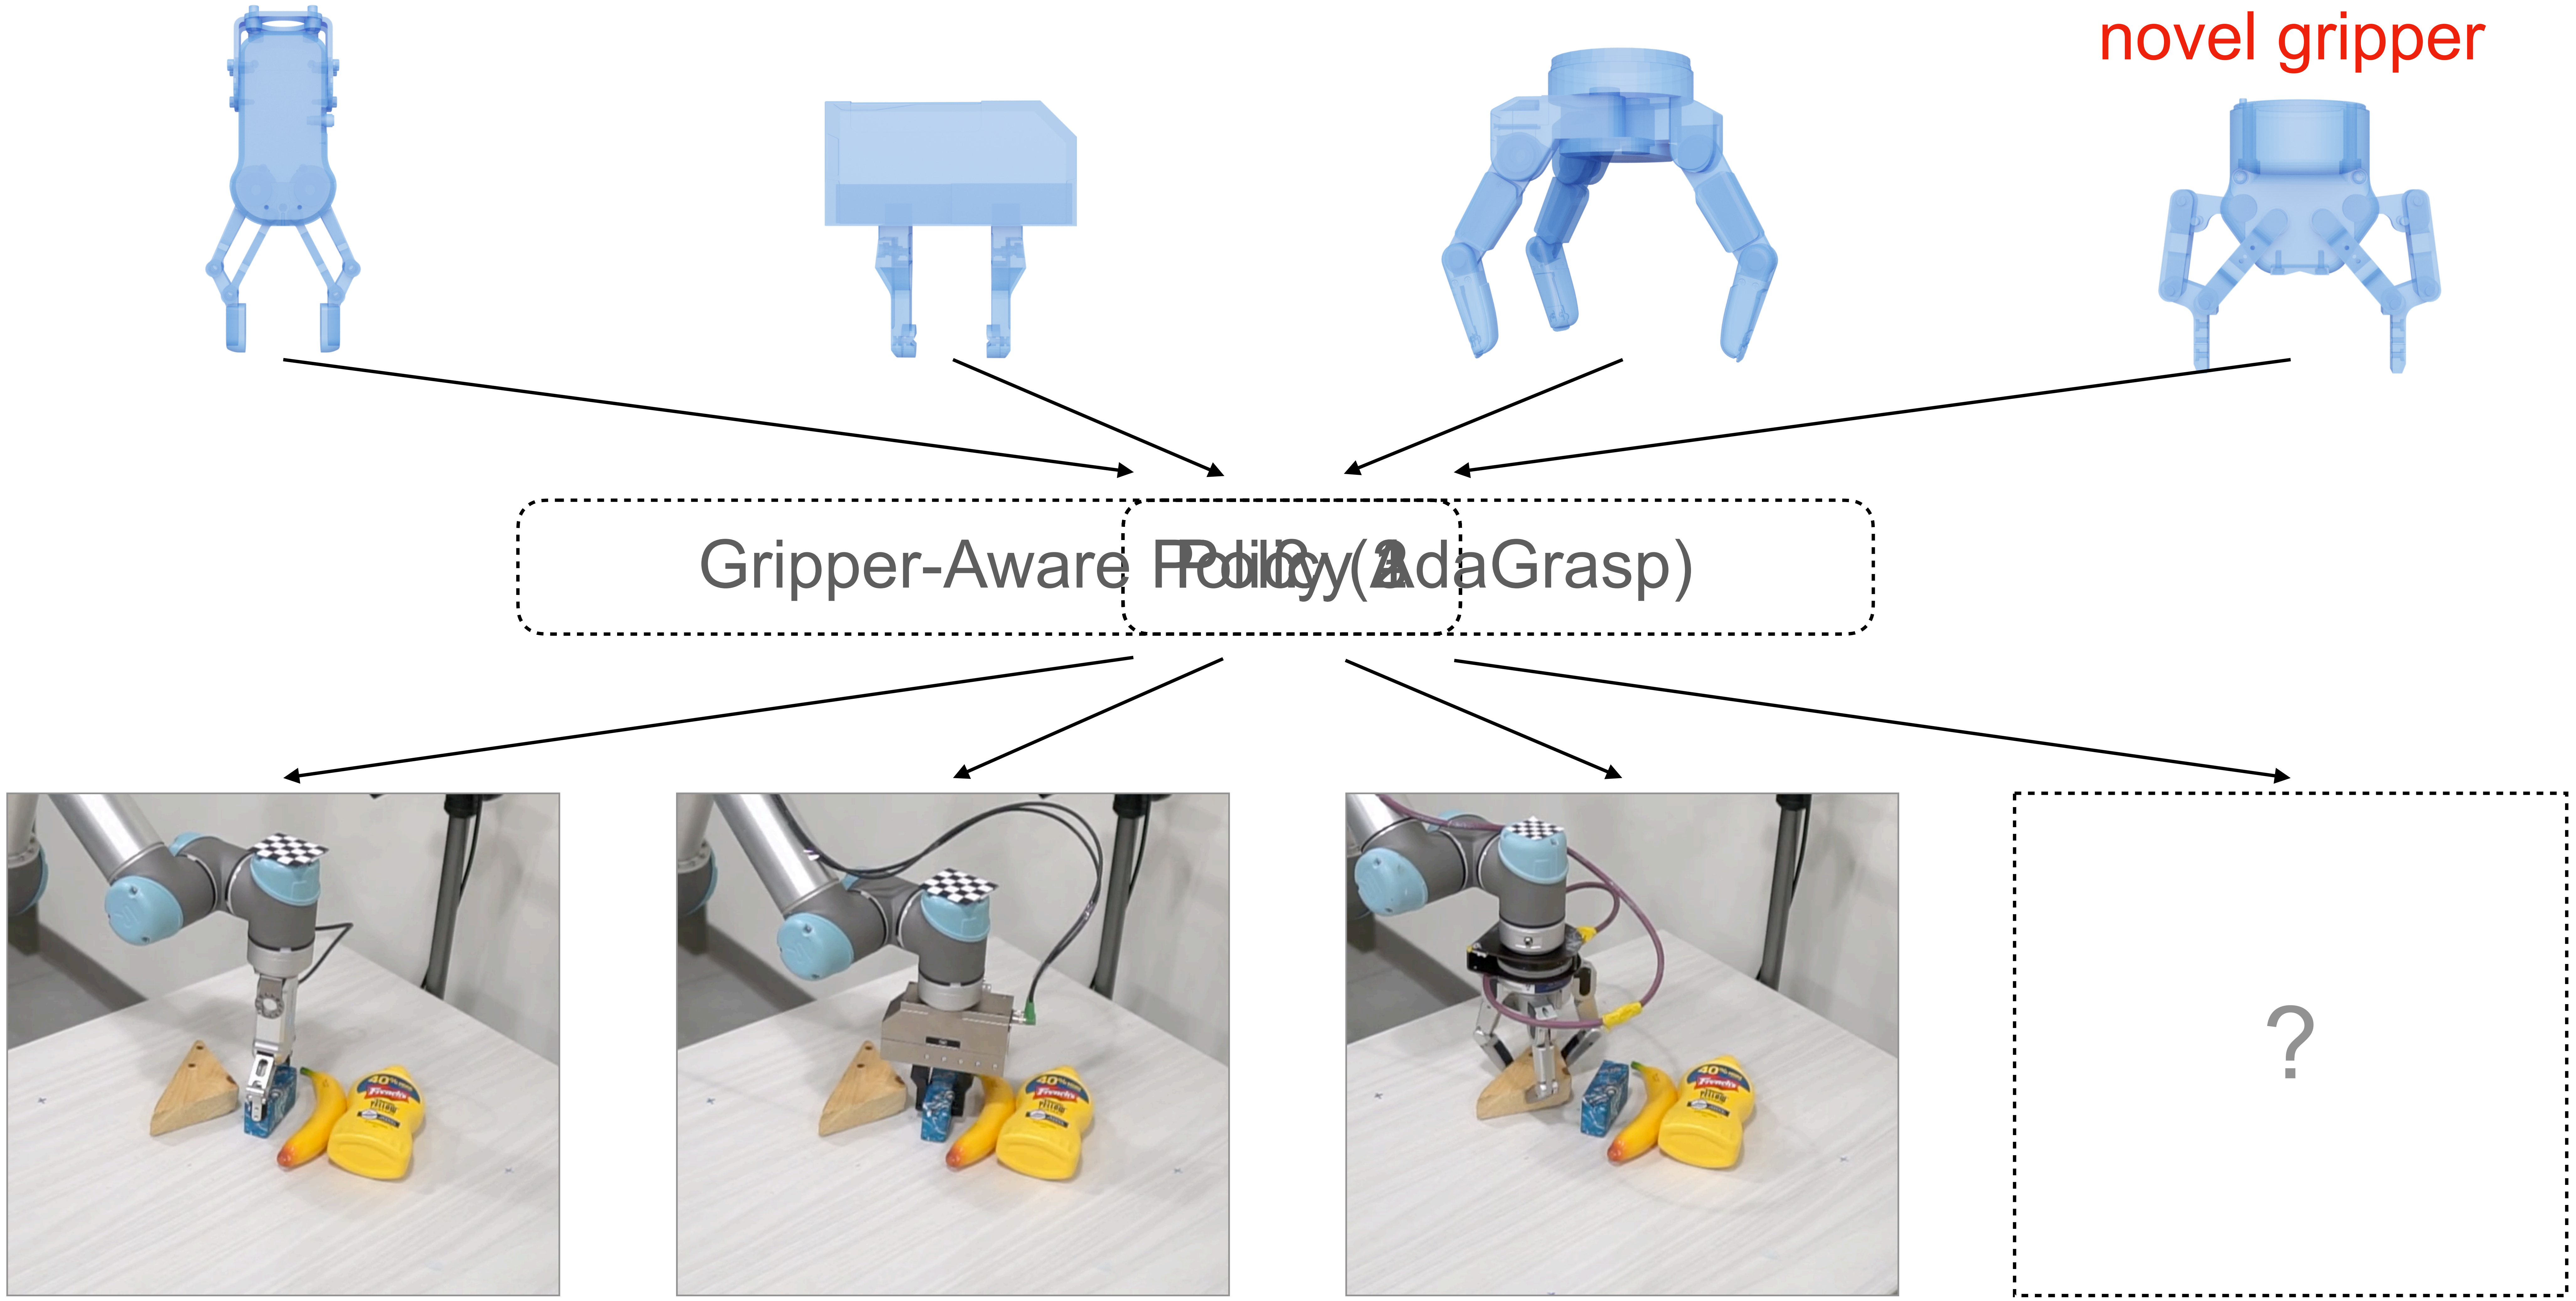

Advantages

- **Adaptability:** quickly adapt to <u>*new*</u> grippers without additional training data.

New Gripper can also be
broken or damaged gripper

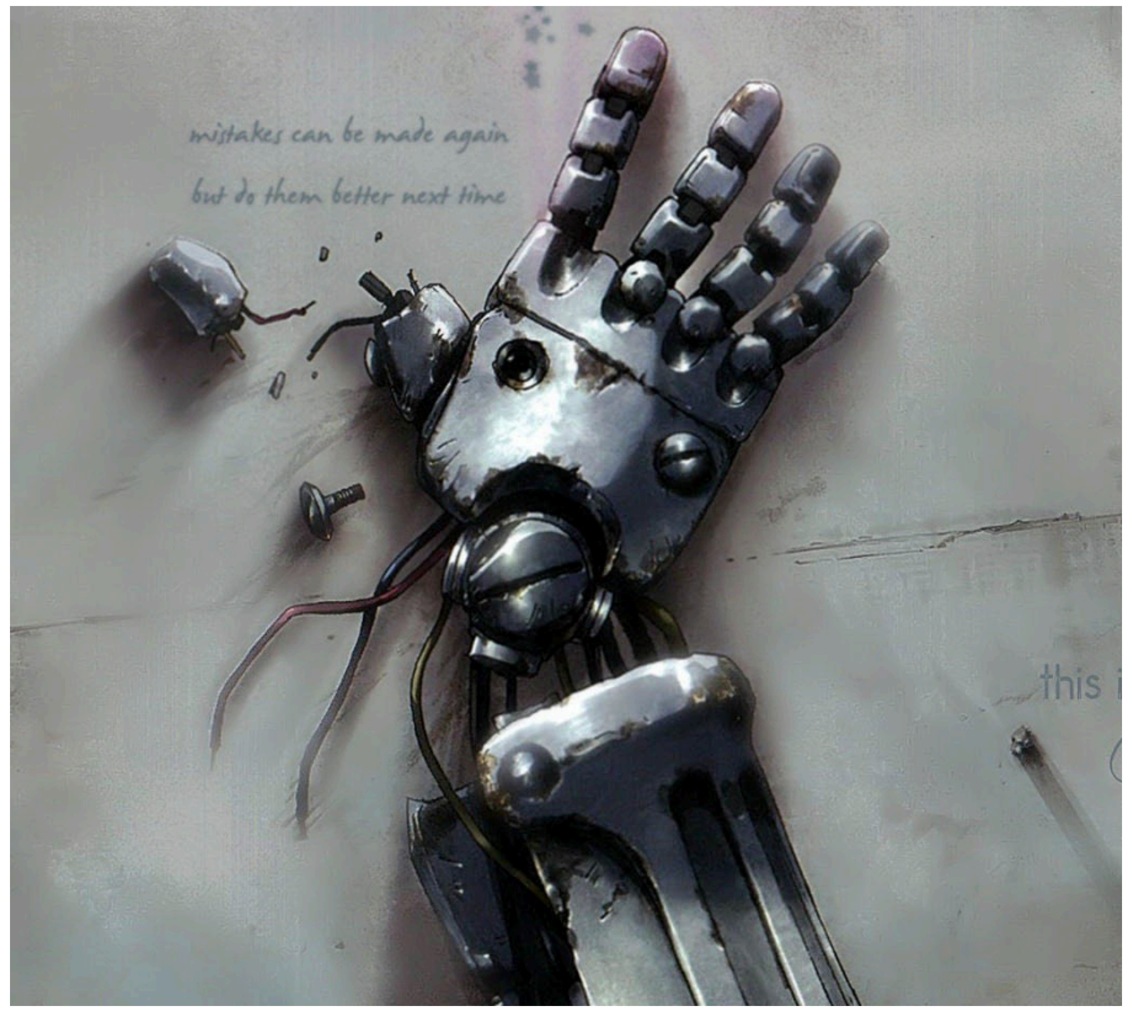

Advantages

- **Adaptability:** quickly adapt to *new* grippers without additional training data.

- **Versatility:** learn to adequately use a diverse set of grippers, the system can effectively improve its versatility on handling a larger variety of objects and tasks.

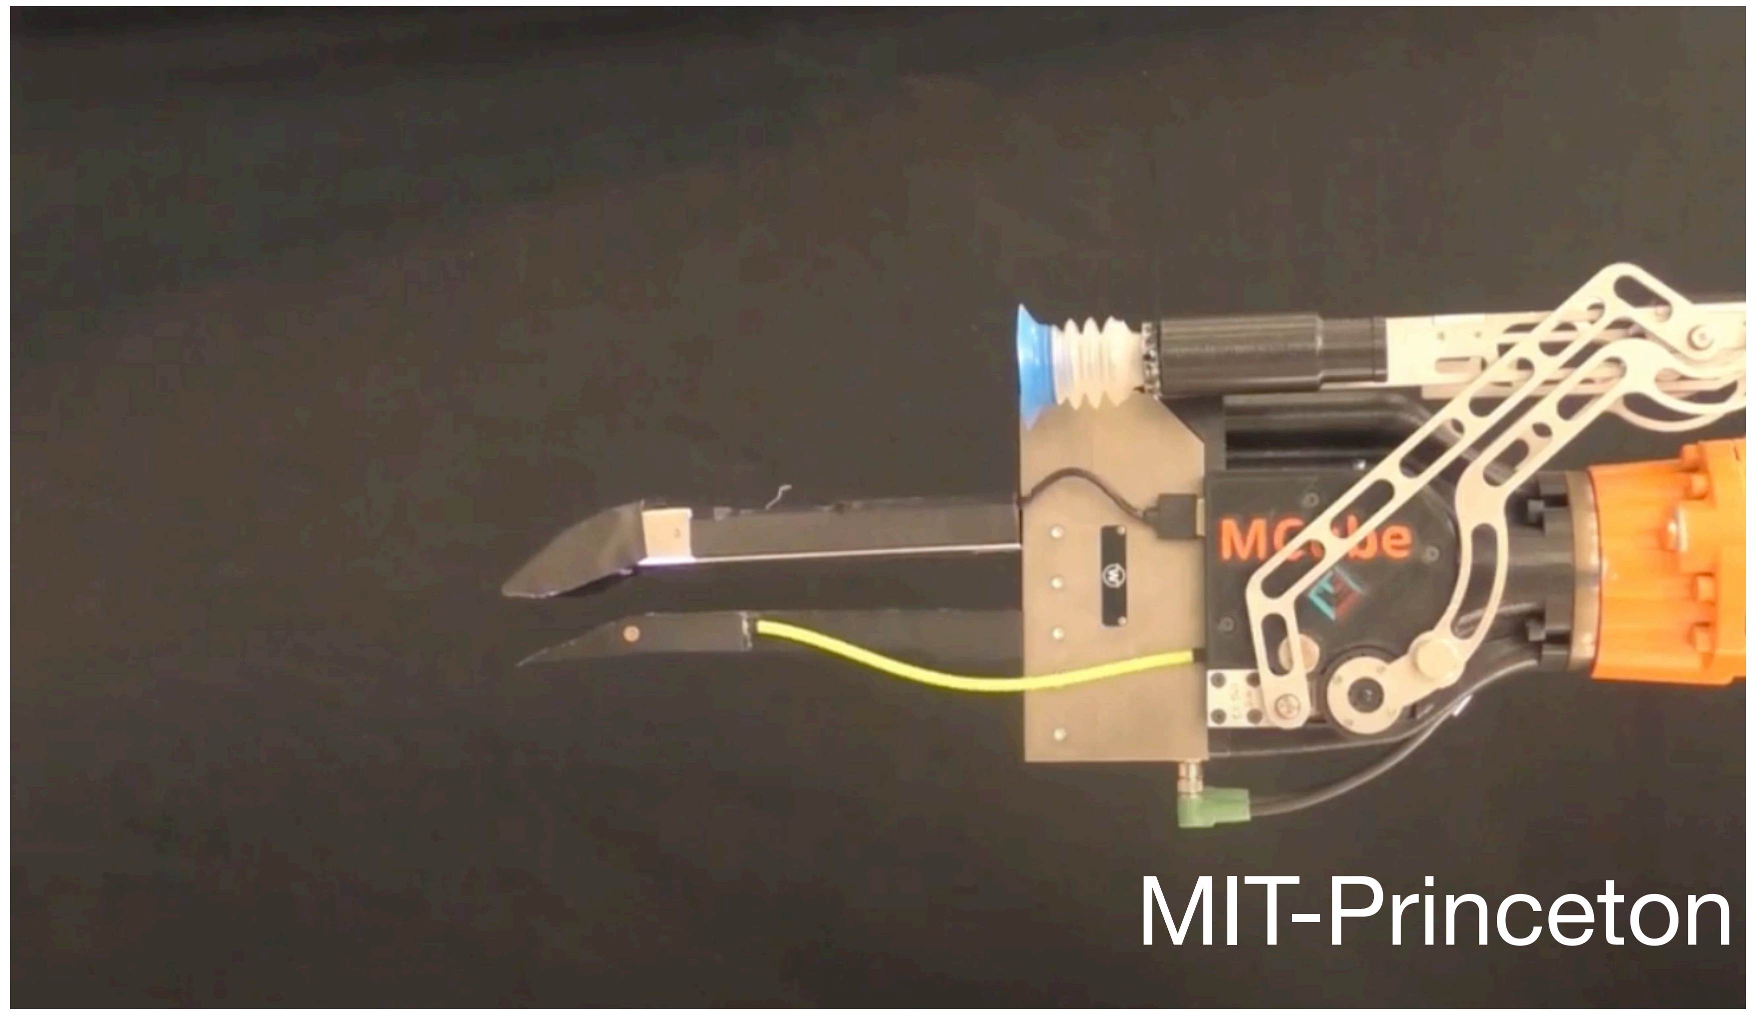

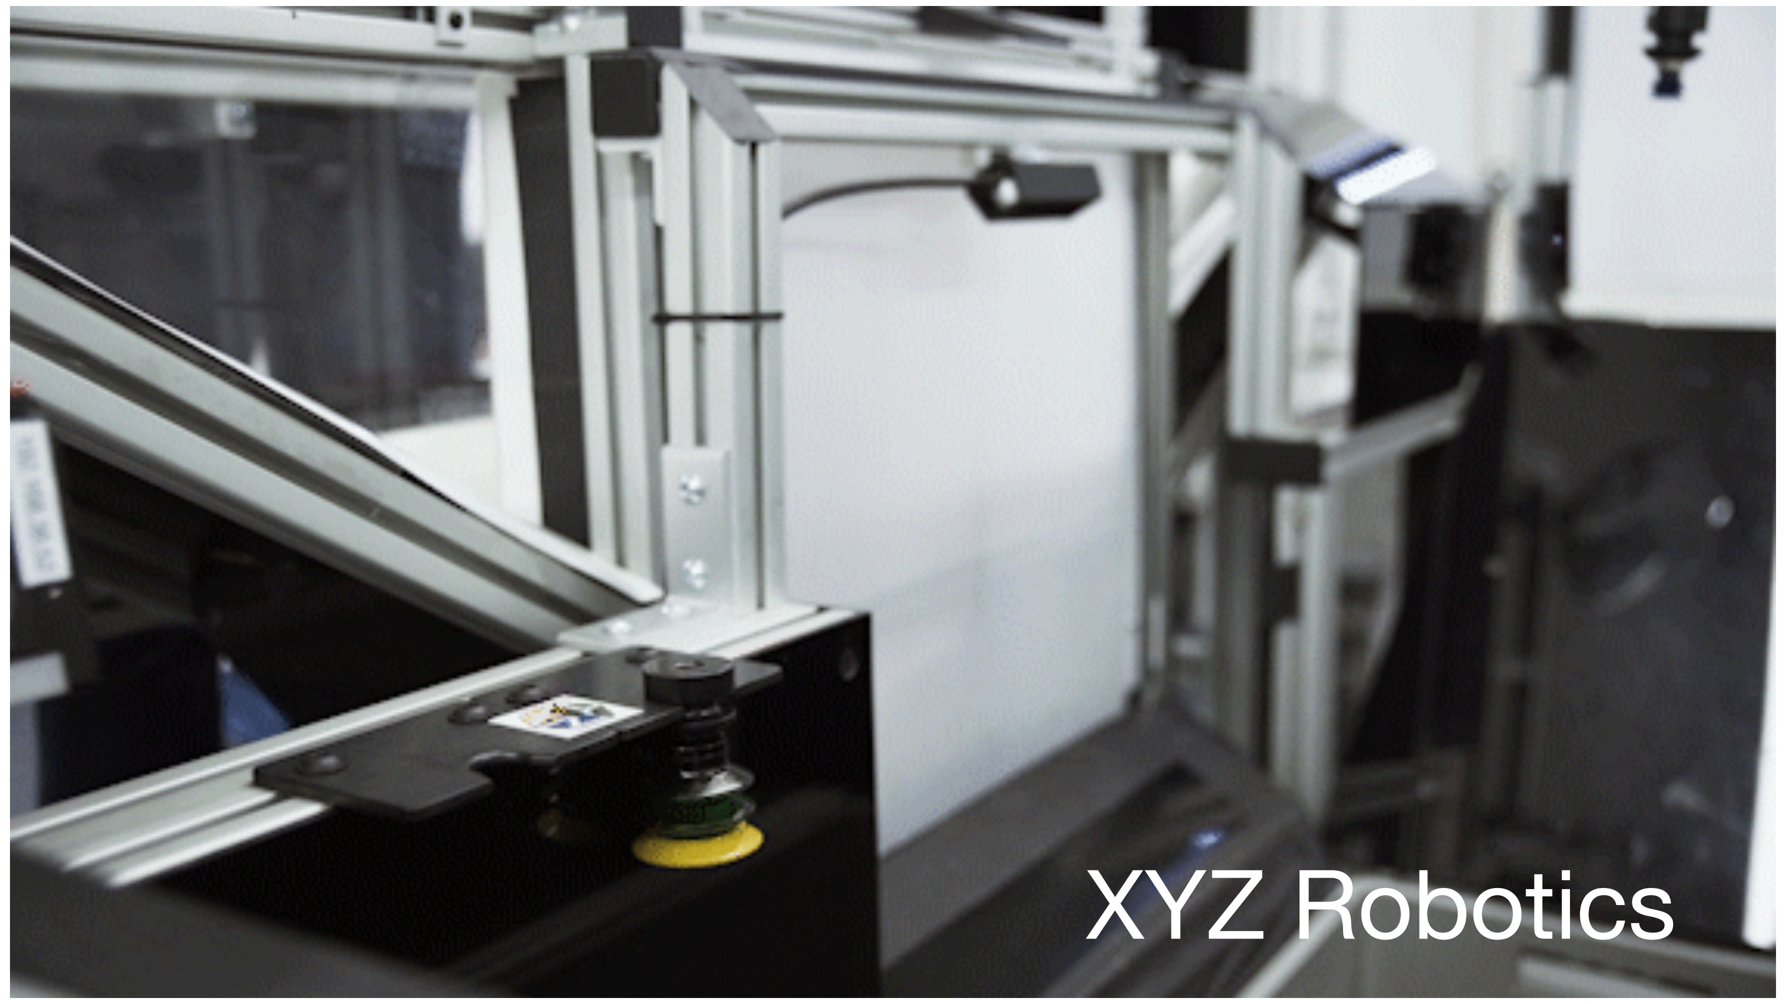

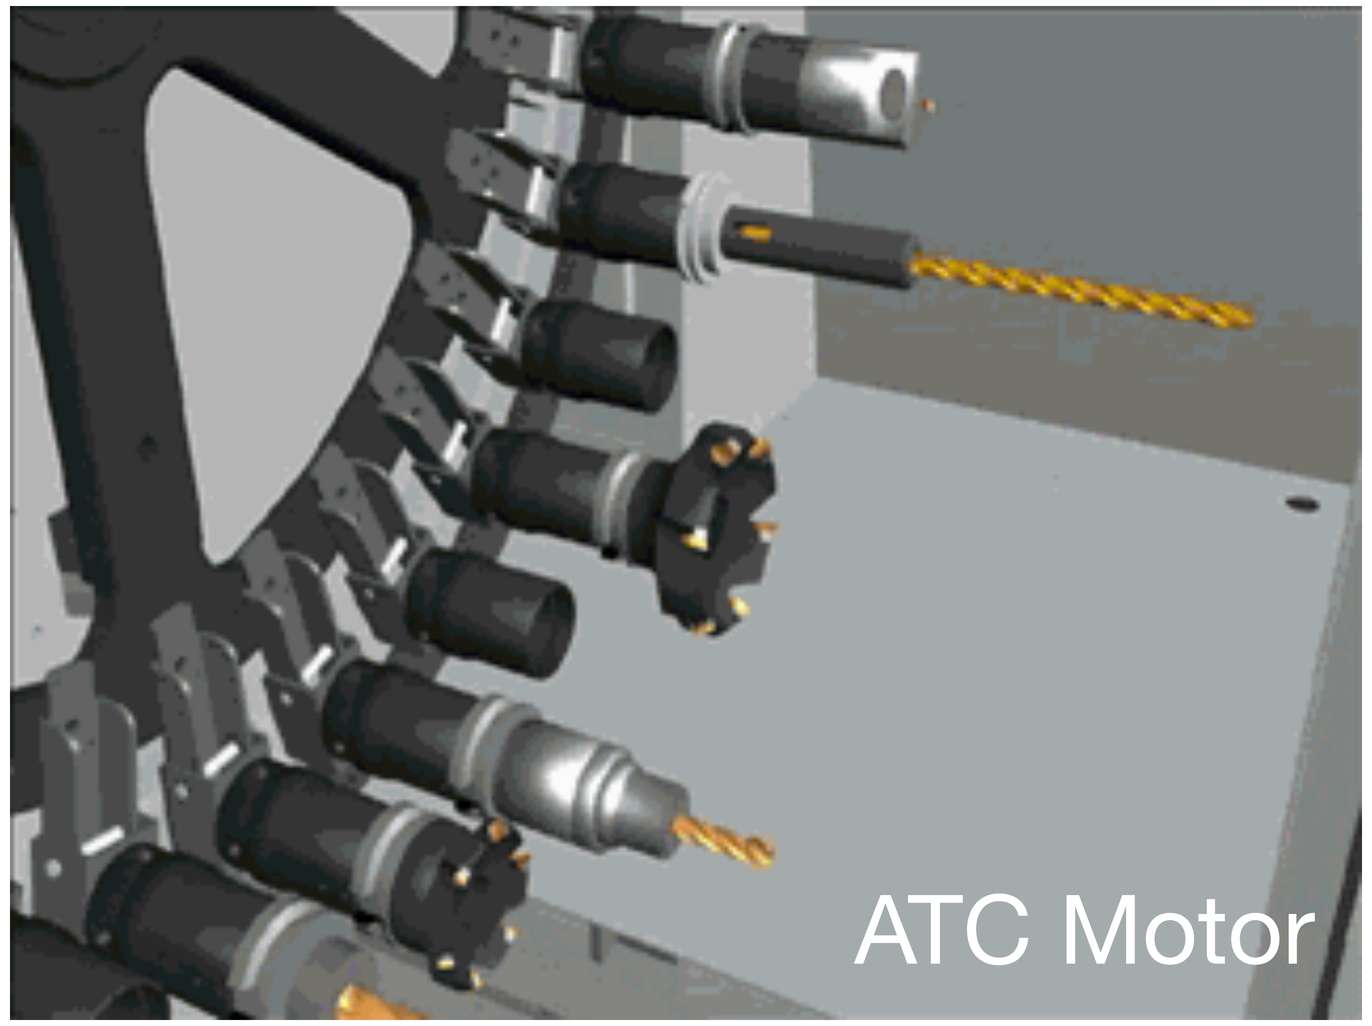

AdaGrasp Overview

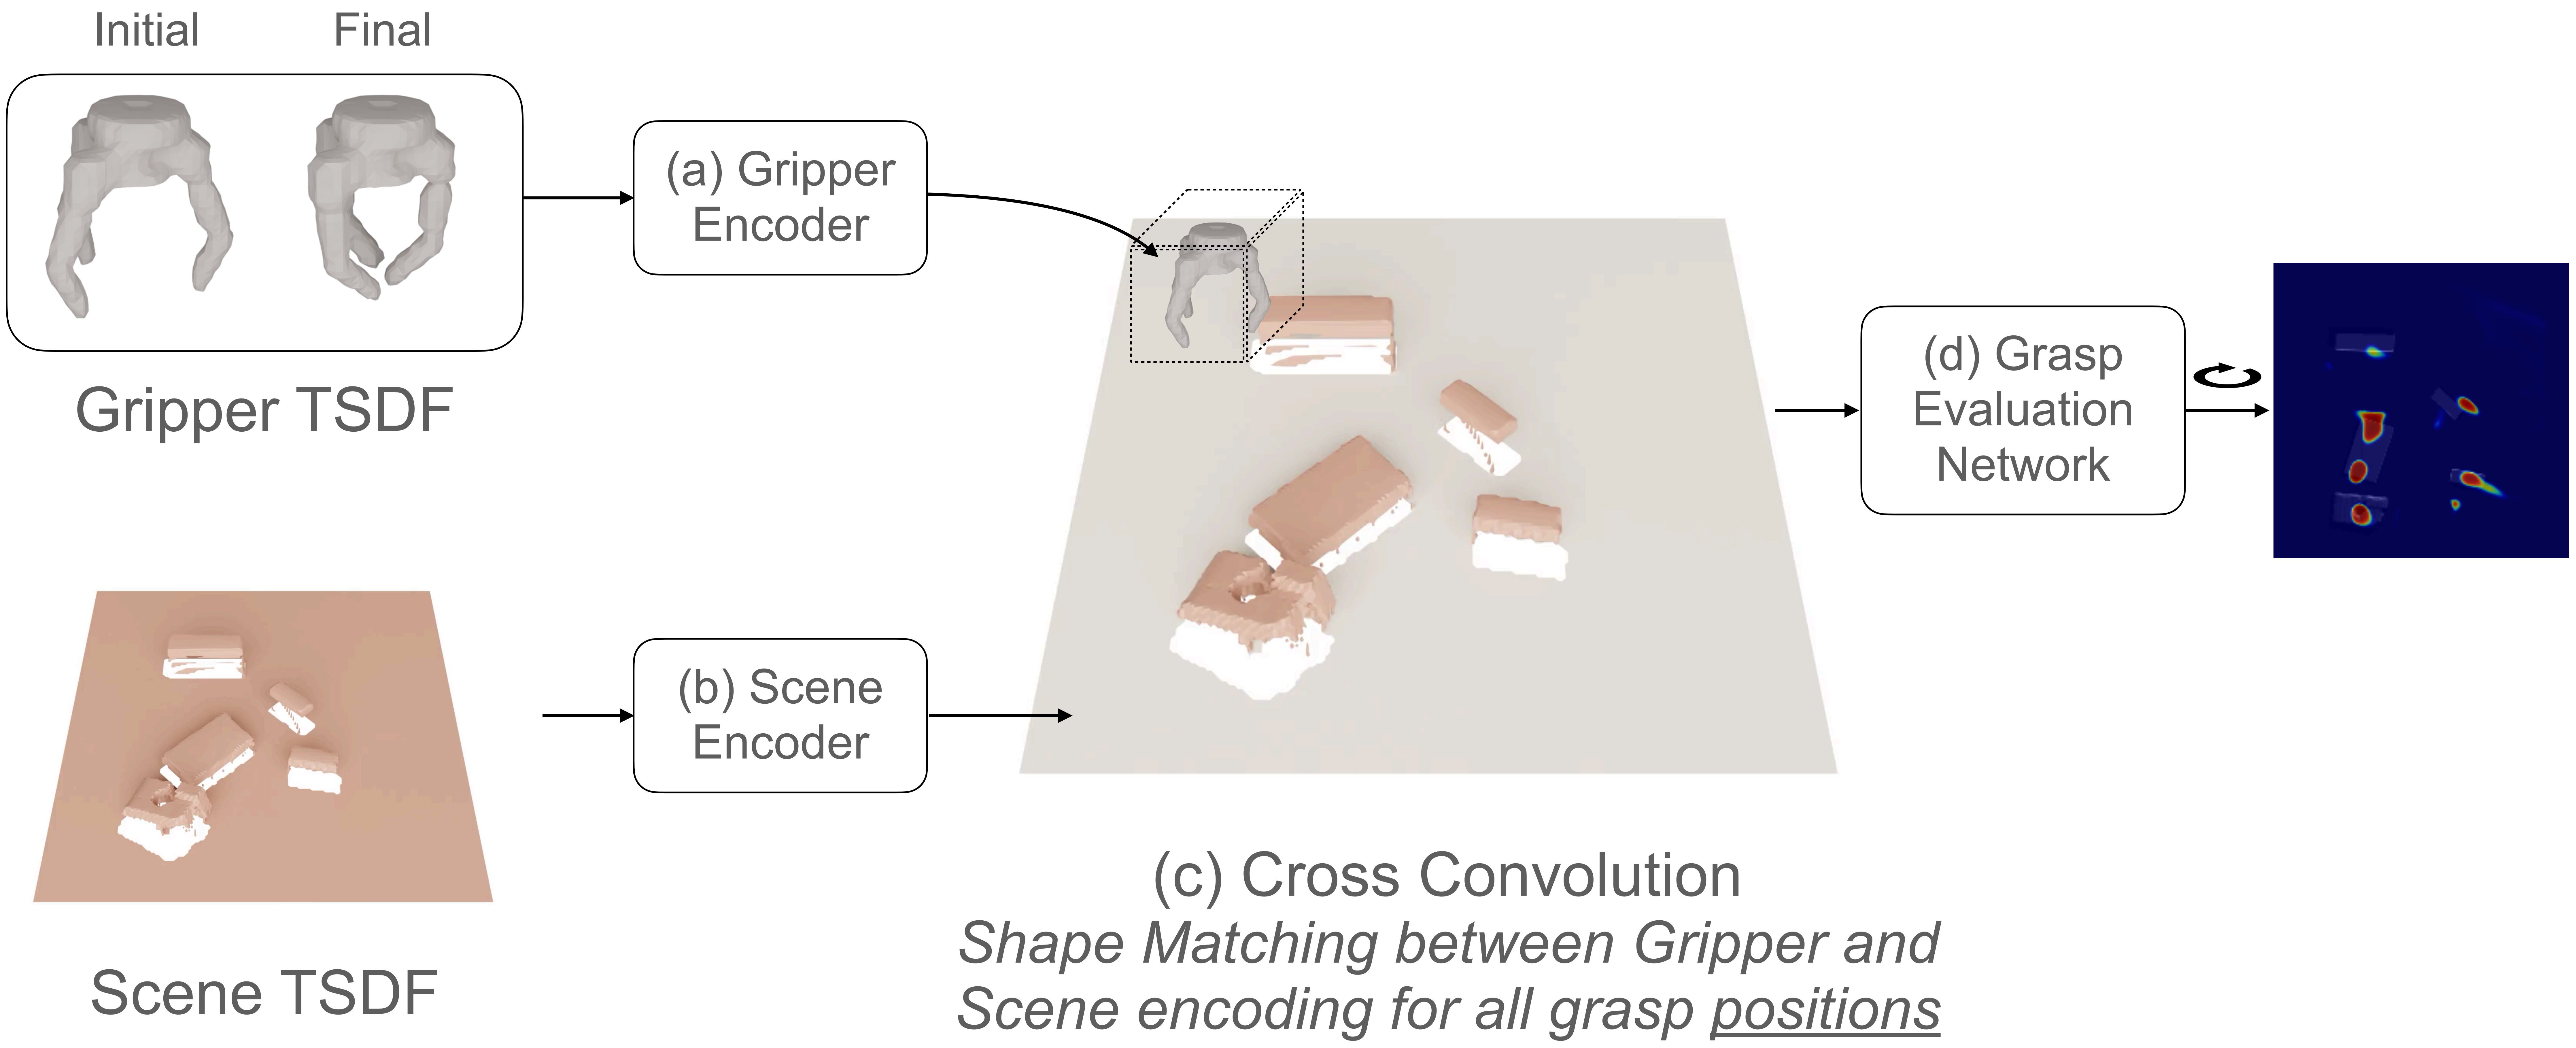

AdaGrasp Overview

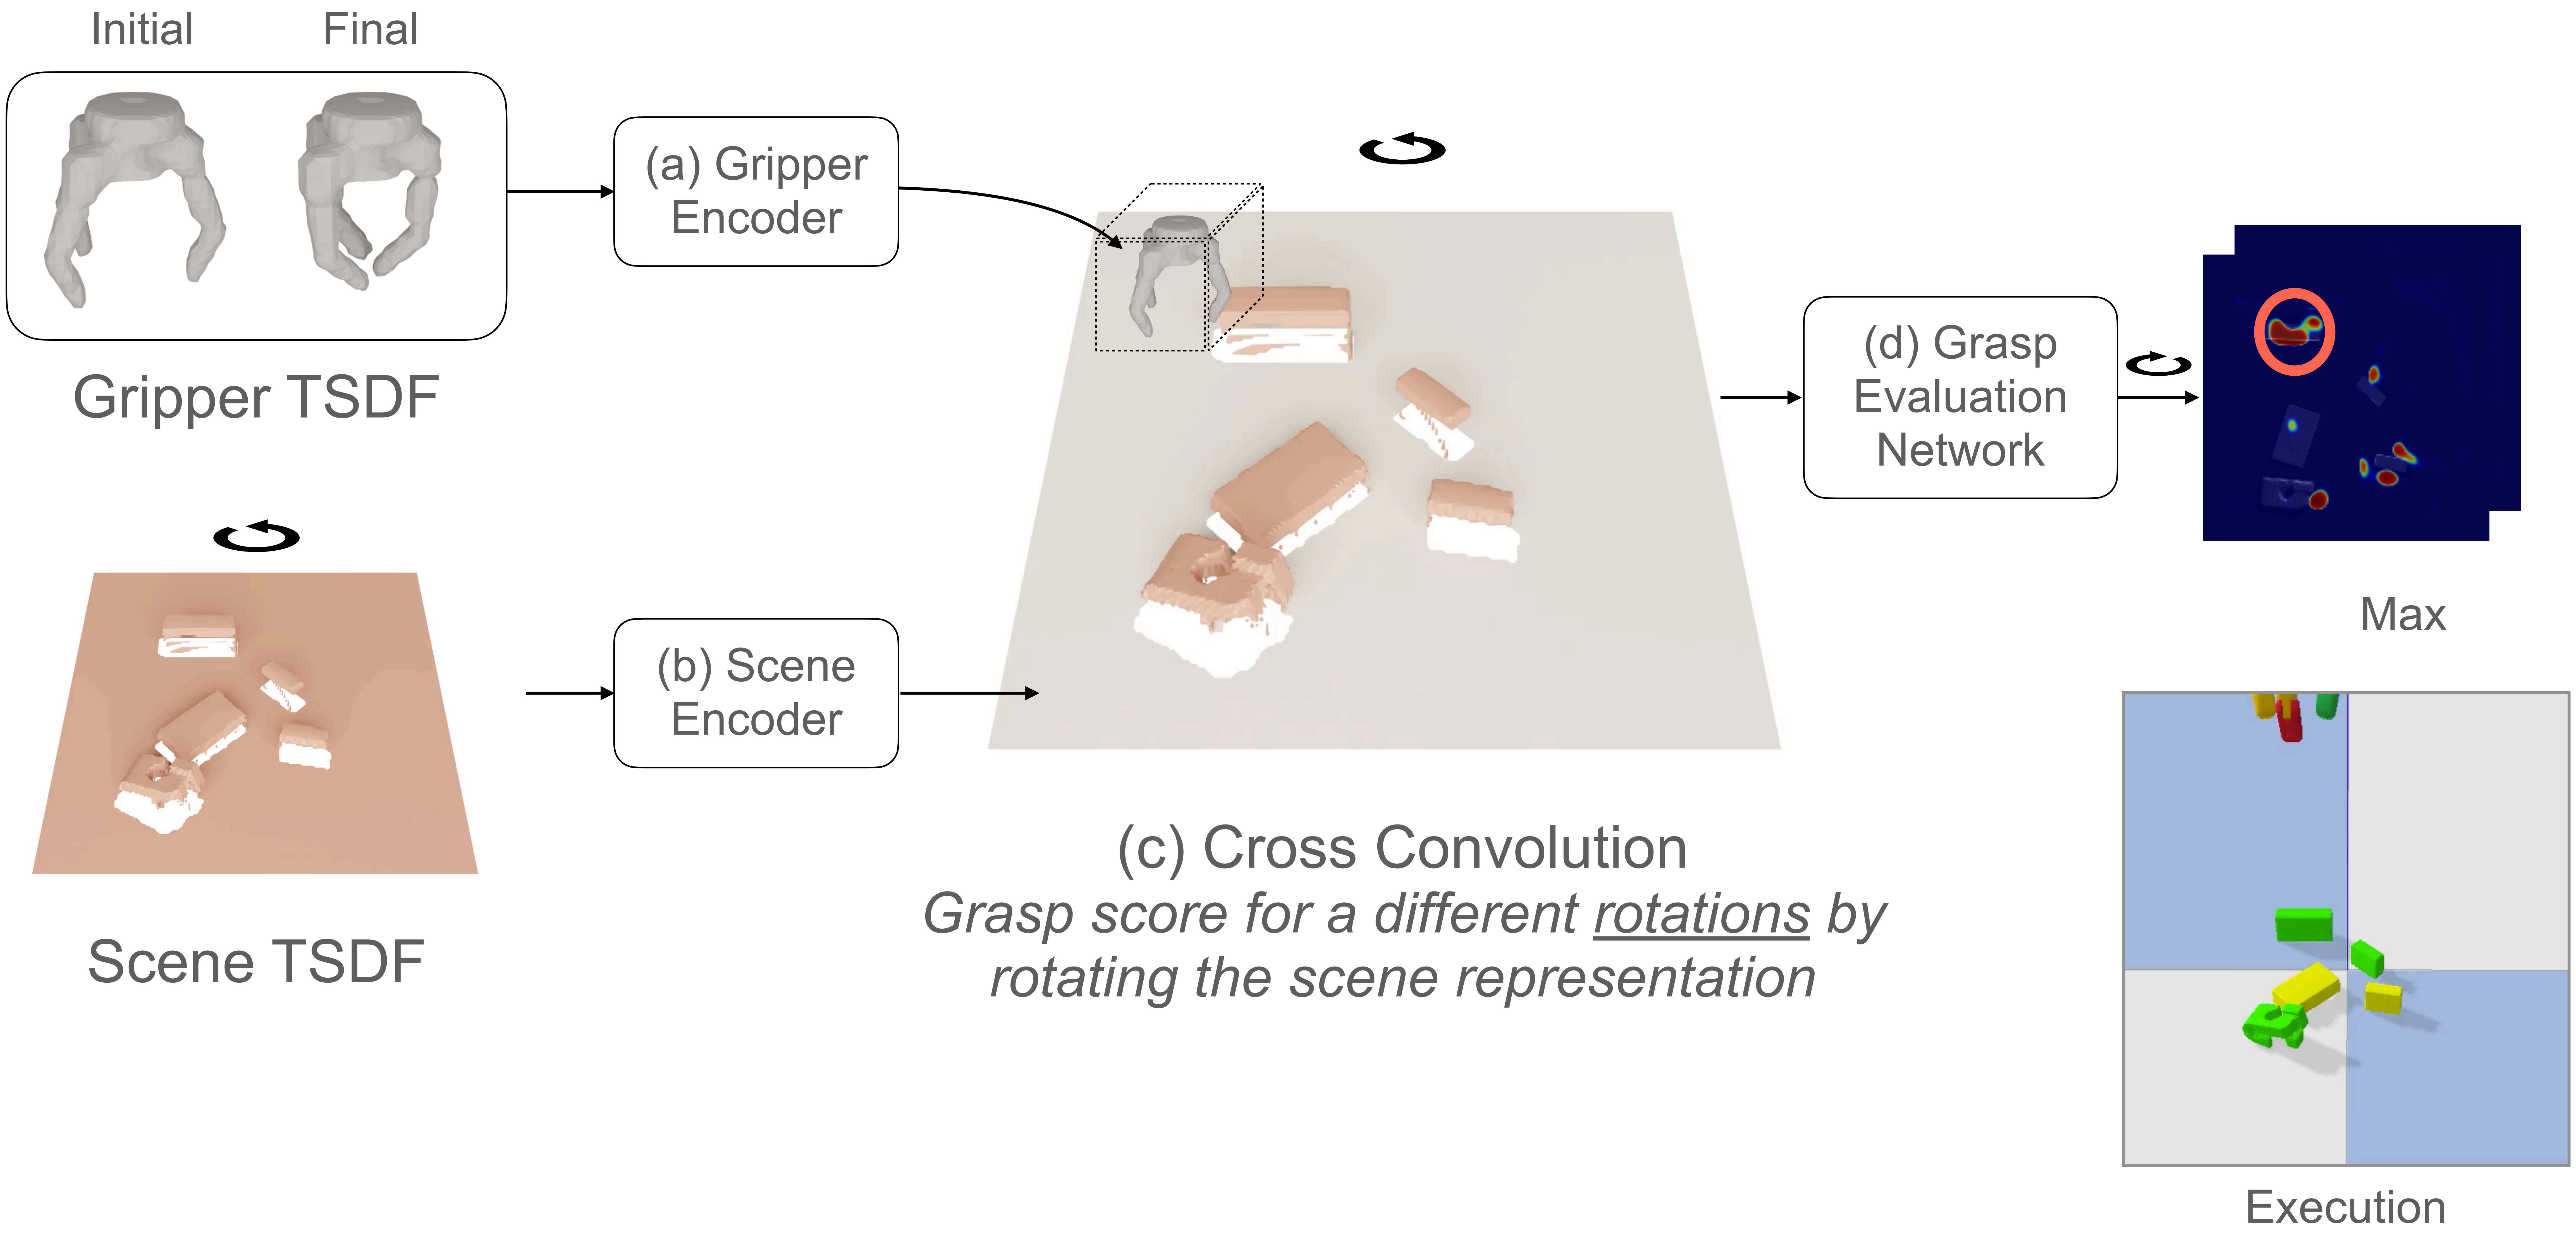

Gripper Selection

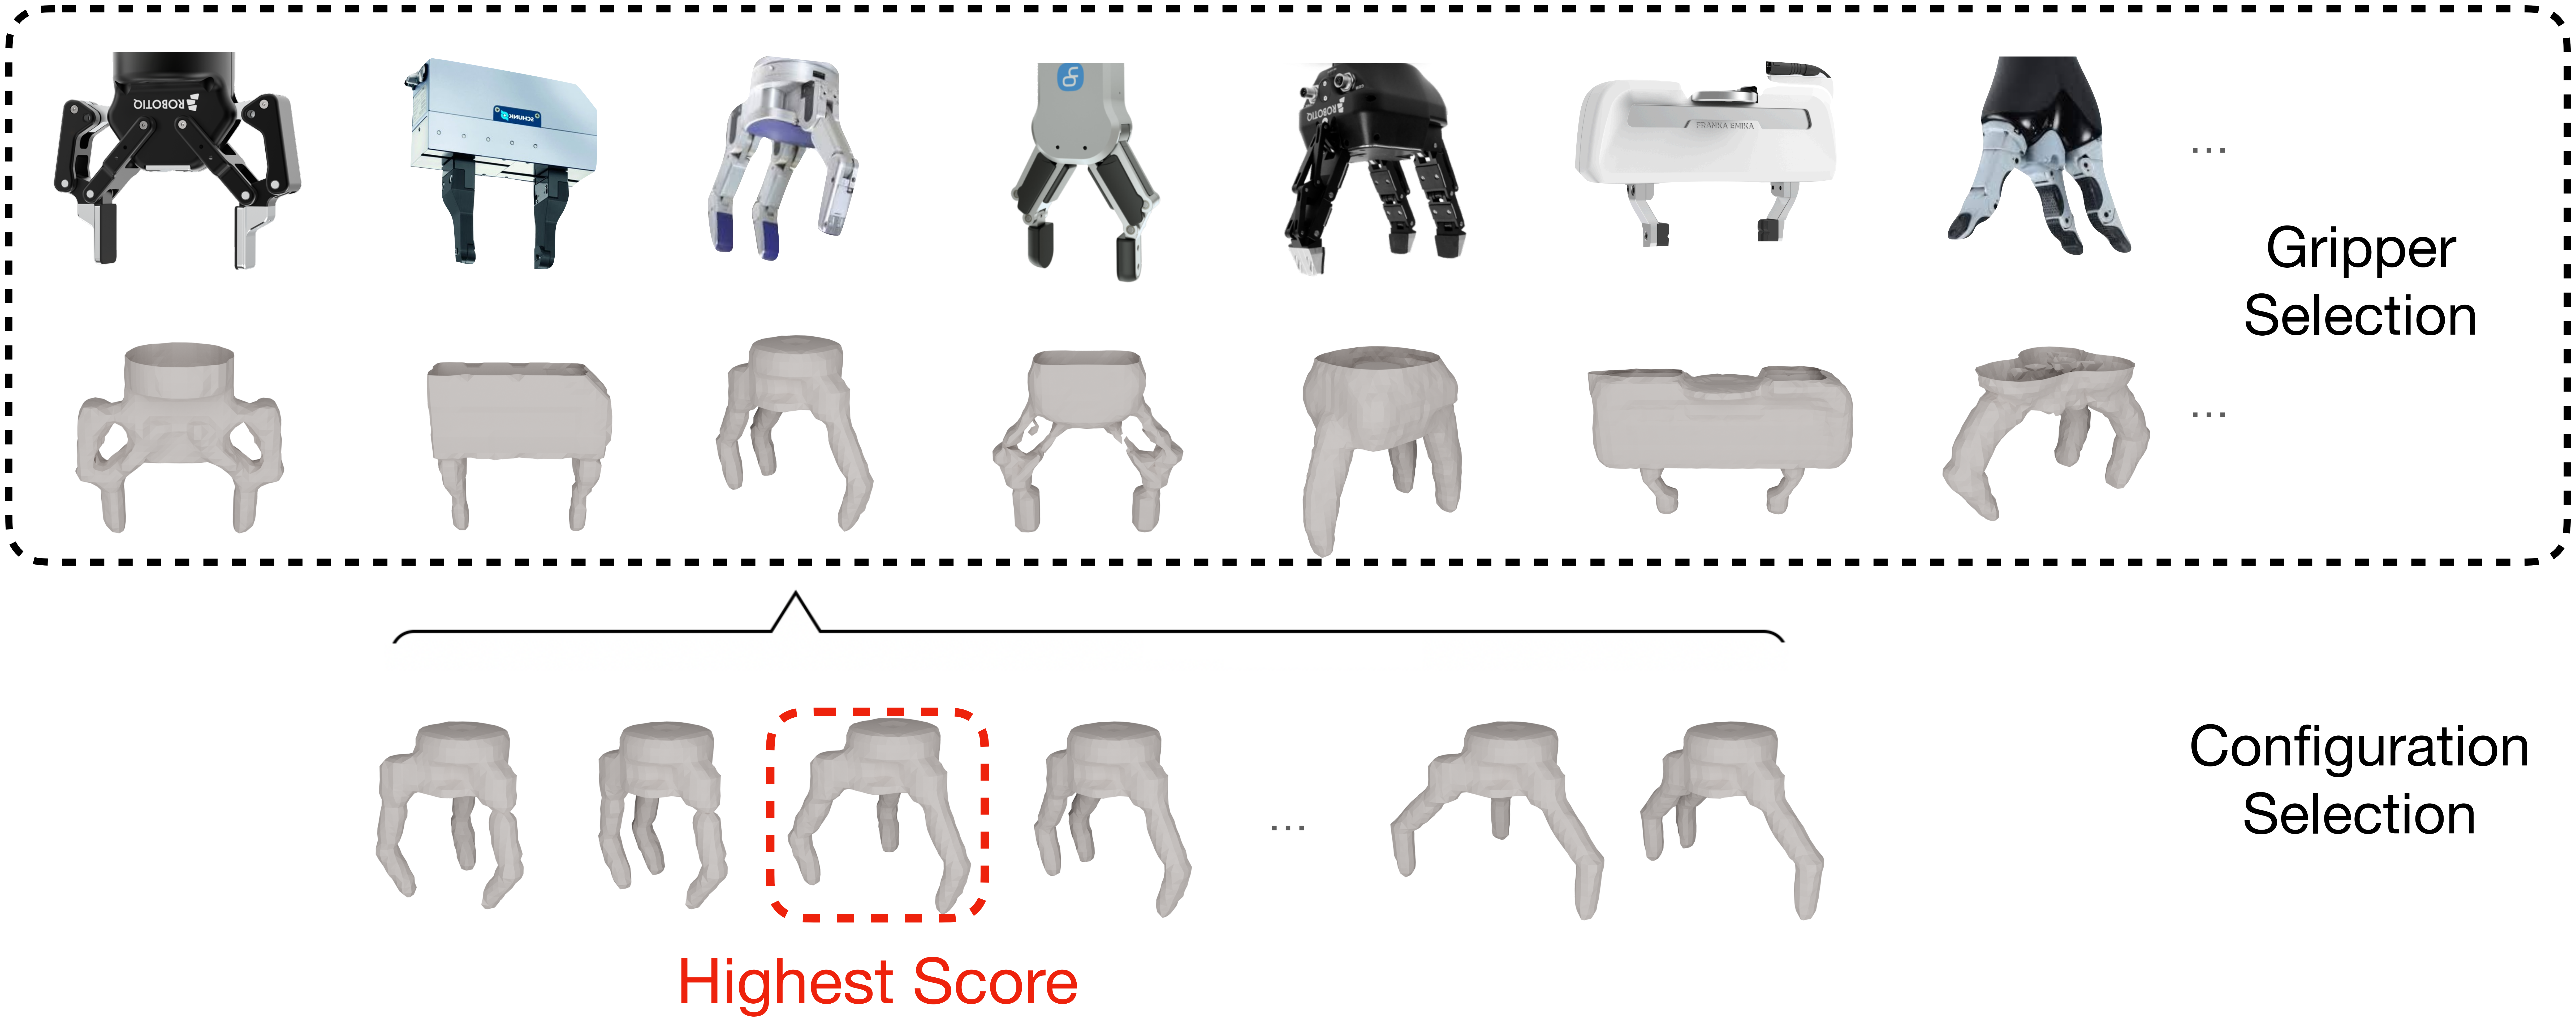

Grasp Score Prediction

Scene

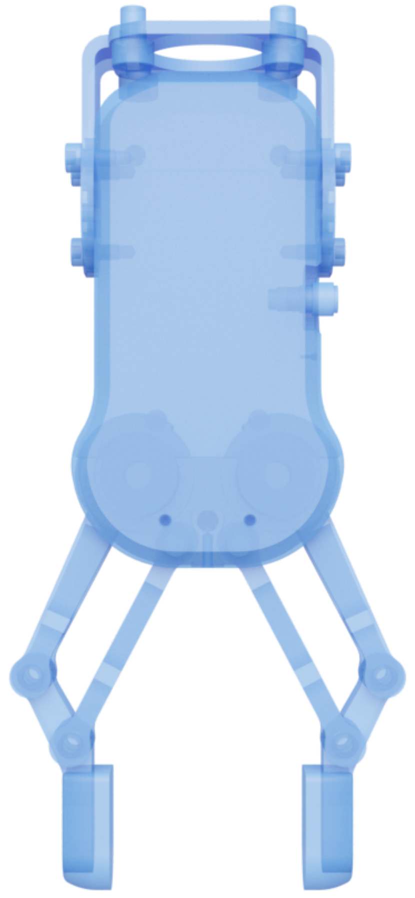

RG2

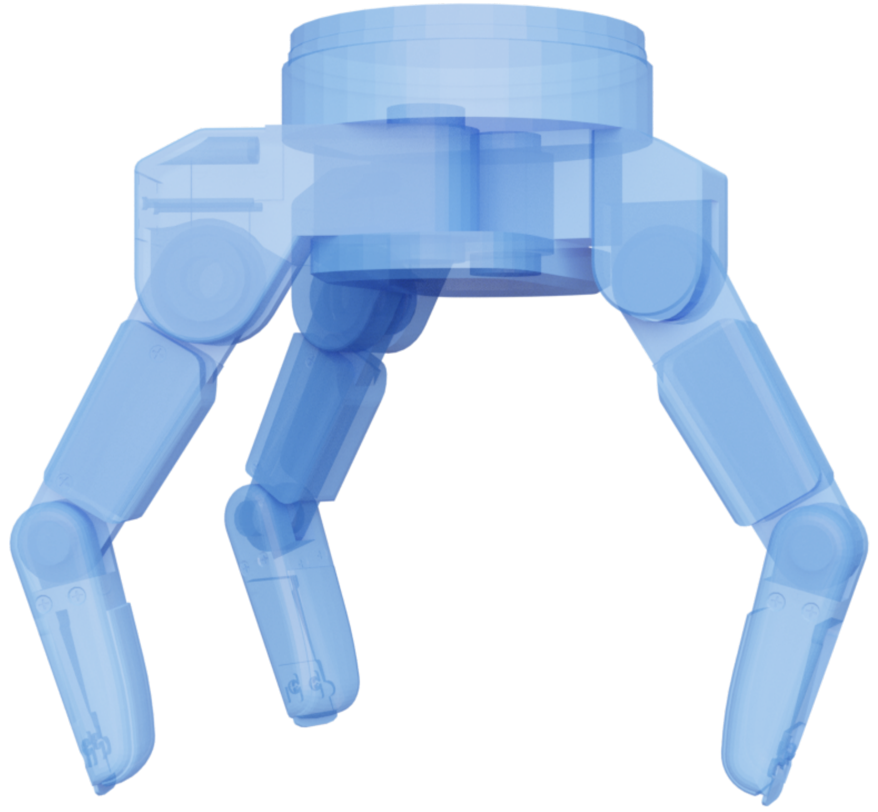

Barrett Hand

Grasp Score Prediction

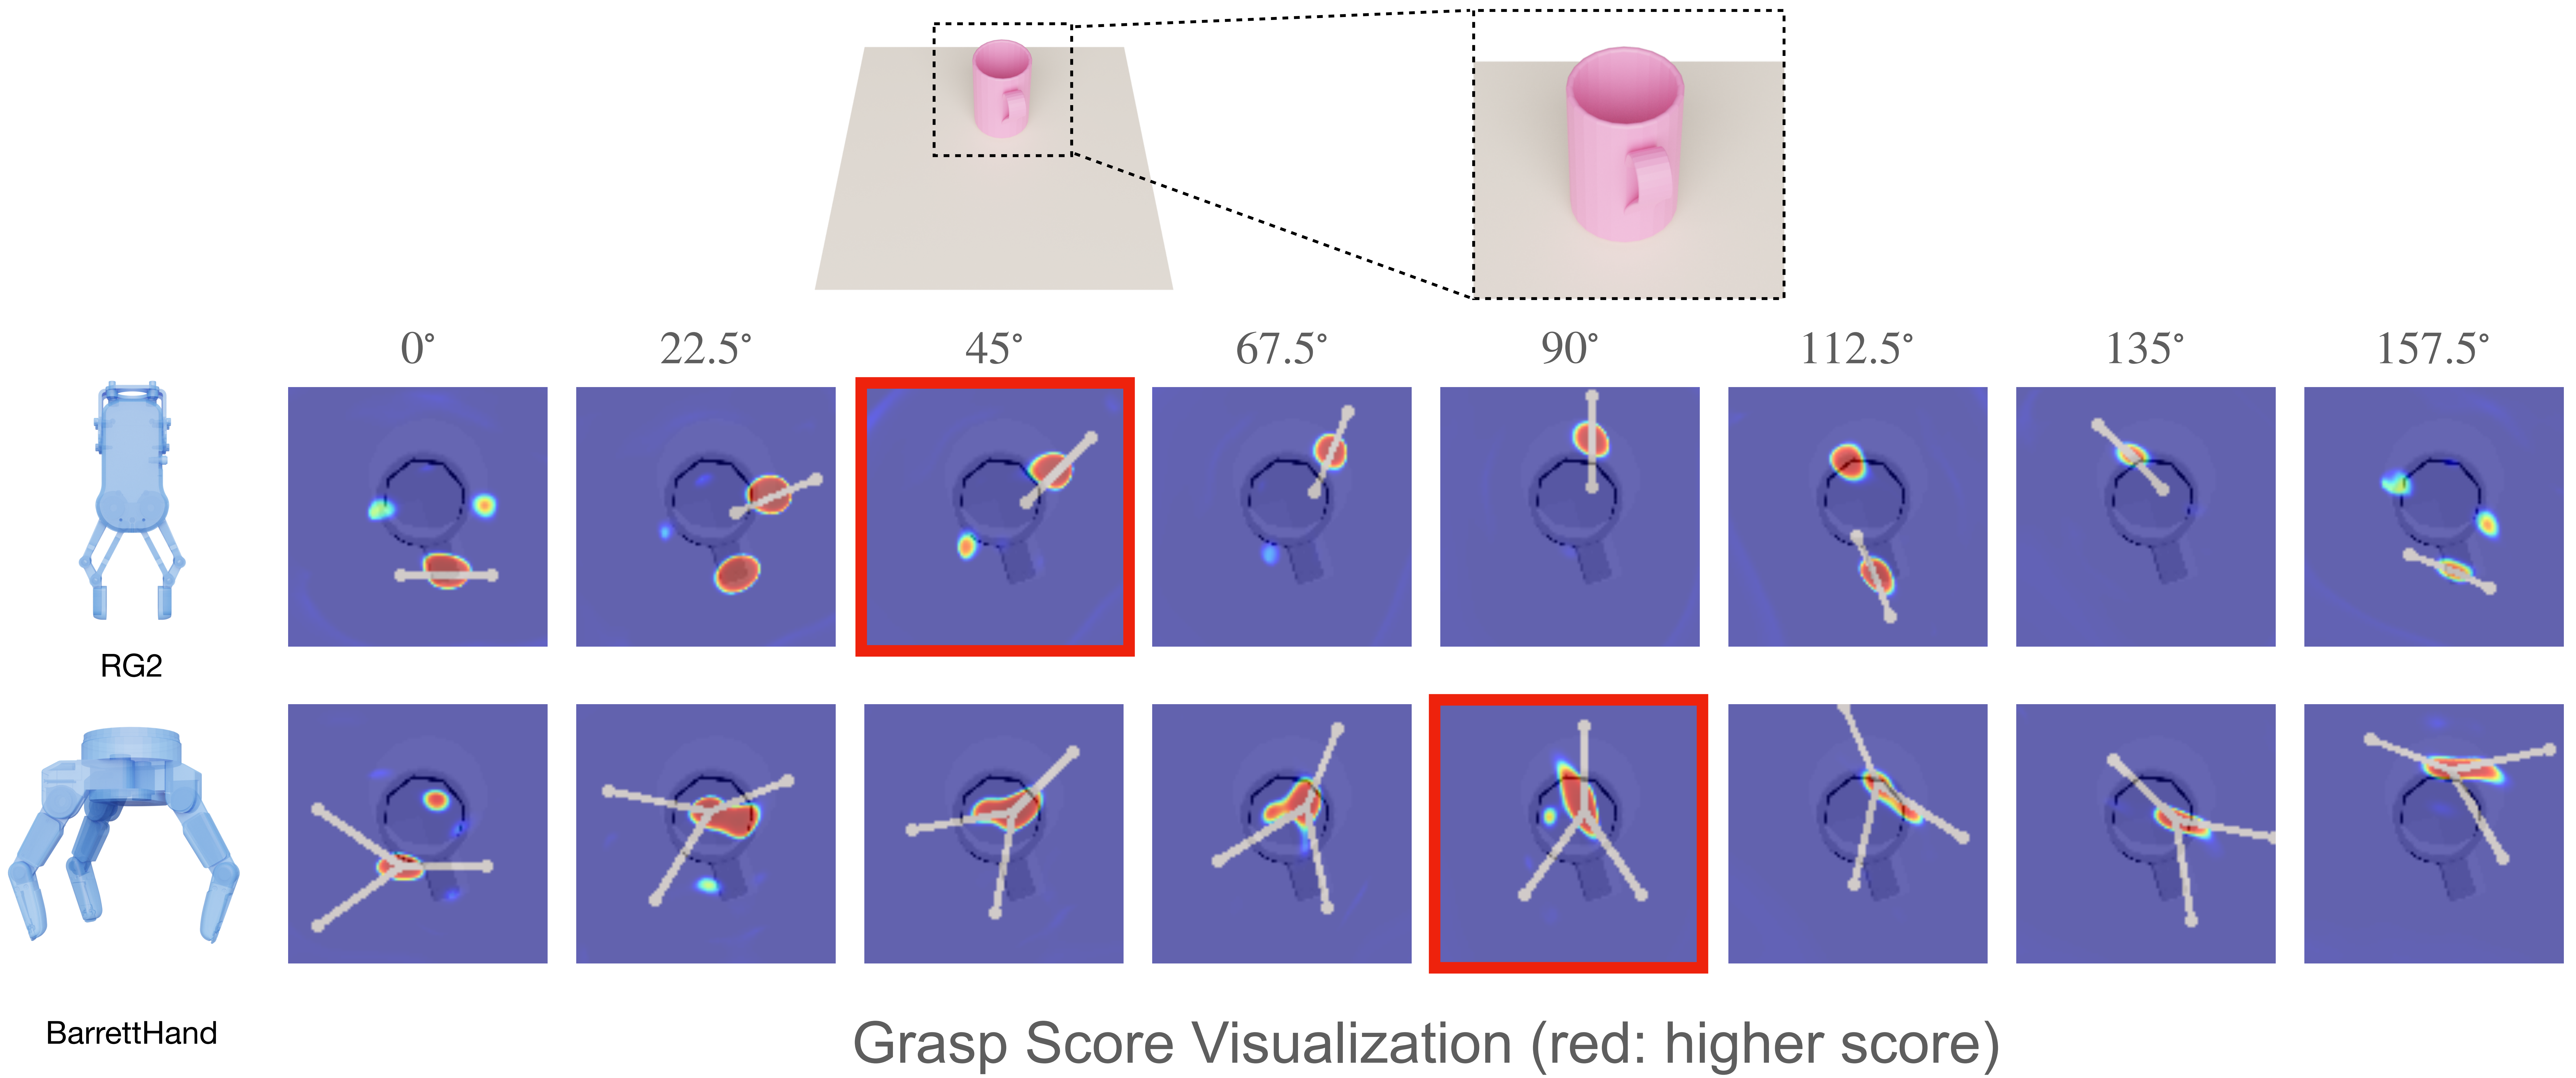

Grasp Score Visualization (red: higher score)

Grasp Score Prediction

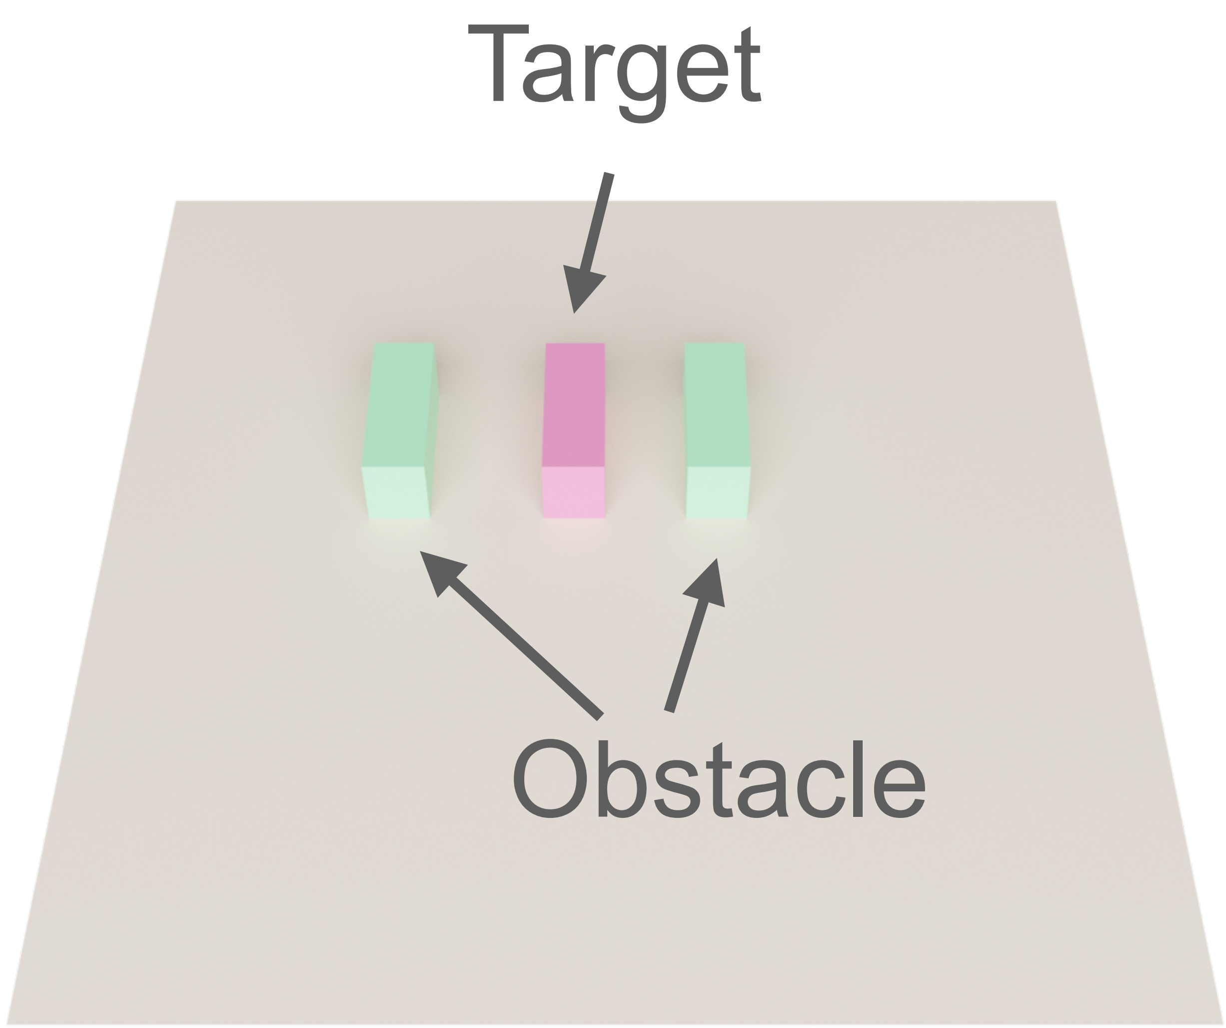

Scene

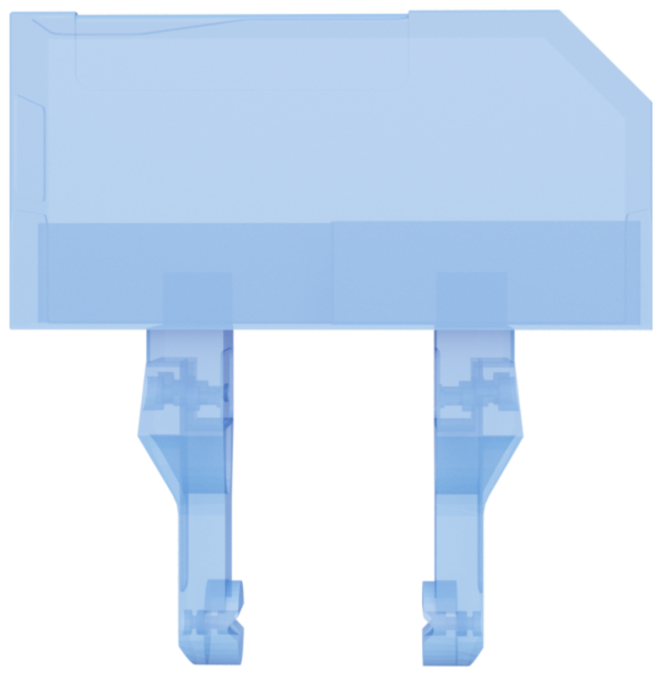

WSG 50
Half-open

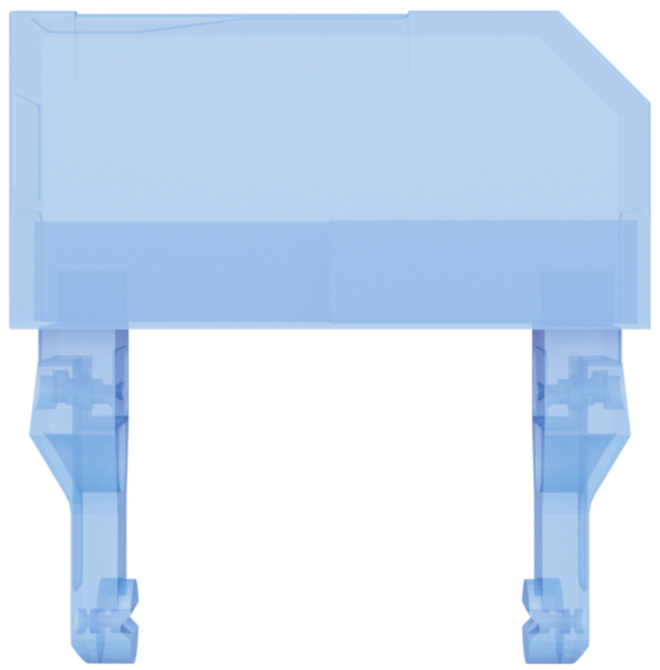

WSG 50
Fully-open

Grasp Score Prediction

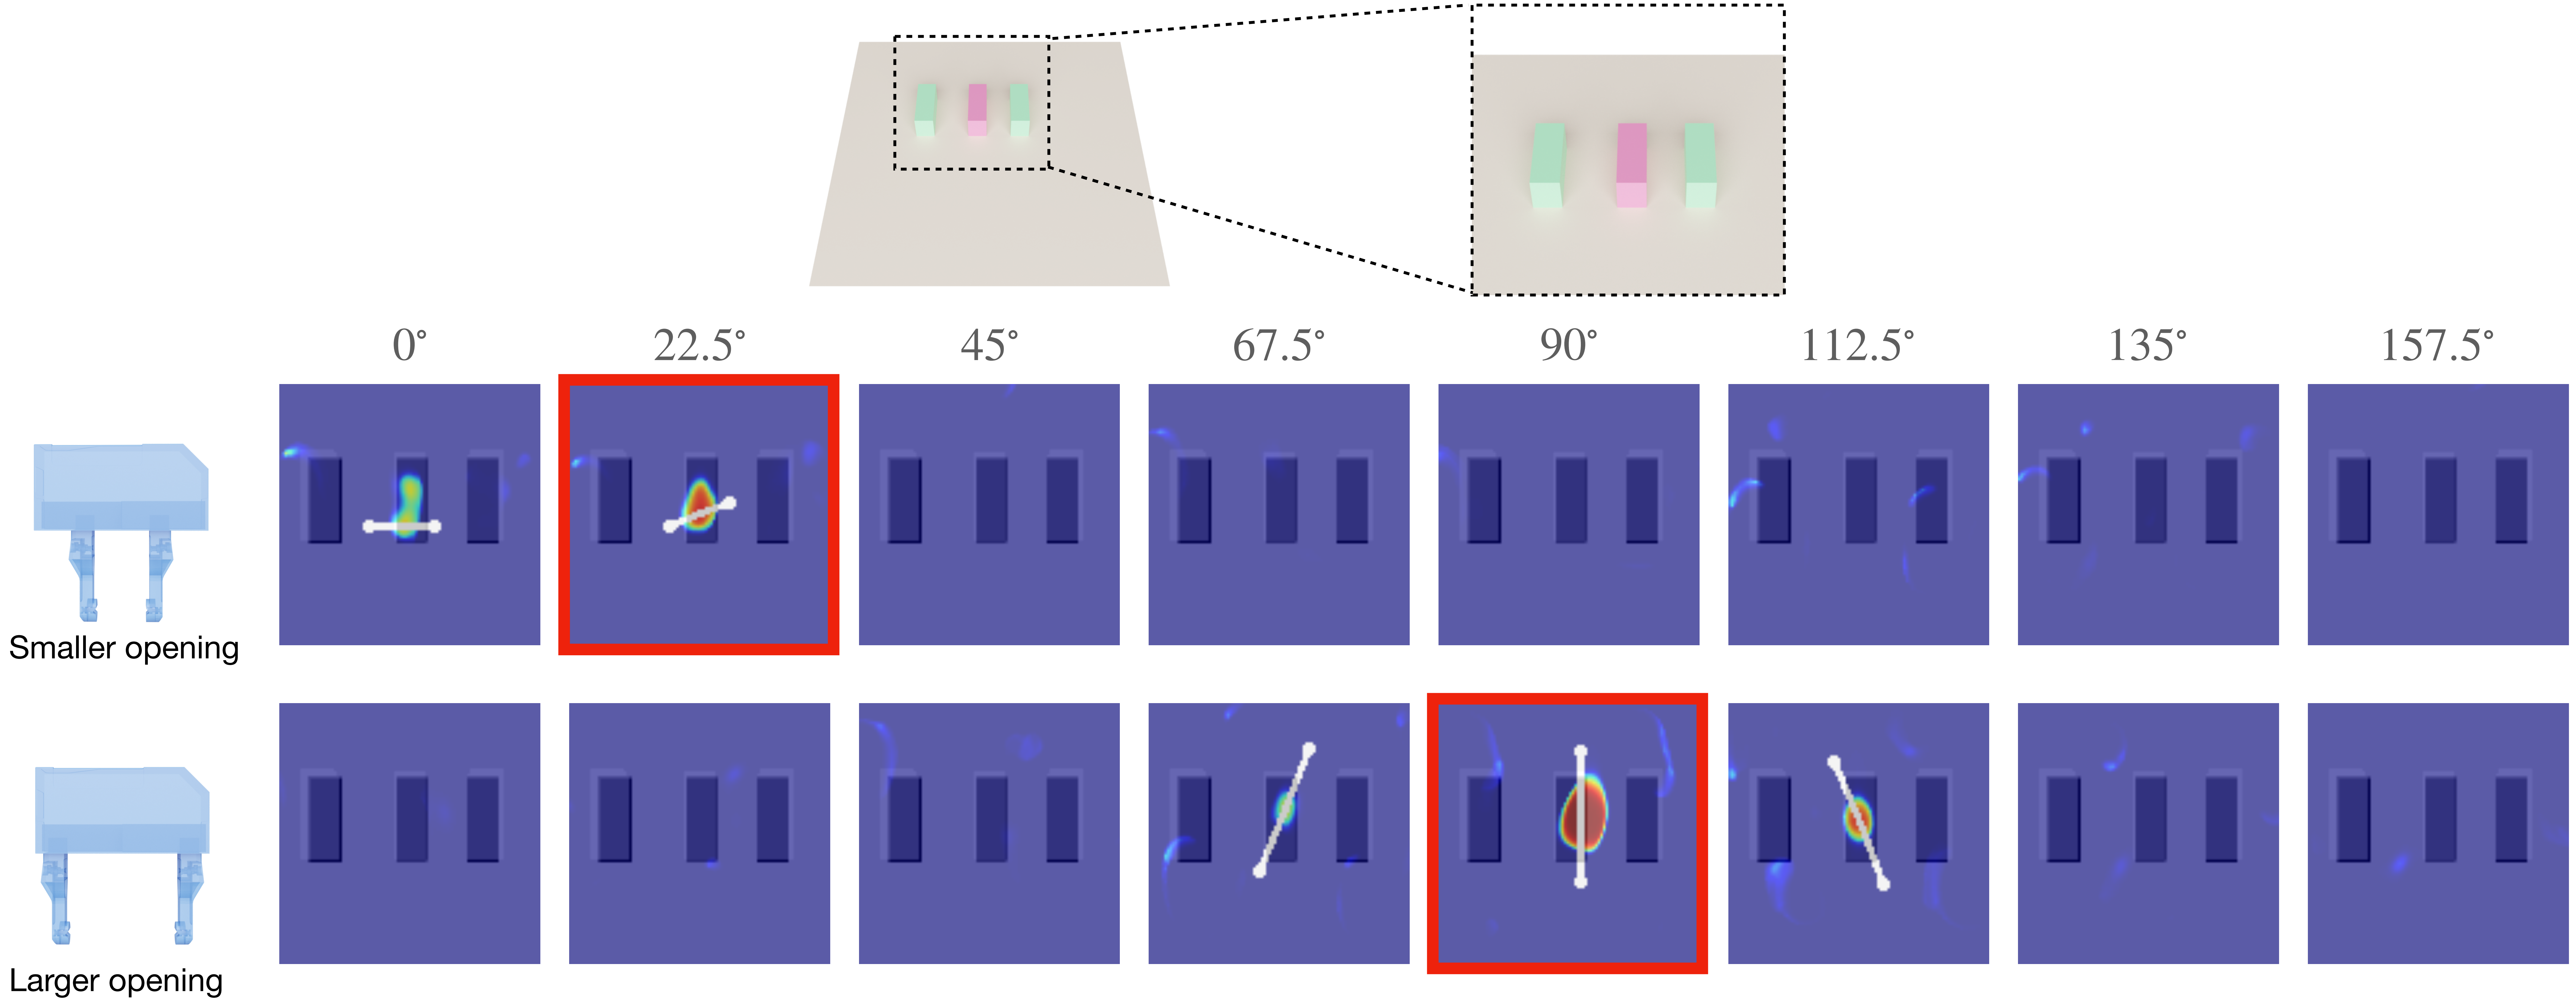

Grasp Score Visualization (red: higher score)

Comparisons to Prior Work

UniGrasp: Learning a Unified Model to Grasp with Multifingered Robotic Hands.
Shao et al.

Key idea: Predicting contact points on object surface that passes force closure check

Comparisons to Prior Work

Robustness to Partial Observation

AdaGrasp
(One camera)

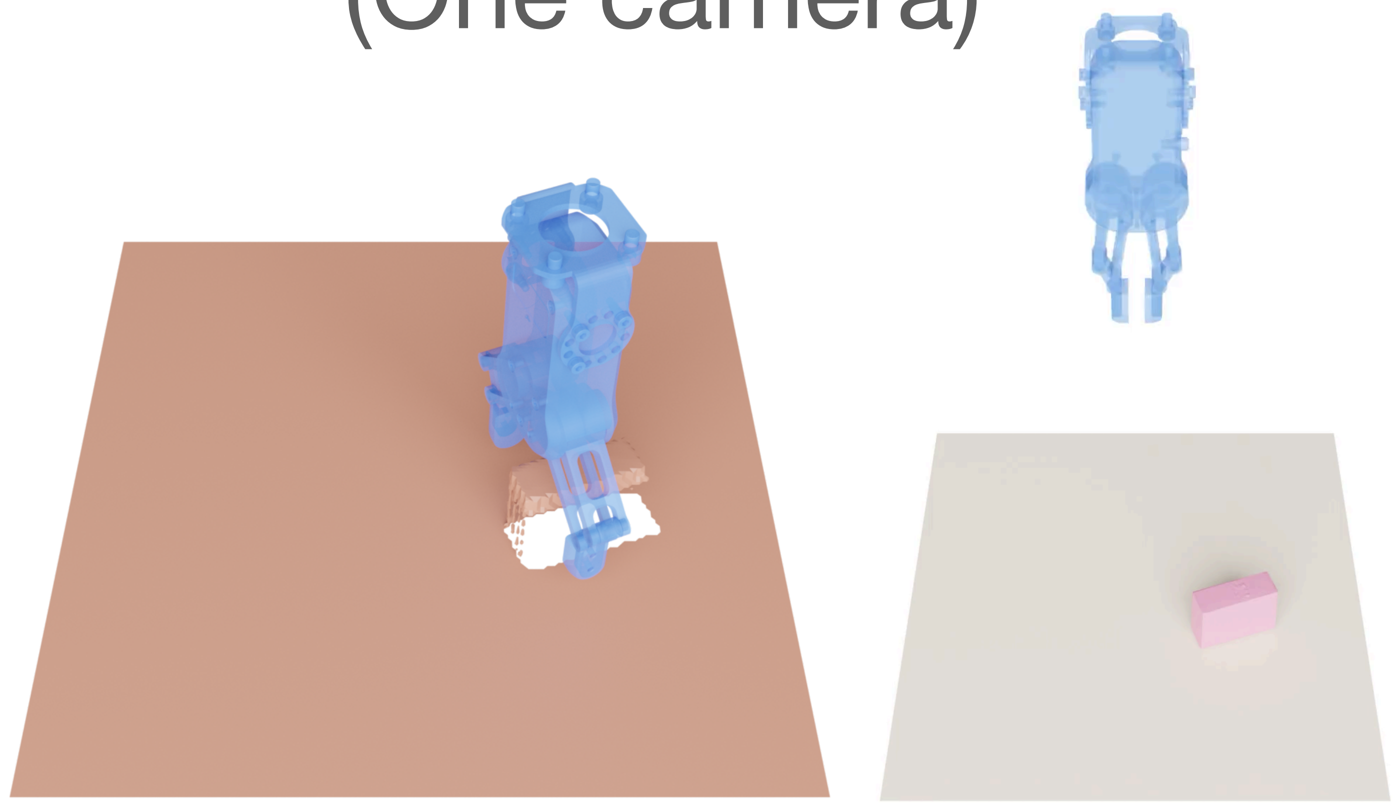

UniGrasp
(One camera)

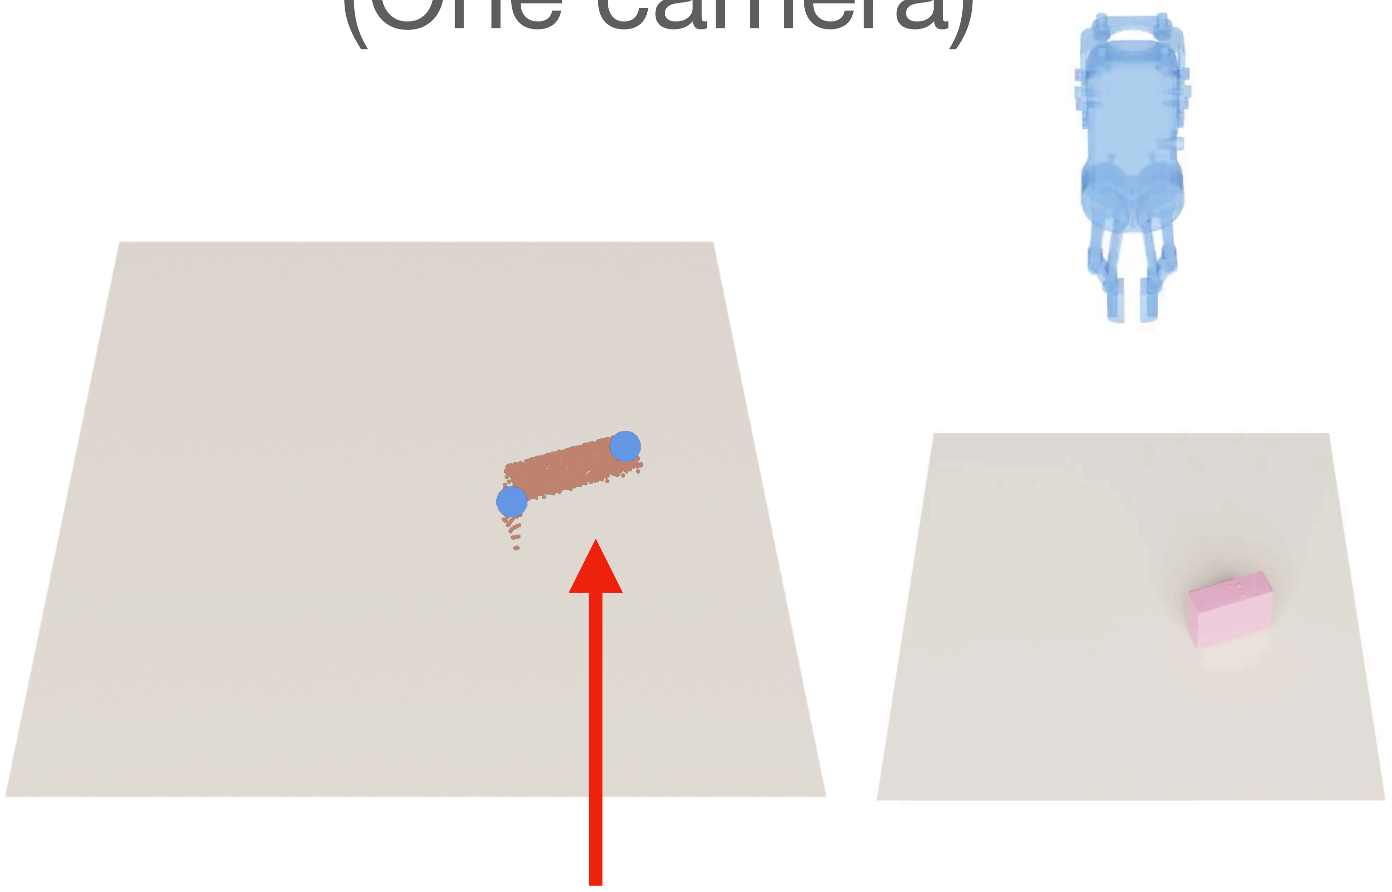

Inaccurate contact point
with partial surface input

UniGrasp
(Four cameras)

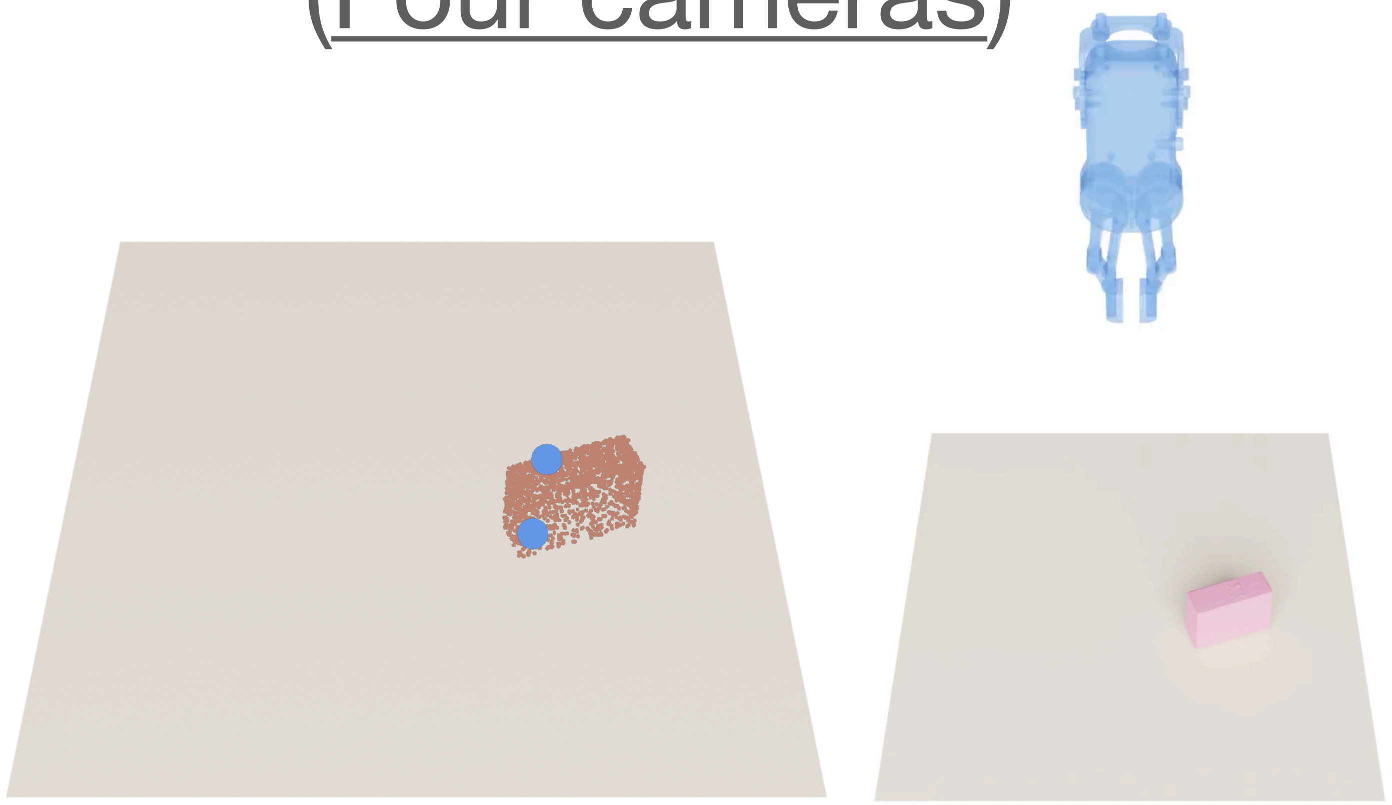

Comparisons to Prior Work

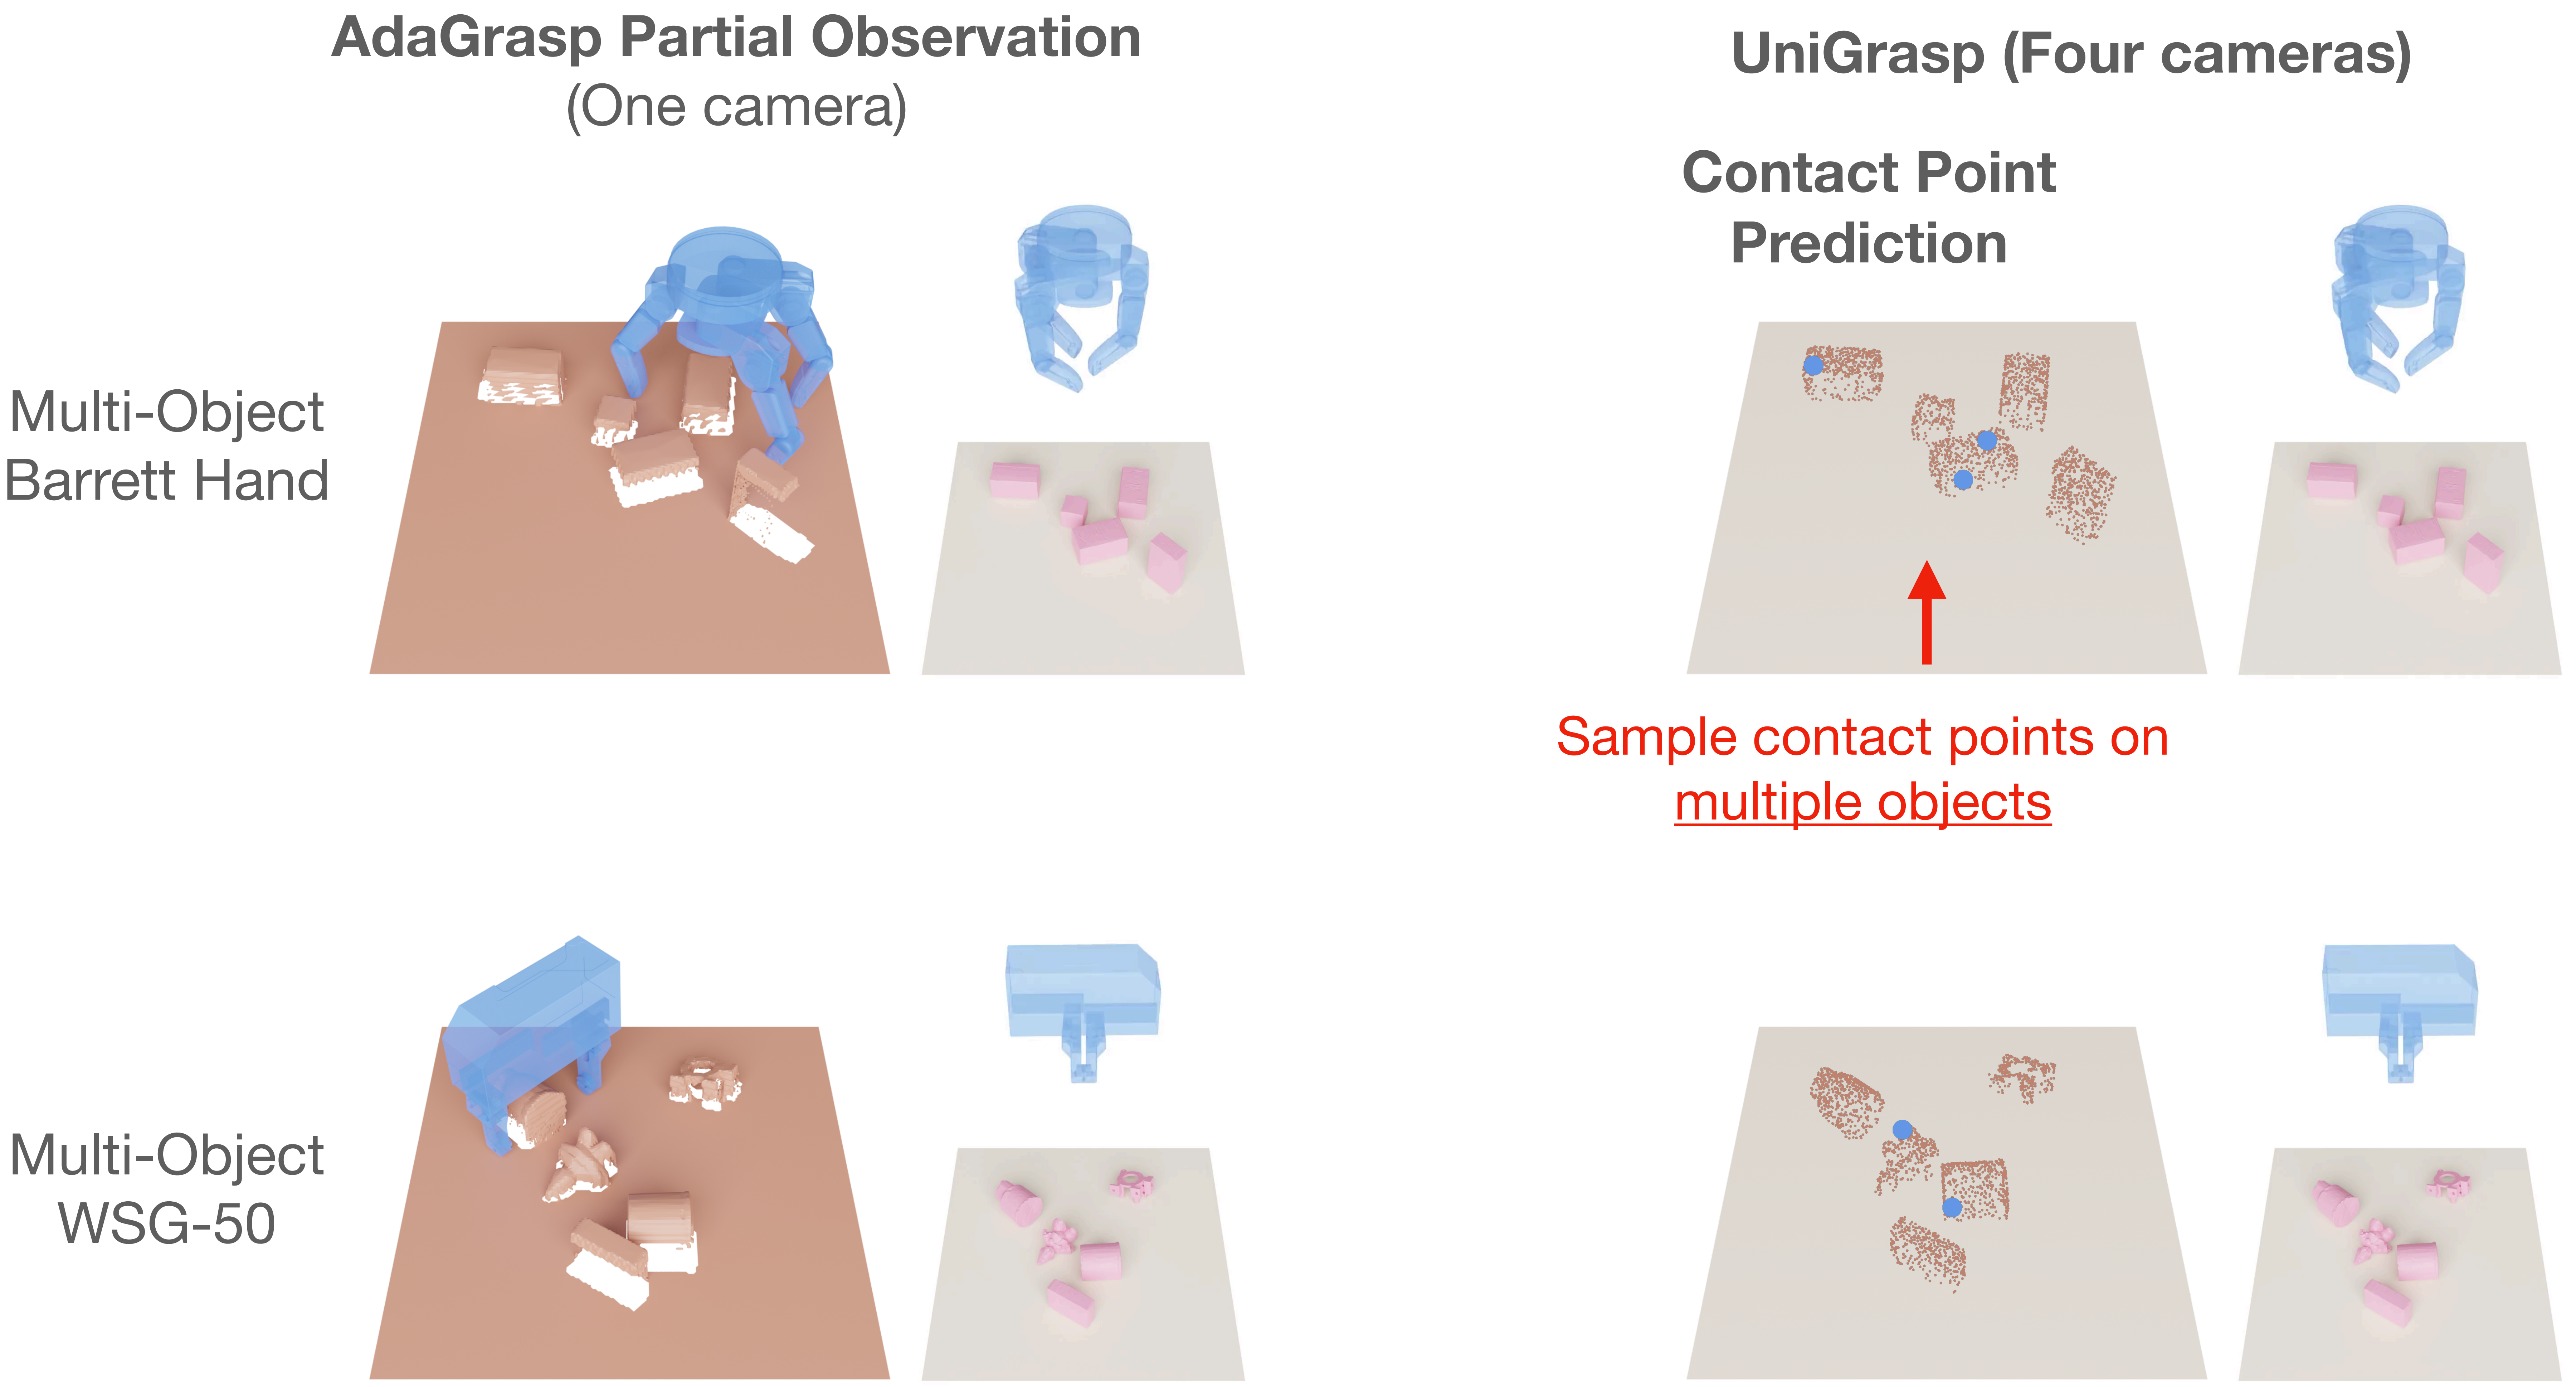

Comparisons to Prior Work

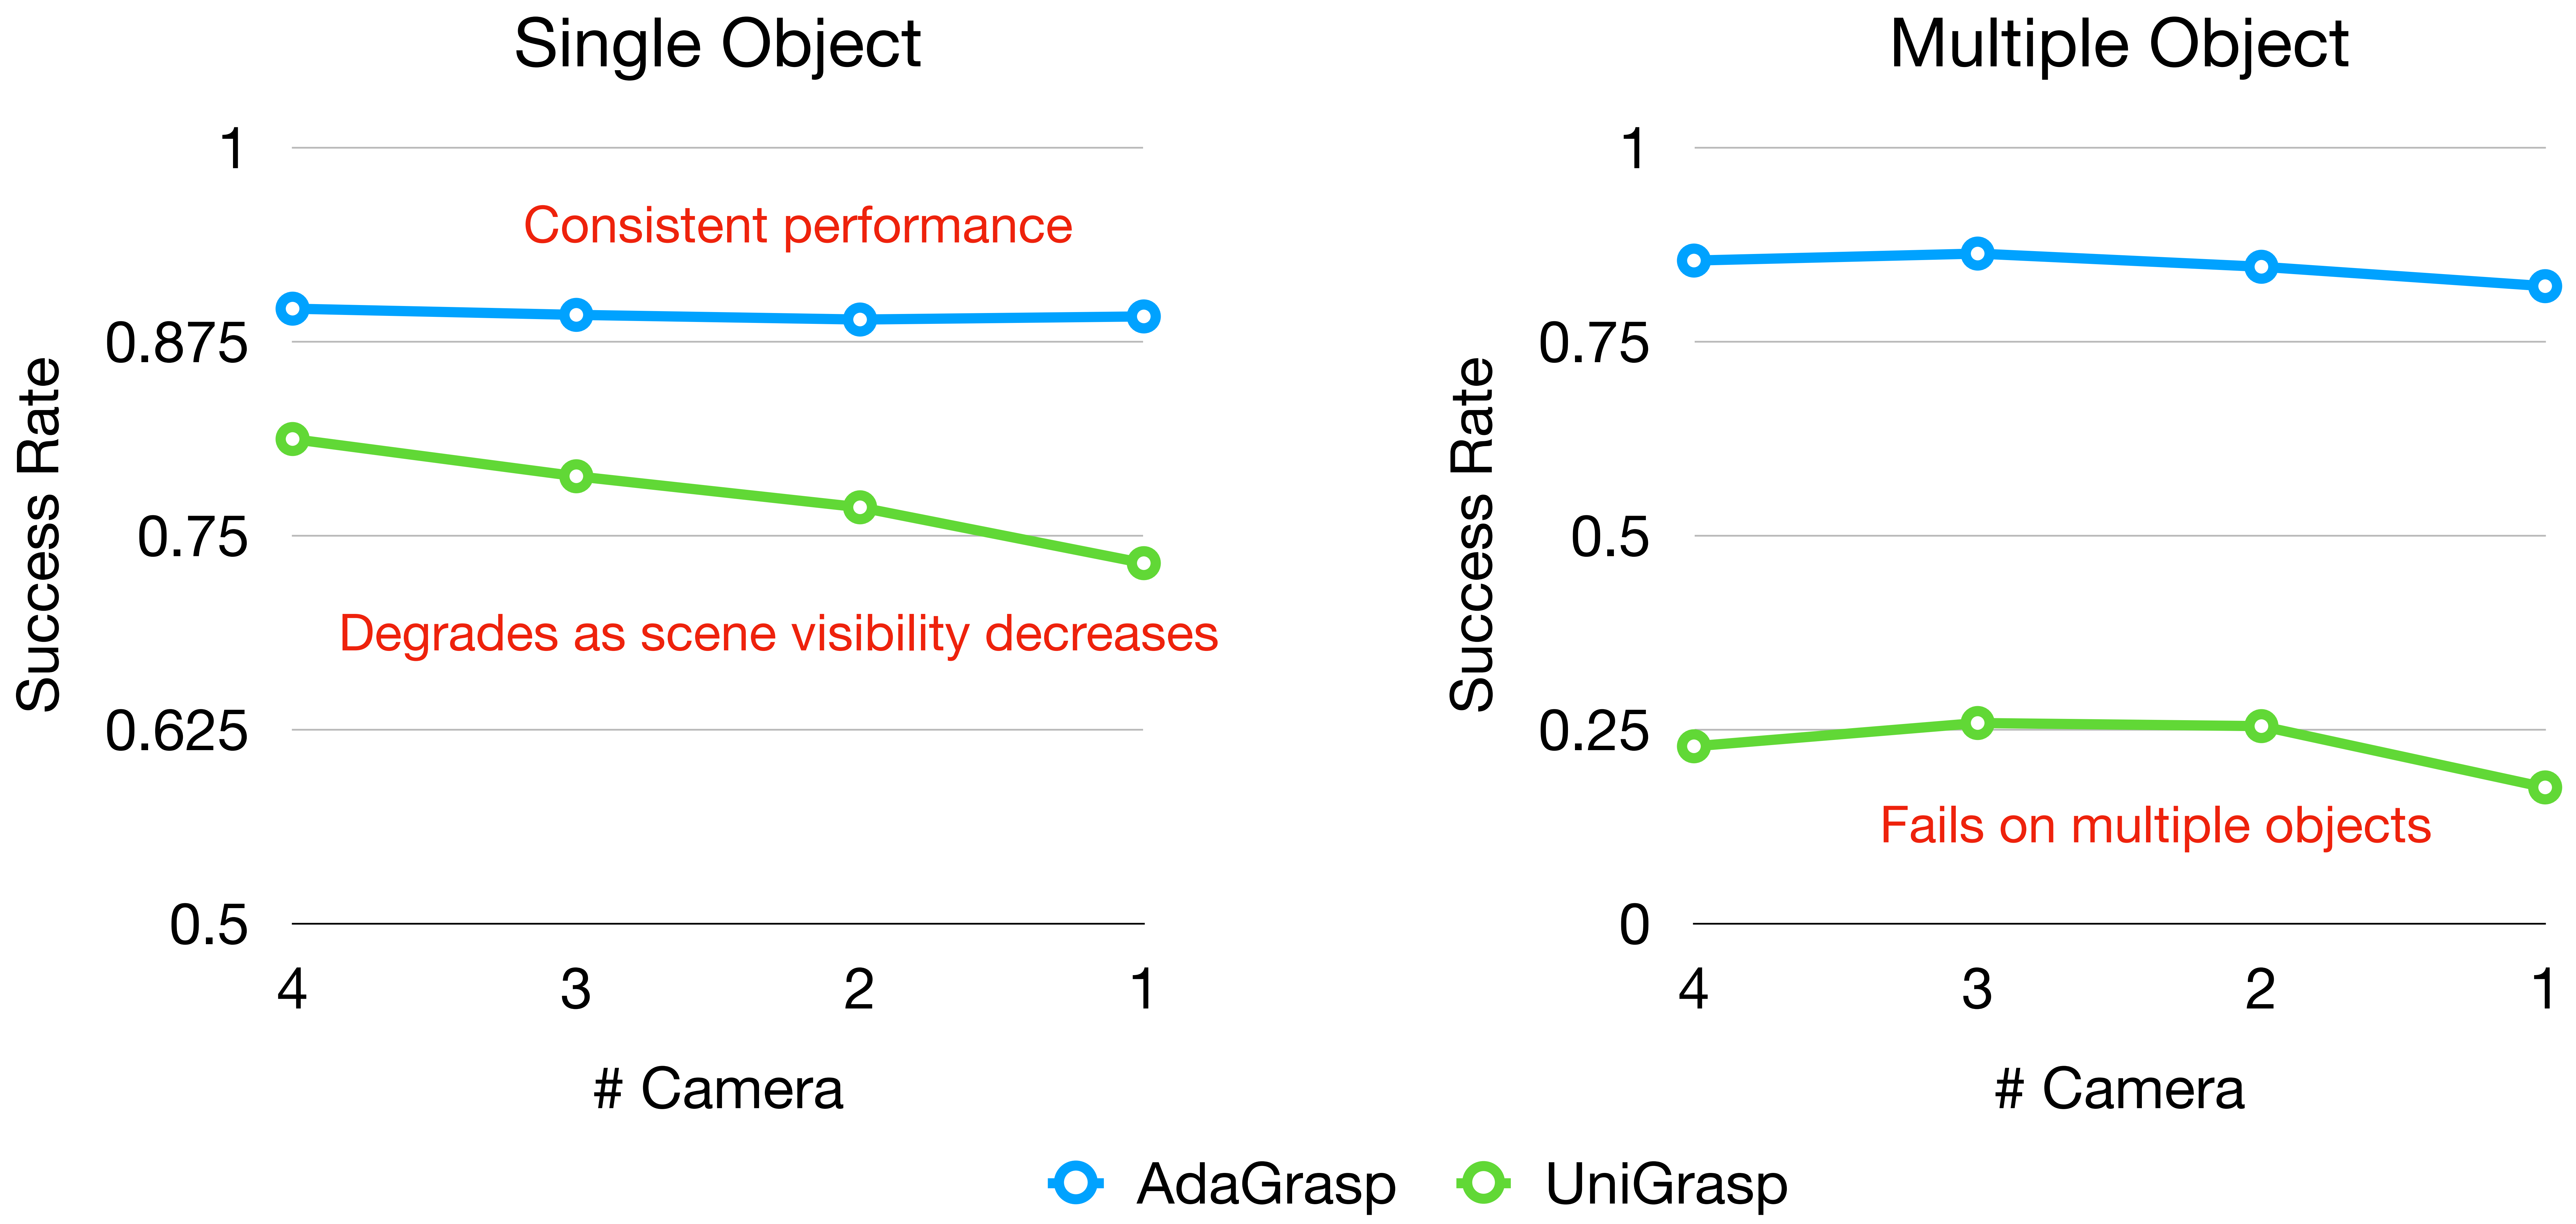

Quantitative Results.

Simulation Multiple Objects

Novel objects not seen during training

Simulation Multi-Object with Obstacles

Novel objects not seen during training. Green: obstacles

Experiments in Real World

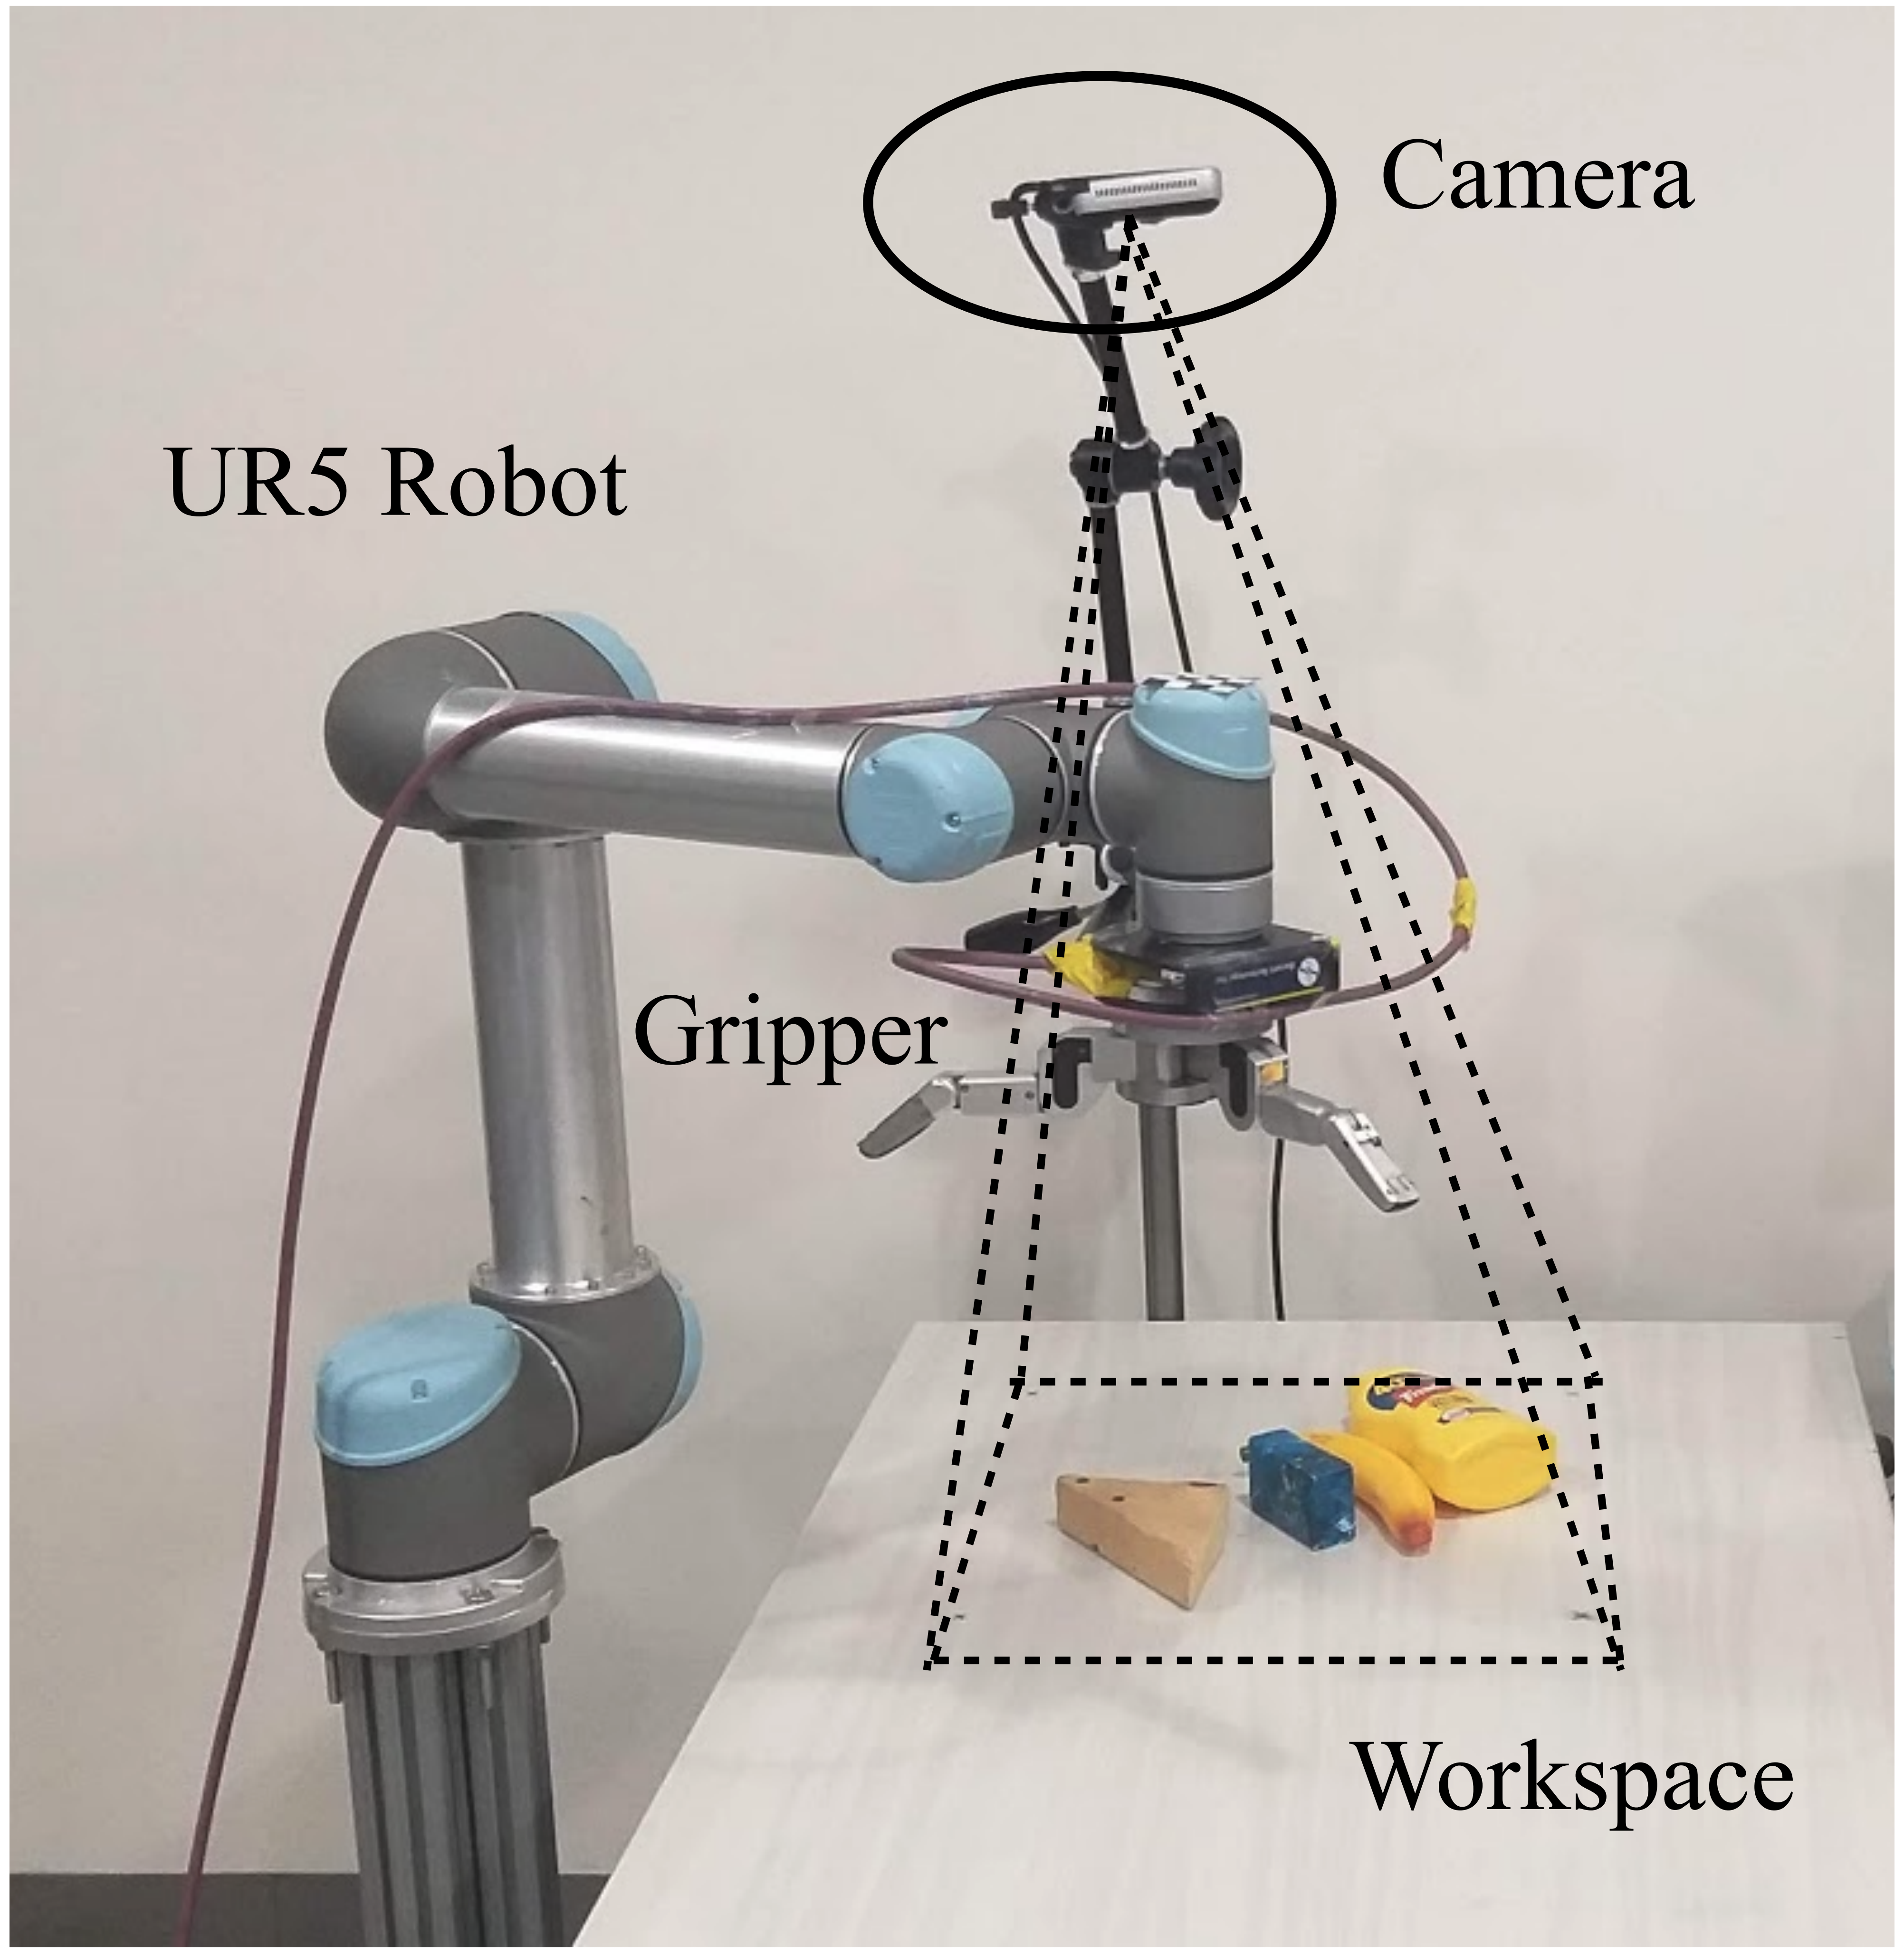

YCB dataset

DexNet 2.0

RG2

WSG 50

Barrett Hand

Barrett Hand-B

Real World Single Object

RG2 | WSG 50 | Barrett Hand-B | Barrett Hand

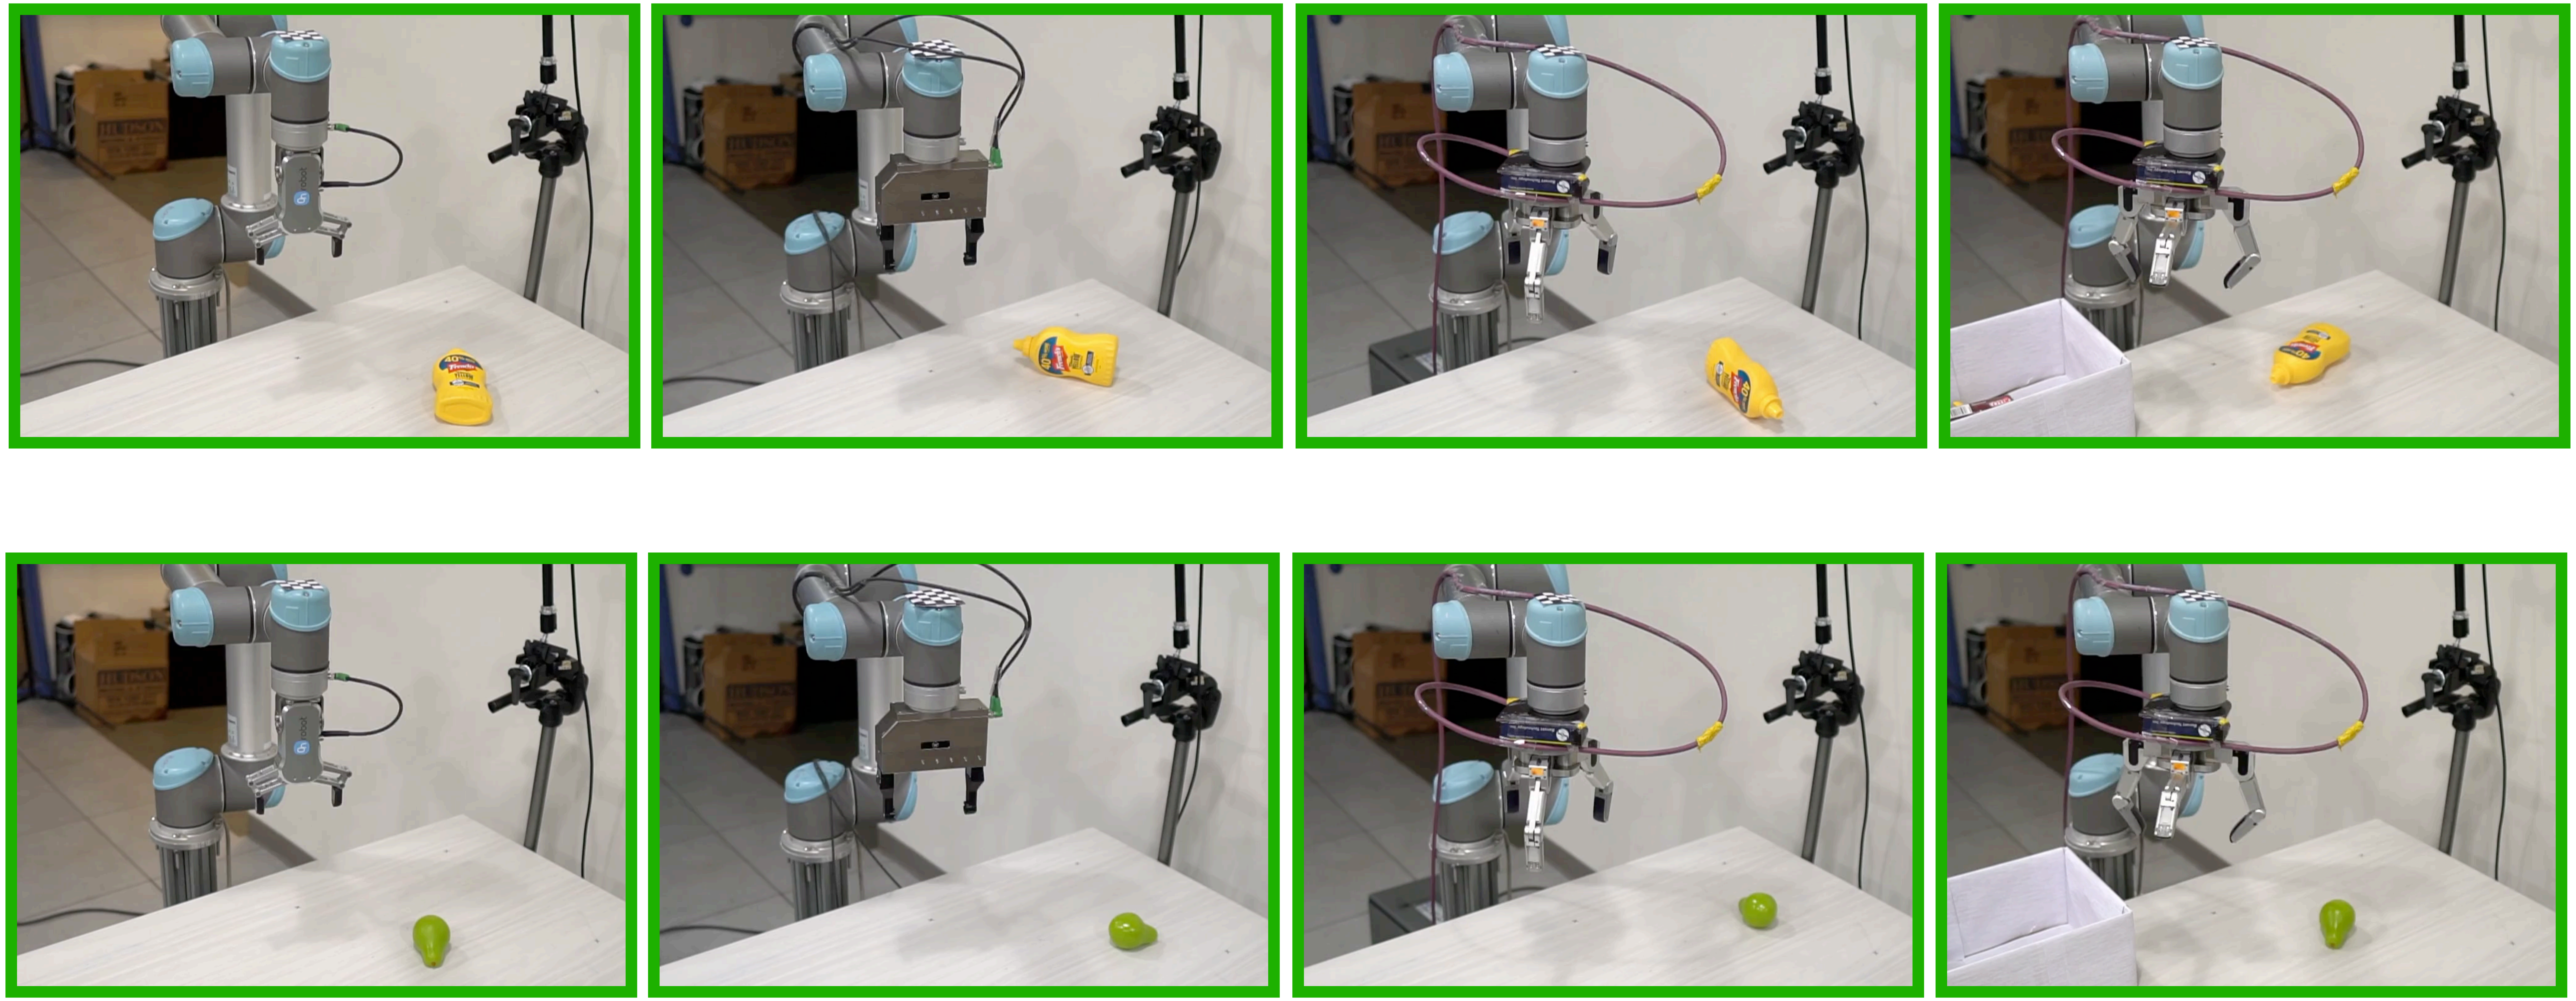

Novel objects & Novel Gripper not Seen during Training

5x speed

Real World <u>Single</u> Object

RG2 | WSG 50 | Barrett Hand-B | Barrett Hand

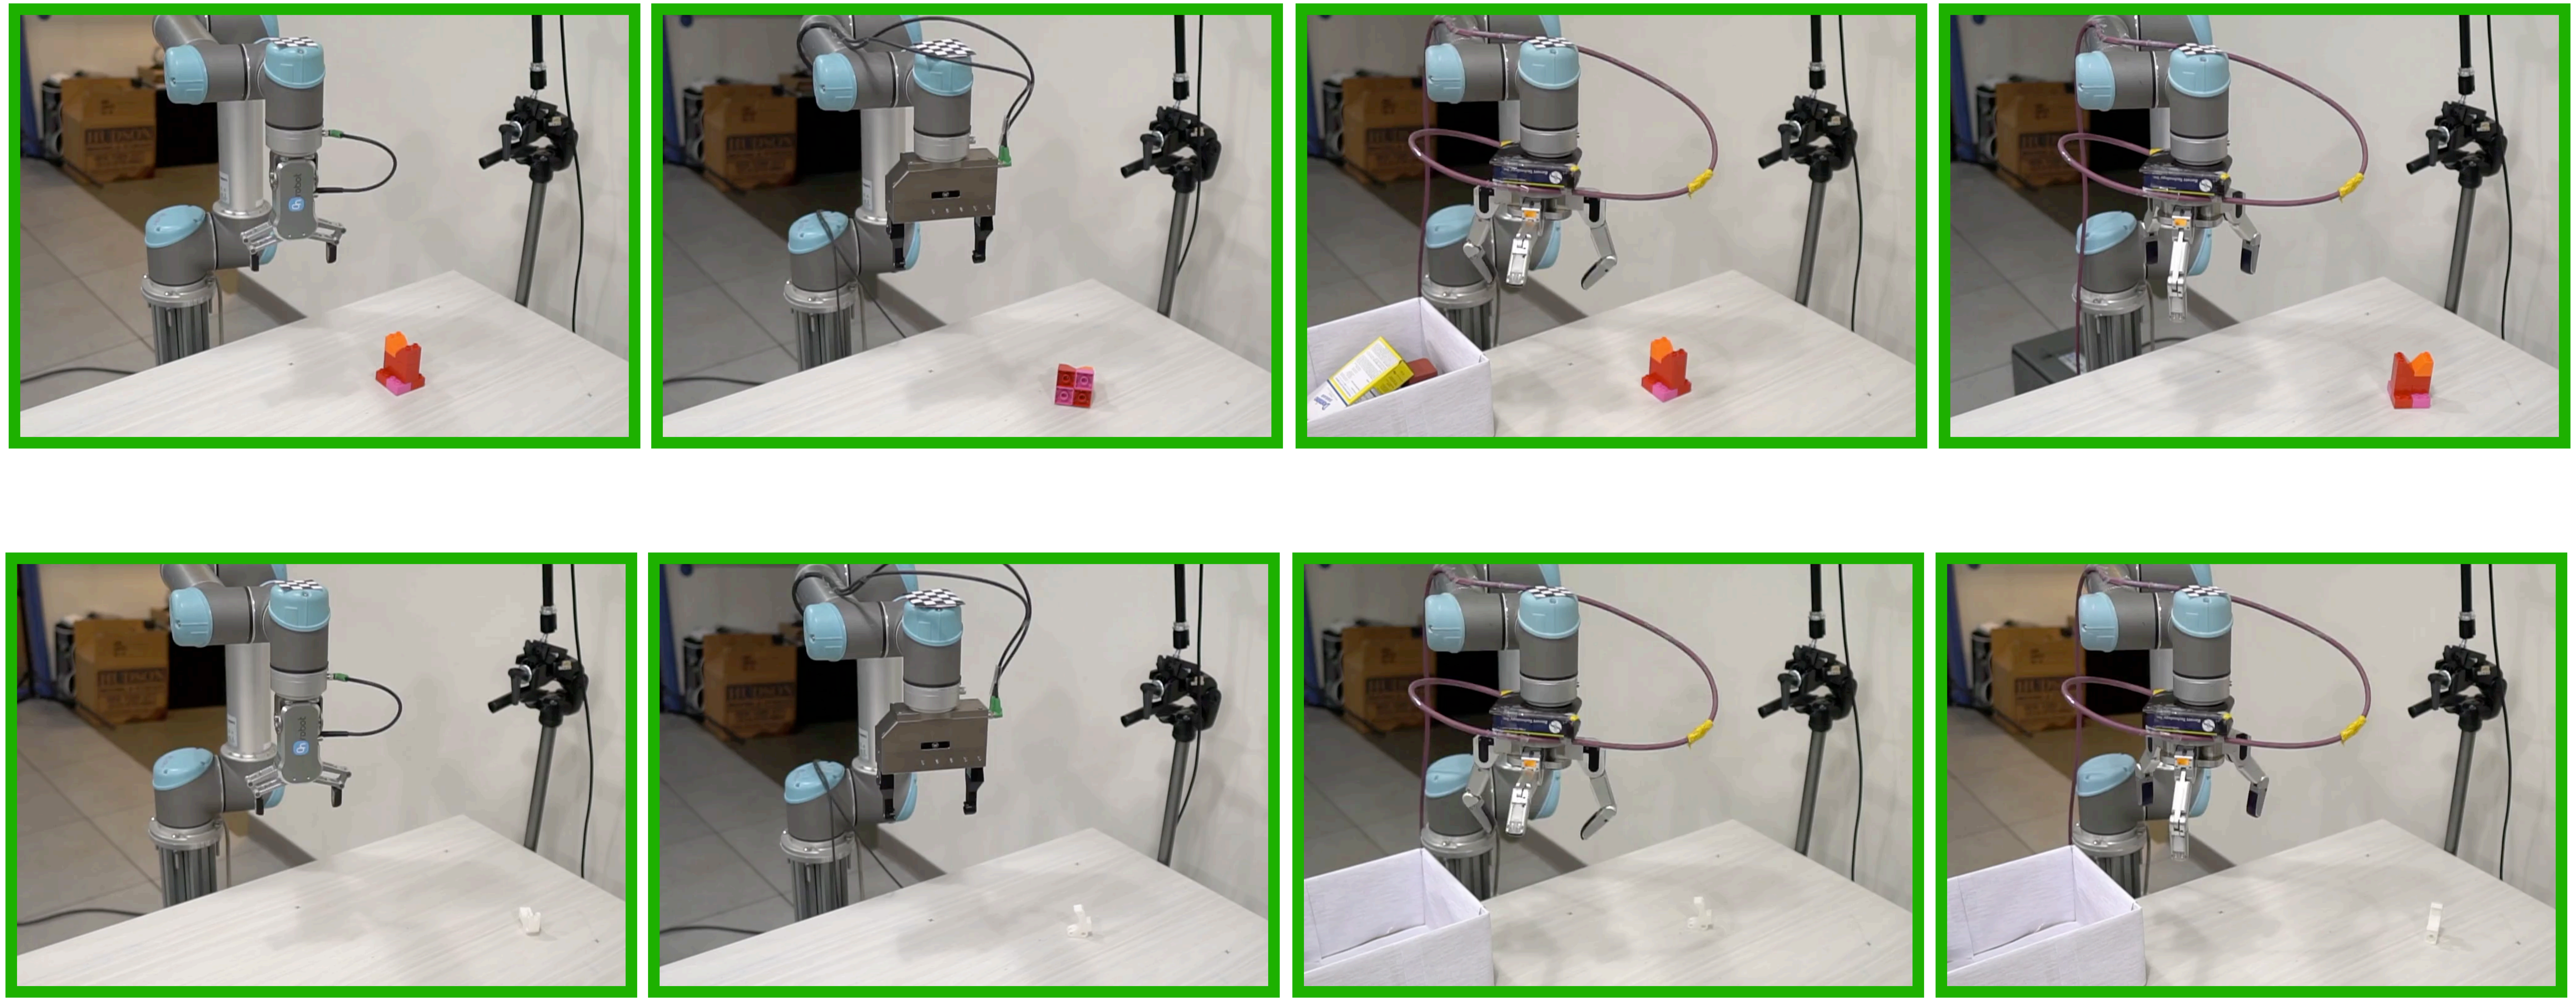

Novel objects & Novel Gripper not Seen during Training

5x speed

Real World Single Object

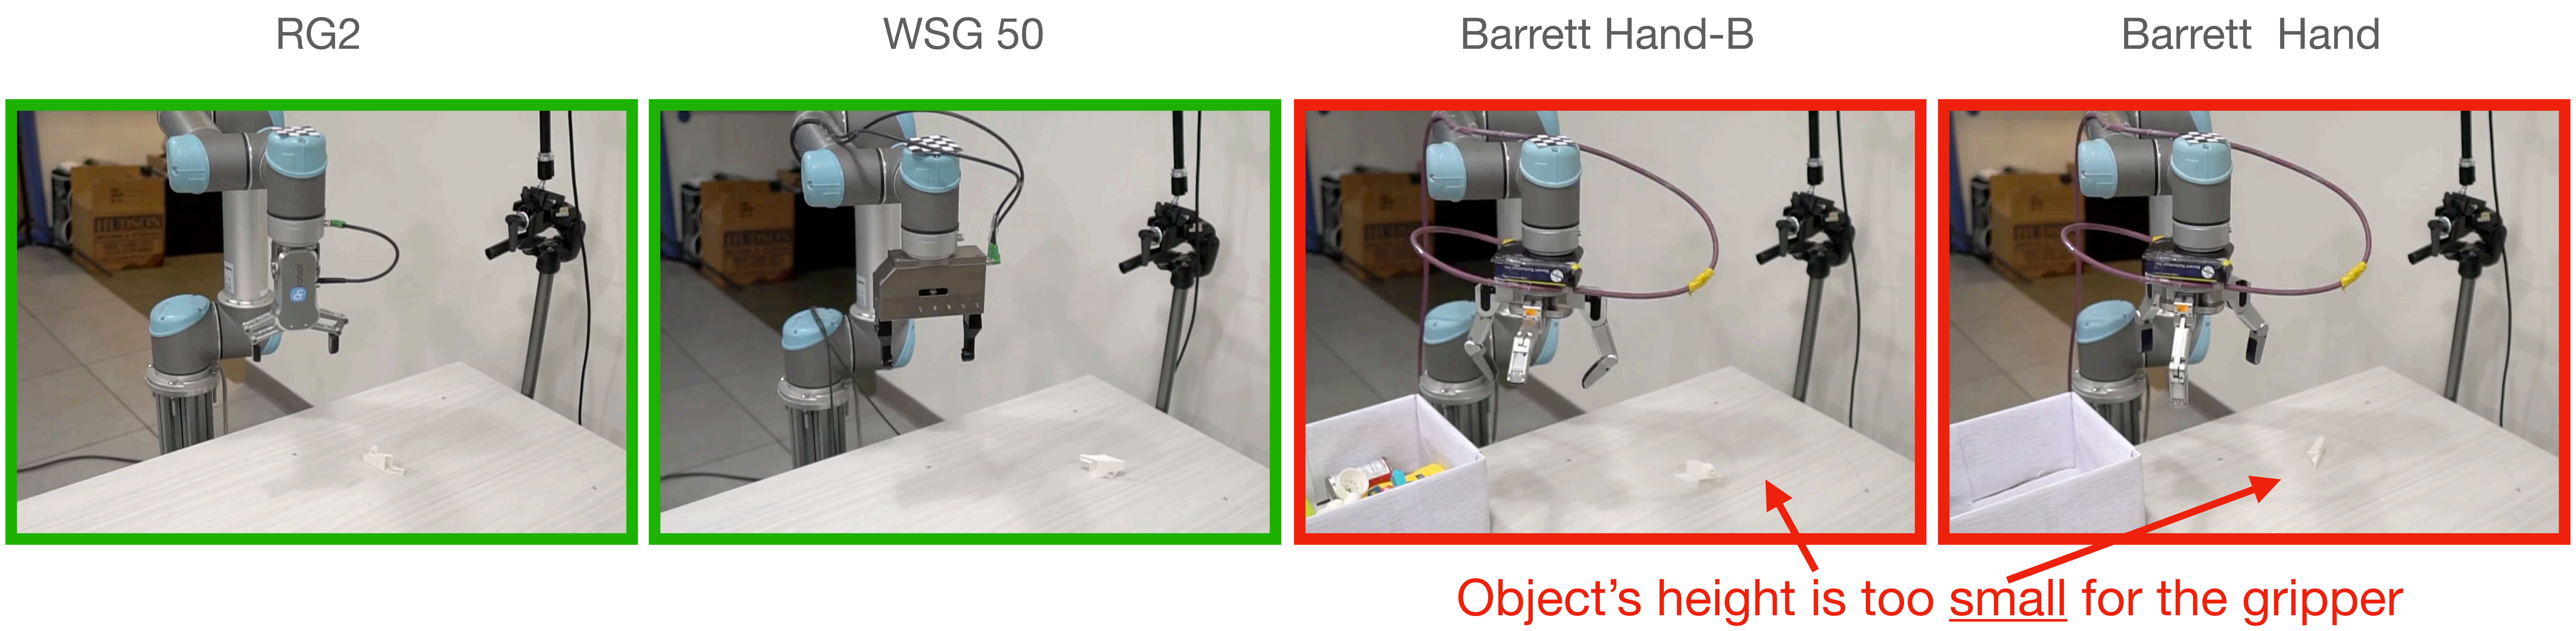

Real World Single Object

RG2 | WSG 50 | Barrett Hand-B | Barrett Hand

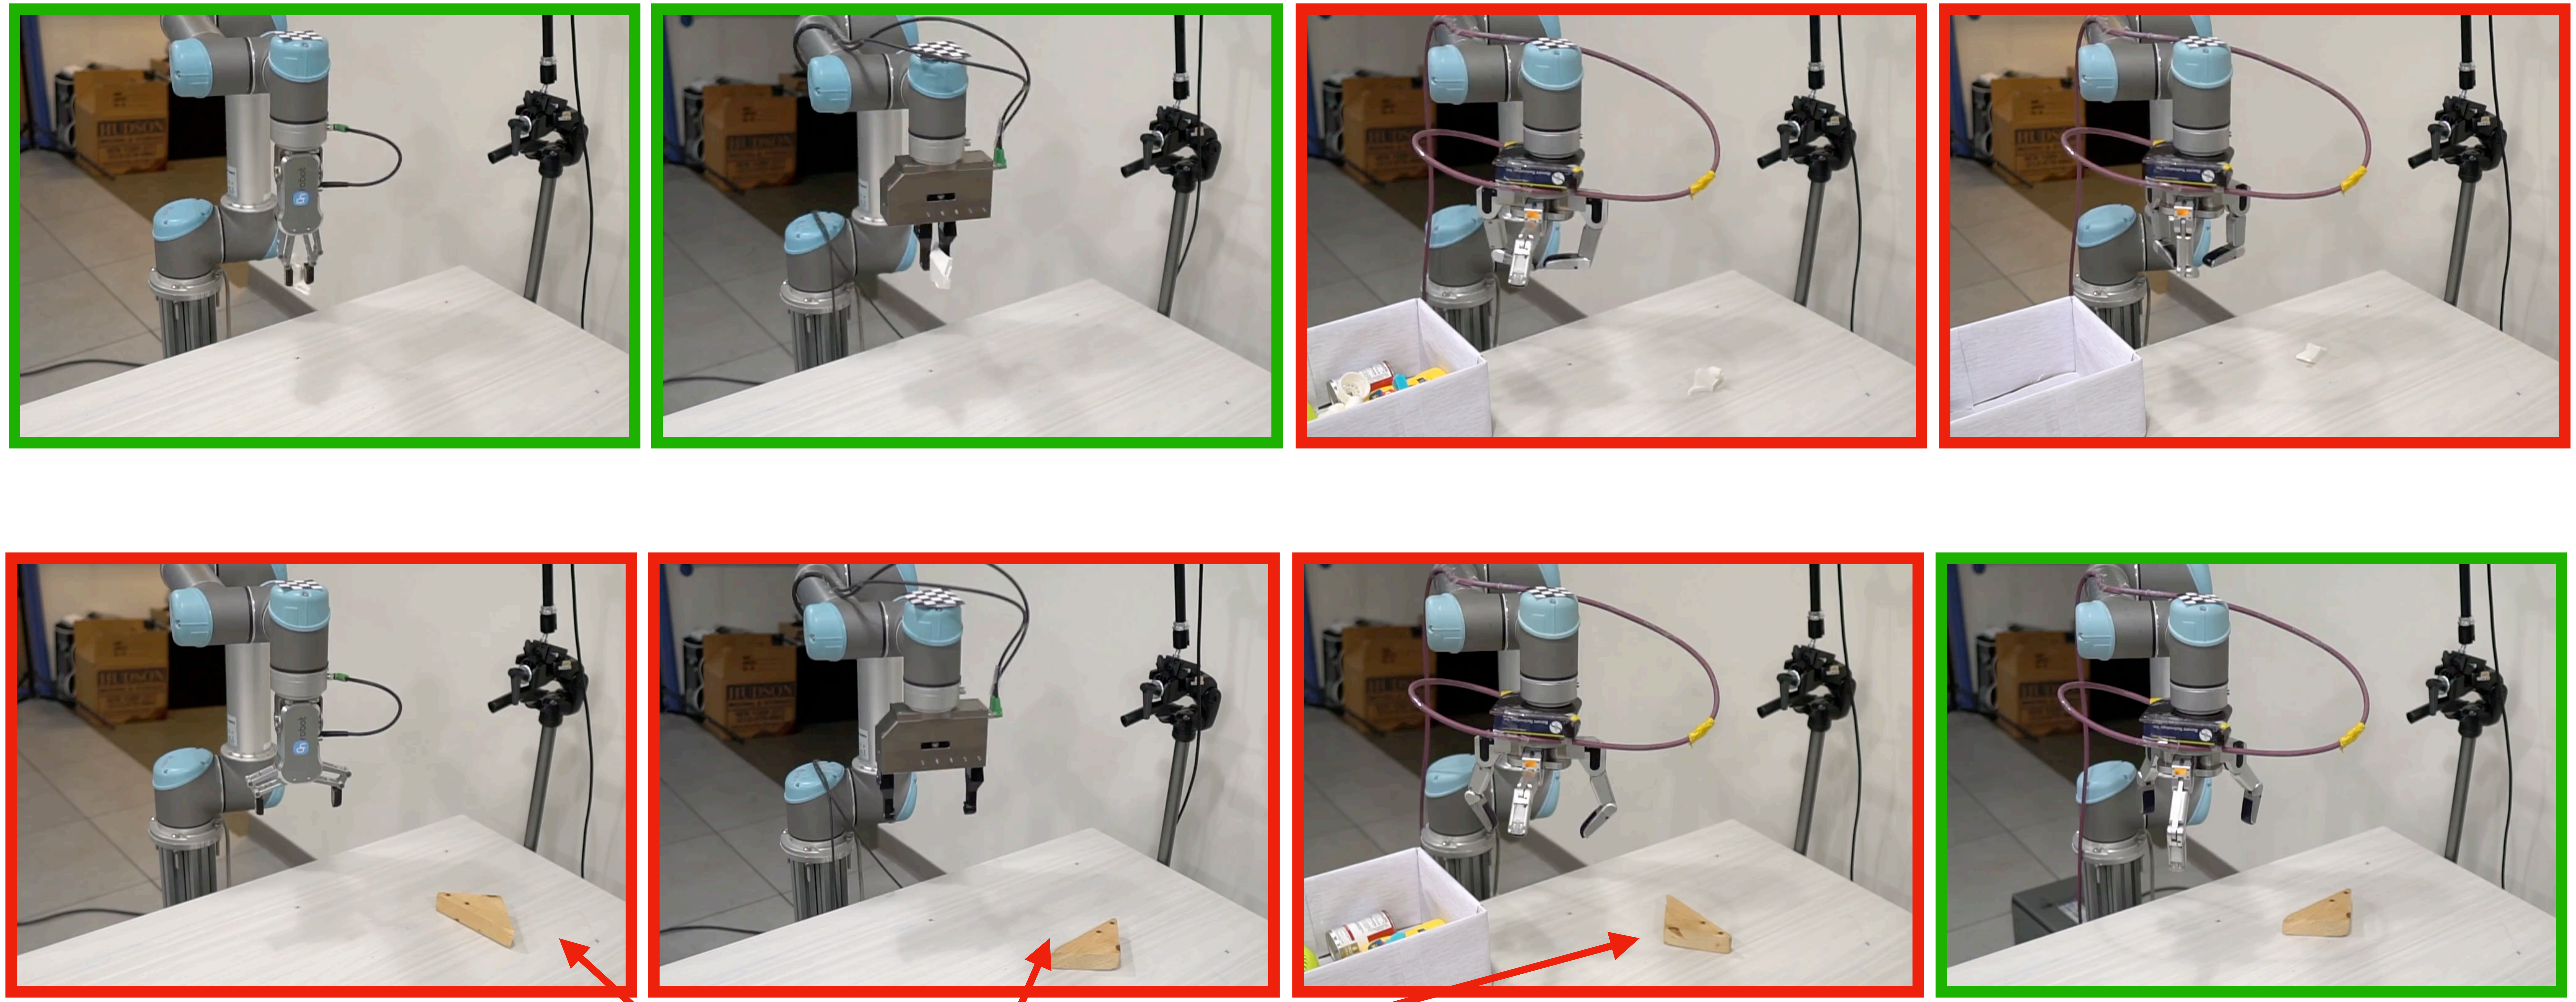

Triangle is hard to grasp for 2-finger grippers

5x speed

Real World Novel Object

RG2 | WSG 50 | Barrett Hand-B | Barrett Hand

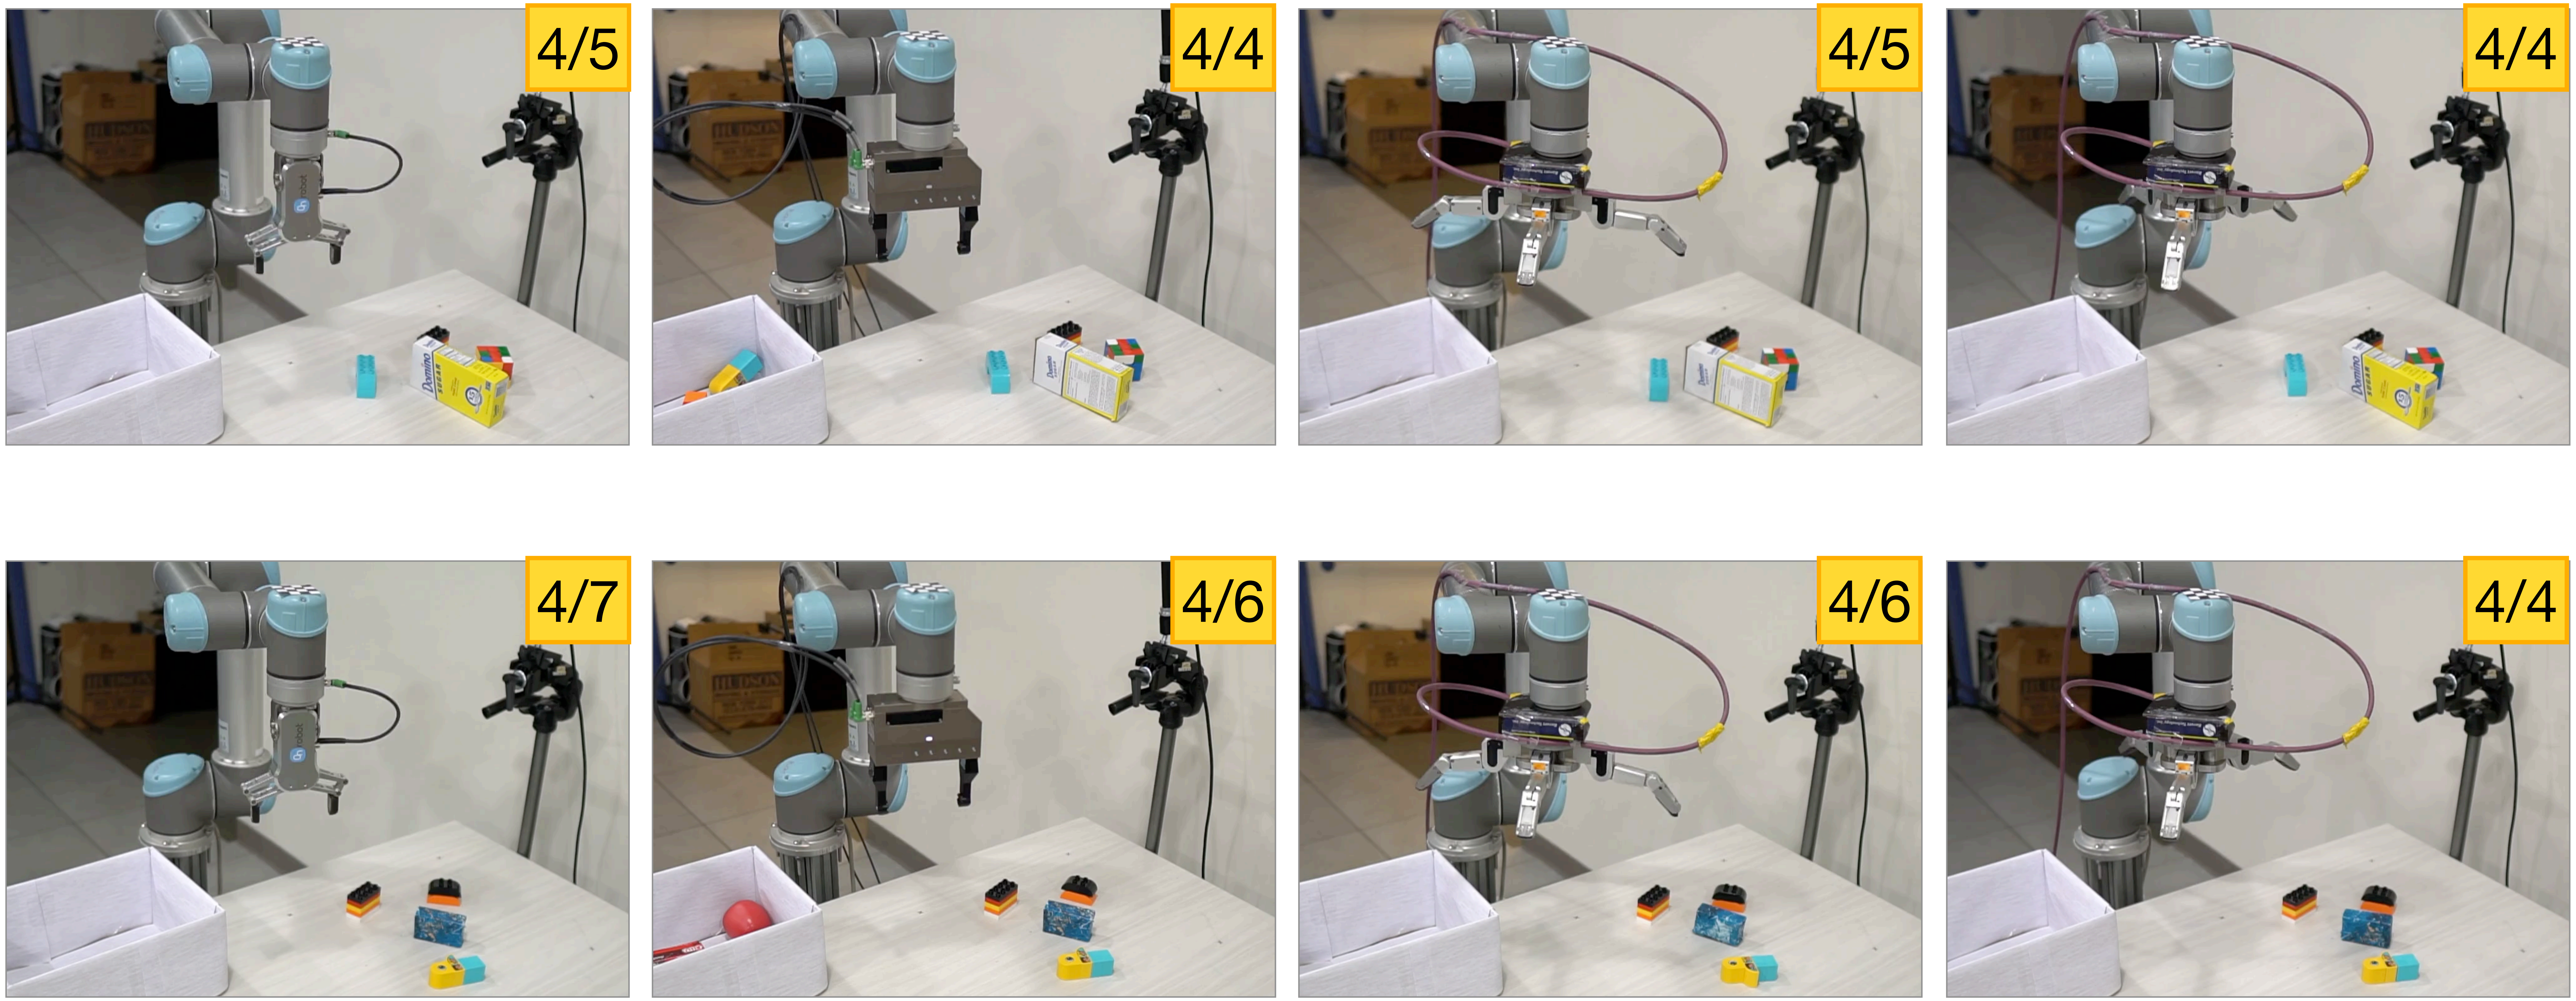

More details: https://adagrasp.cs.columbia.edu/

Self-Adaptive Manipulation

Treating a robot's hardware as an integral part of robot behavior —

the learned manipulation policies should be conditioned on their hardware

also inform how hardware should be improved.

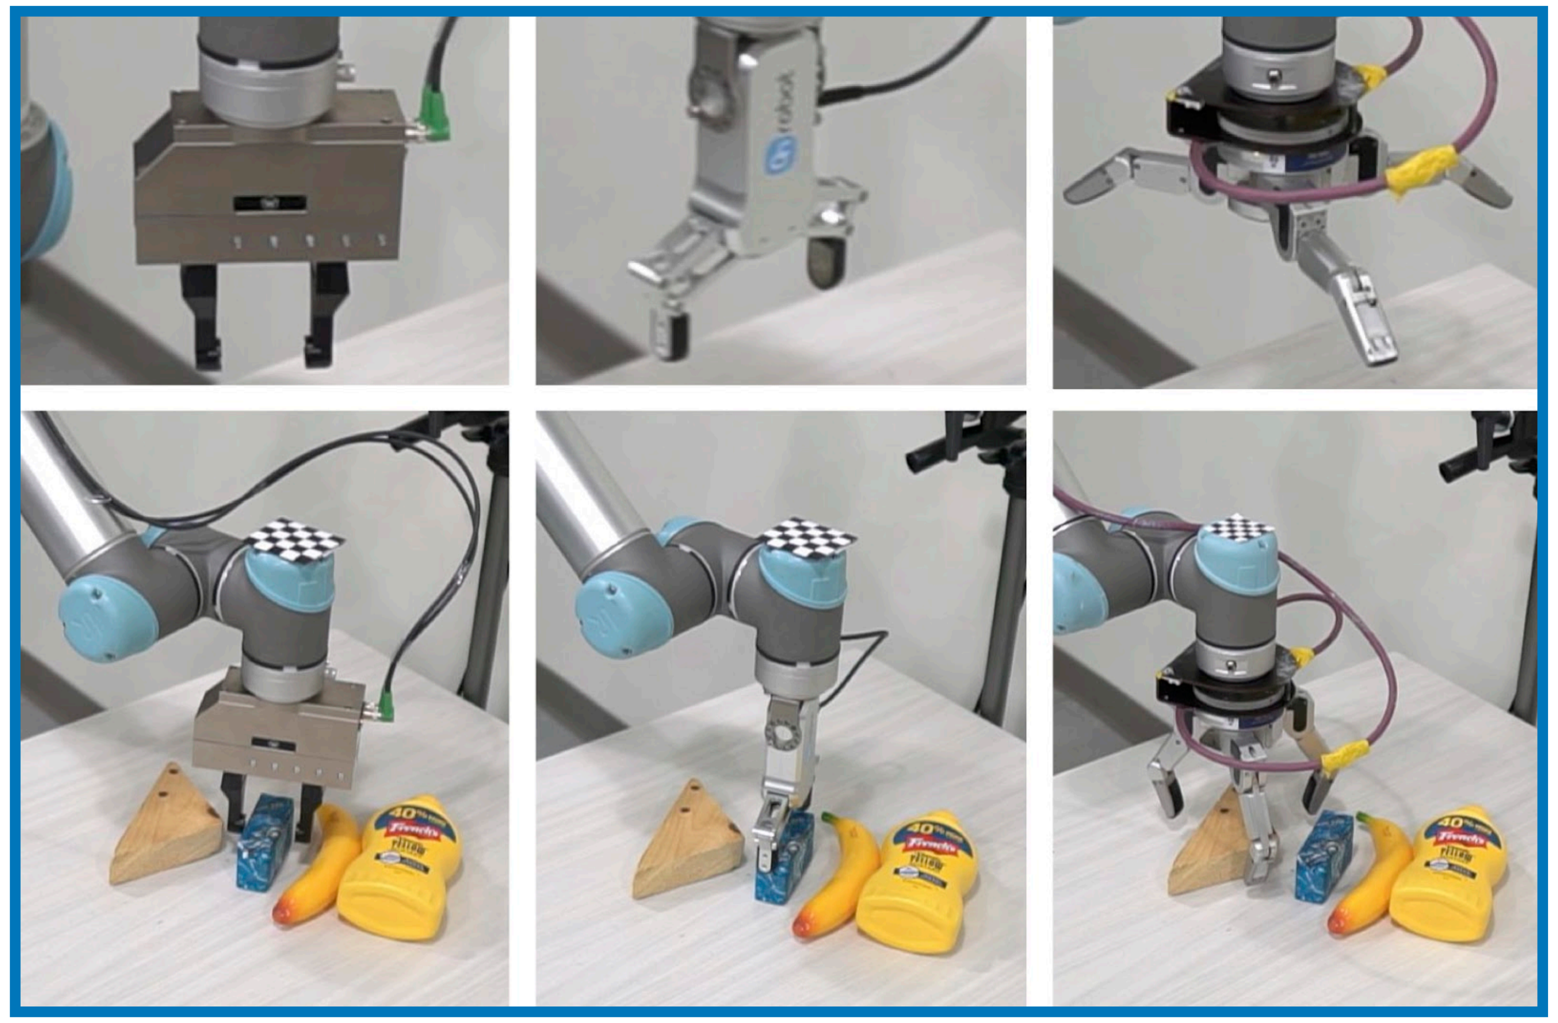

Adagrasp *(ICRA'21)*

AdaGrasp: Learning an Adaptive Gripper-Aware Grasping Policy
Zhenjia Xu, Beichun Qi, Shubham Agrawal, Shuran Song
ICRA 2021

Self-Adaptive Manipulation

Treating a robot's hardware as an integral part of robot behavior —

the learned manipulation policies should be conditioned on their hardware

also inform how hardware should be improved.

Fit2Form: 3D Generative Model for Robot Gripper Form Design
Conference on Robotic Learning (CoRL) 2020

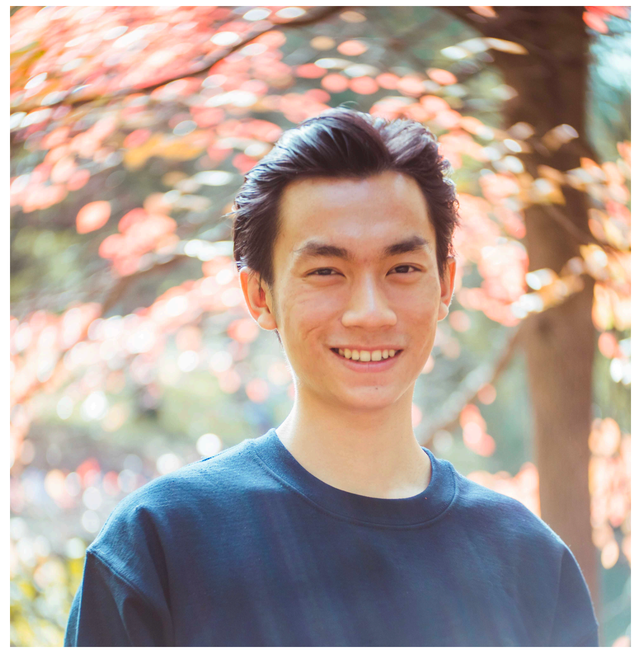

Huy Ha

Shubham Agrawal

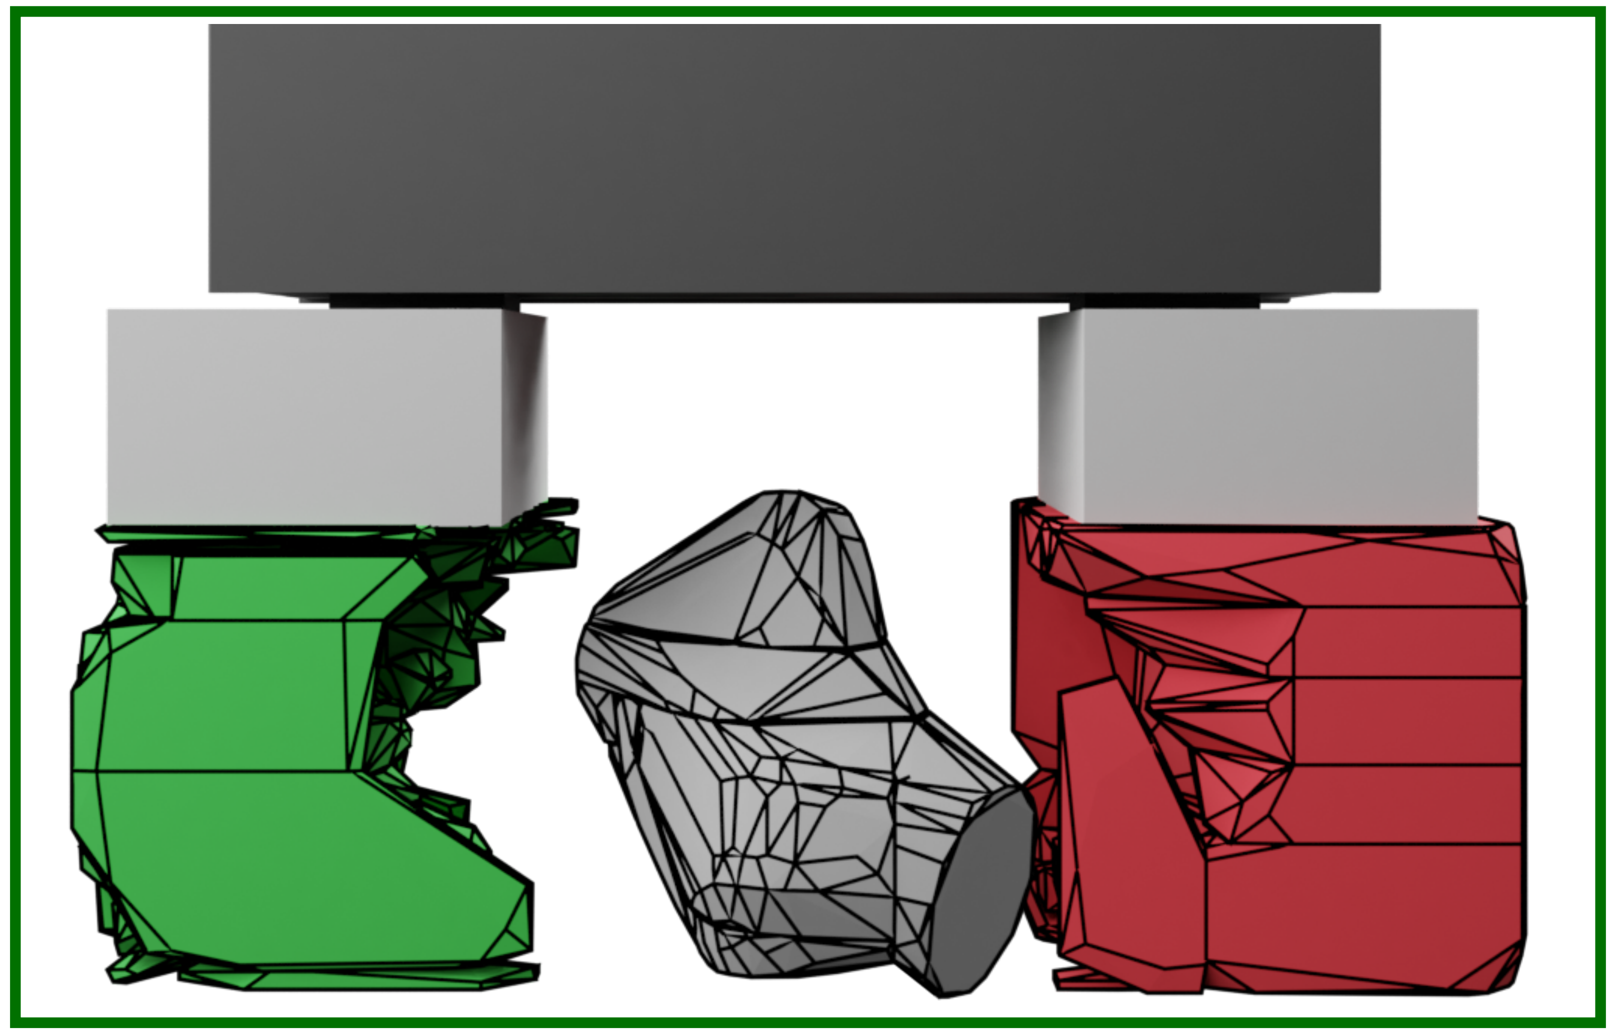

Fit2Form *(CoRL'20)*

Motivation: Specialized tools in Natural

Motivation

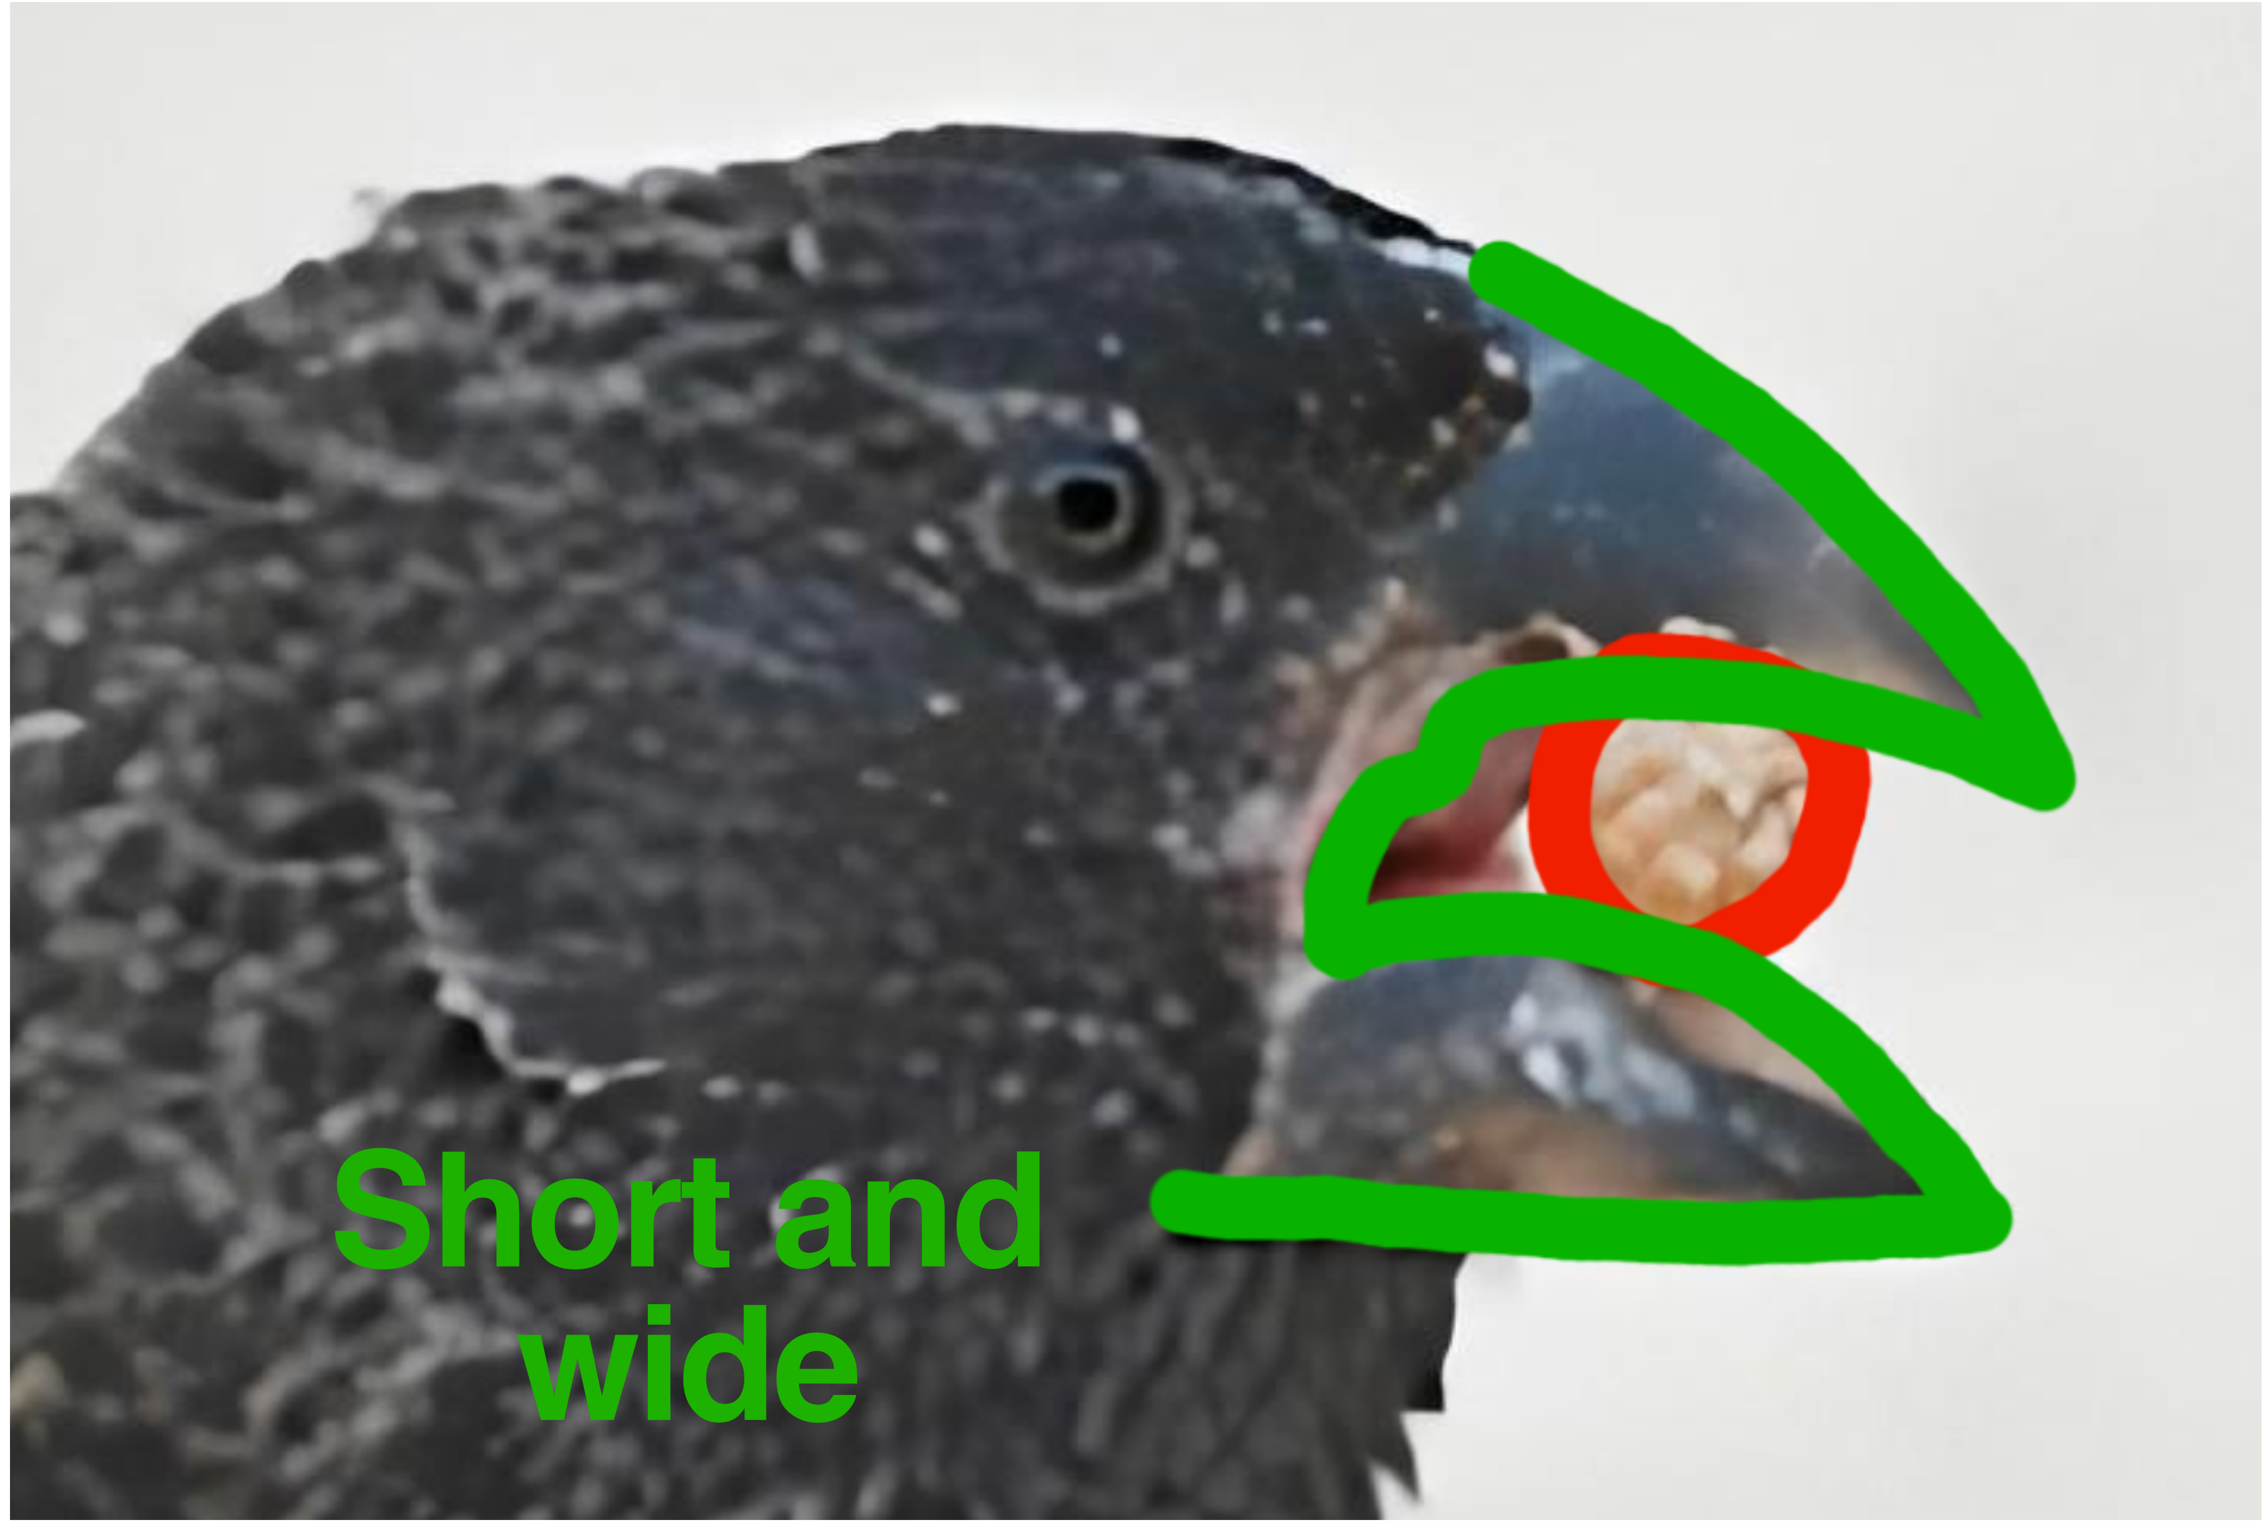

You need the right tool for the right job

Photograph by Author Morris (left) and Peter R. Grant (right)

Task specific gripper

General purpose gripper

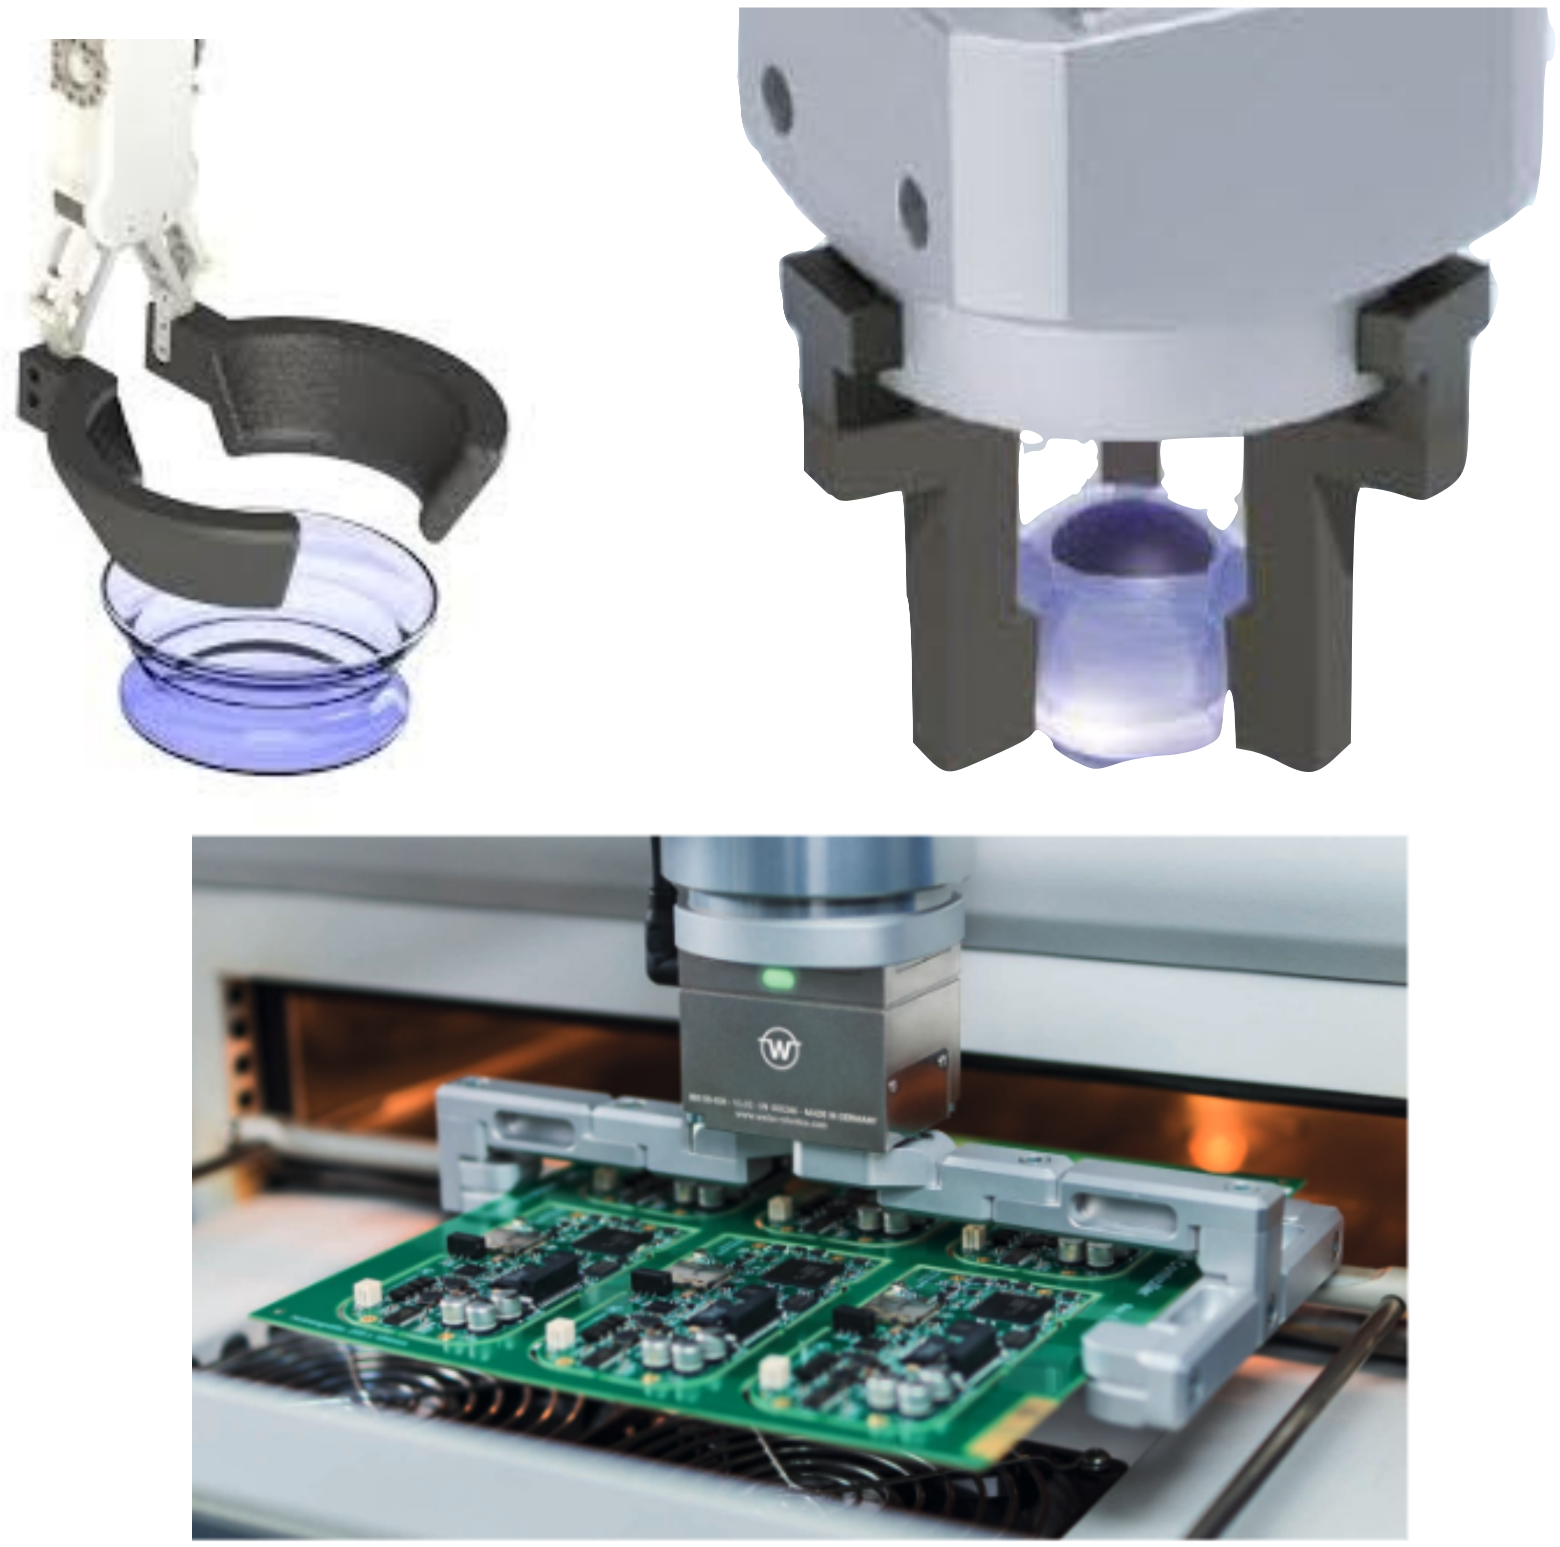

Task specific gripper

✓ Improves performance and robustness

× Manual design process is costly and time-consuming

<u>Can we use machine learning algorithms to automate this design process?</u>

Gripper Finger Form Design

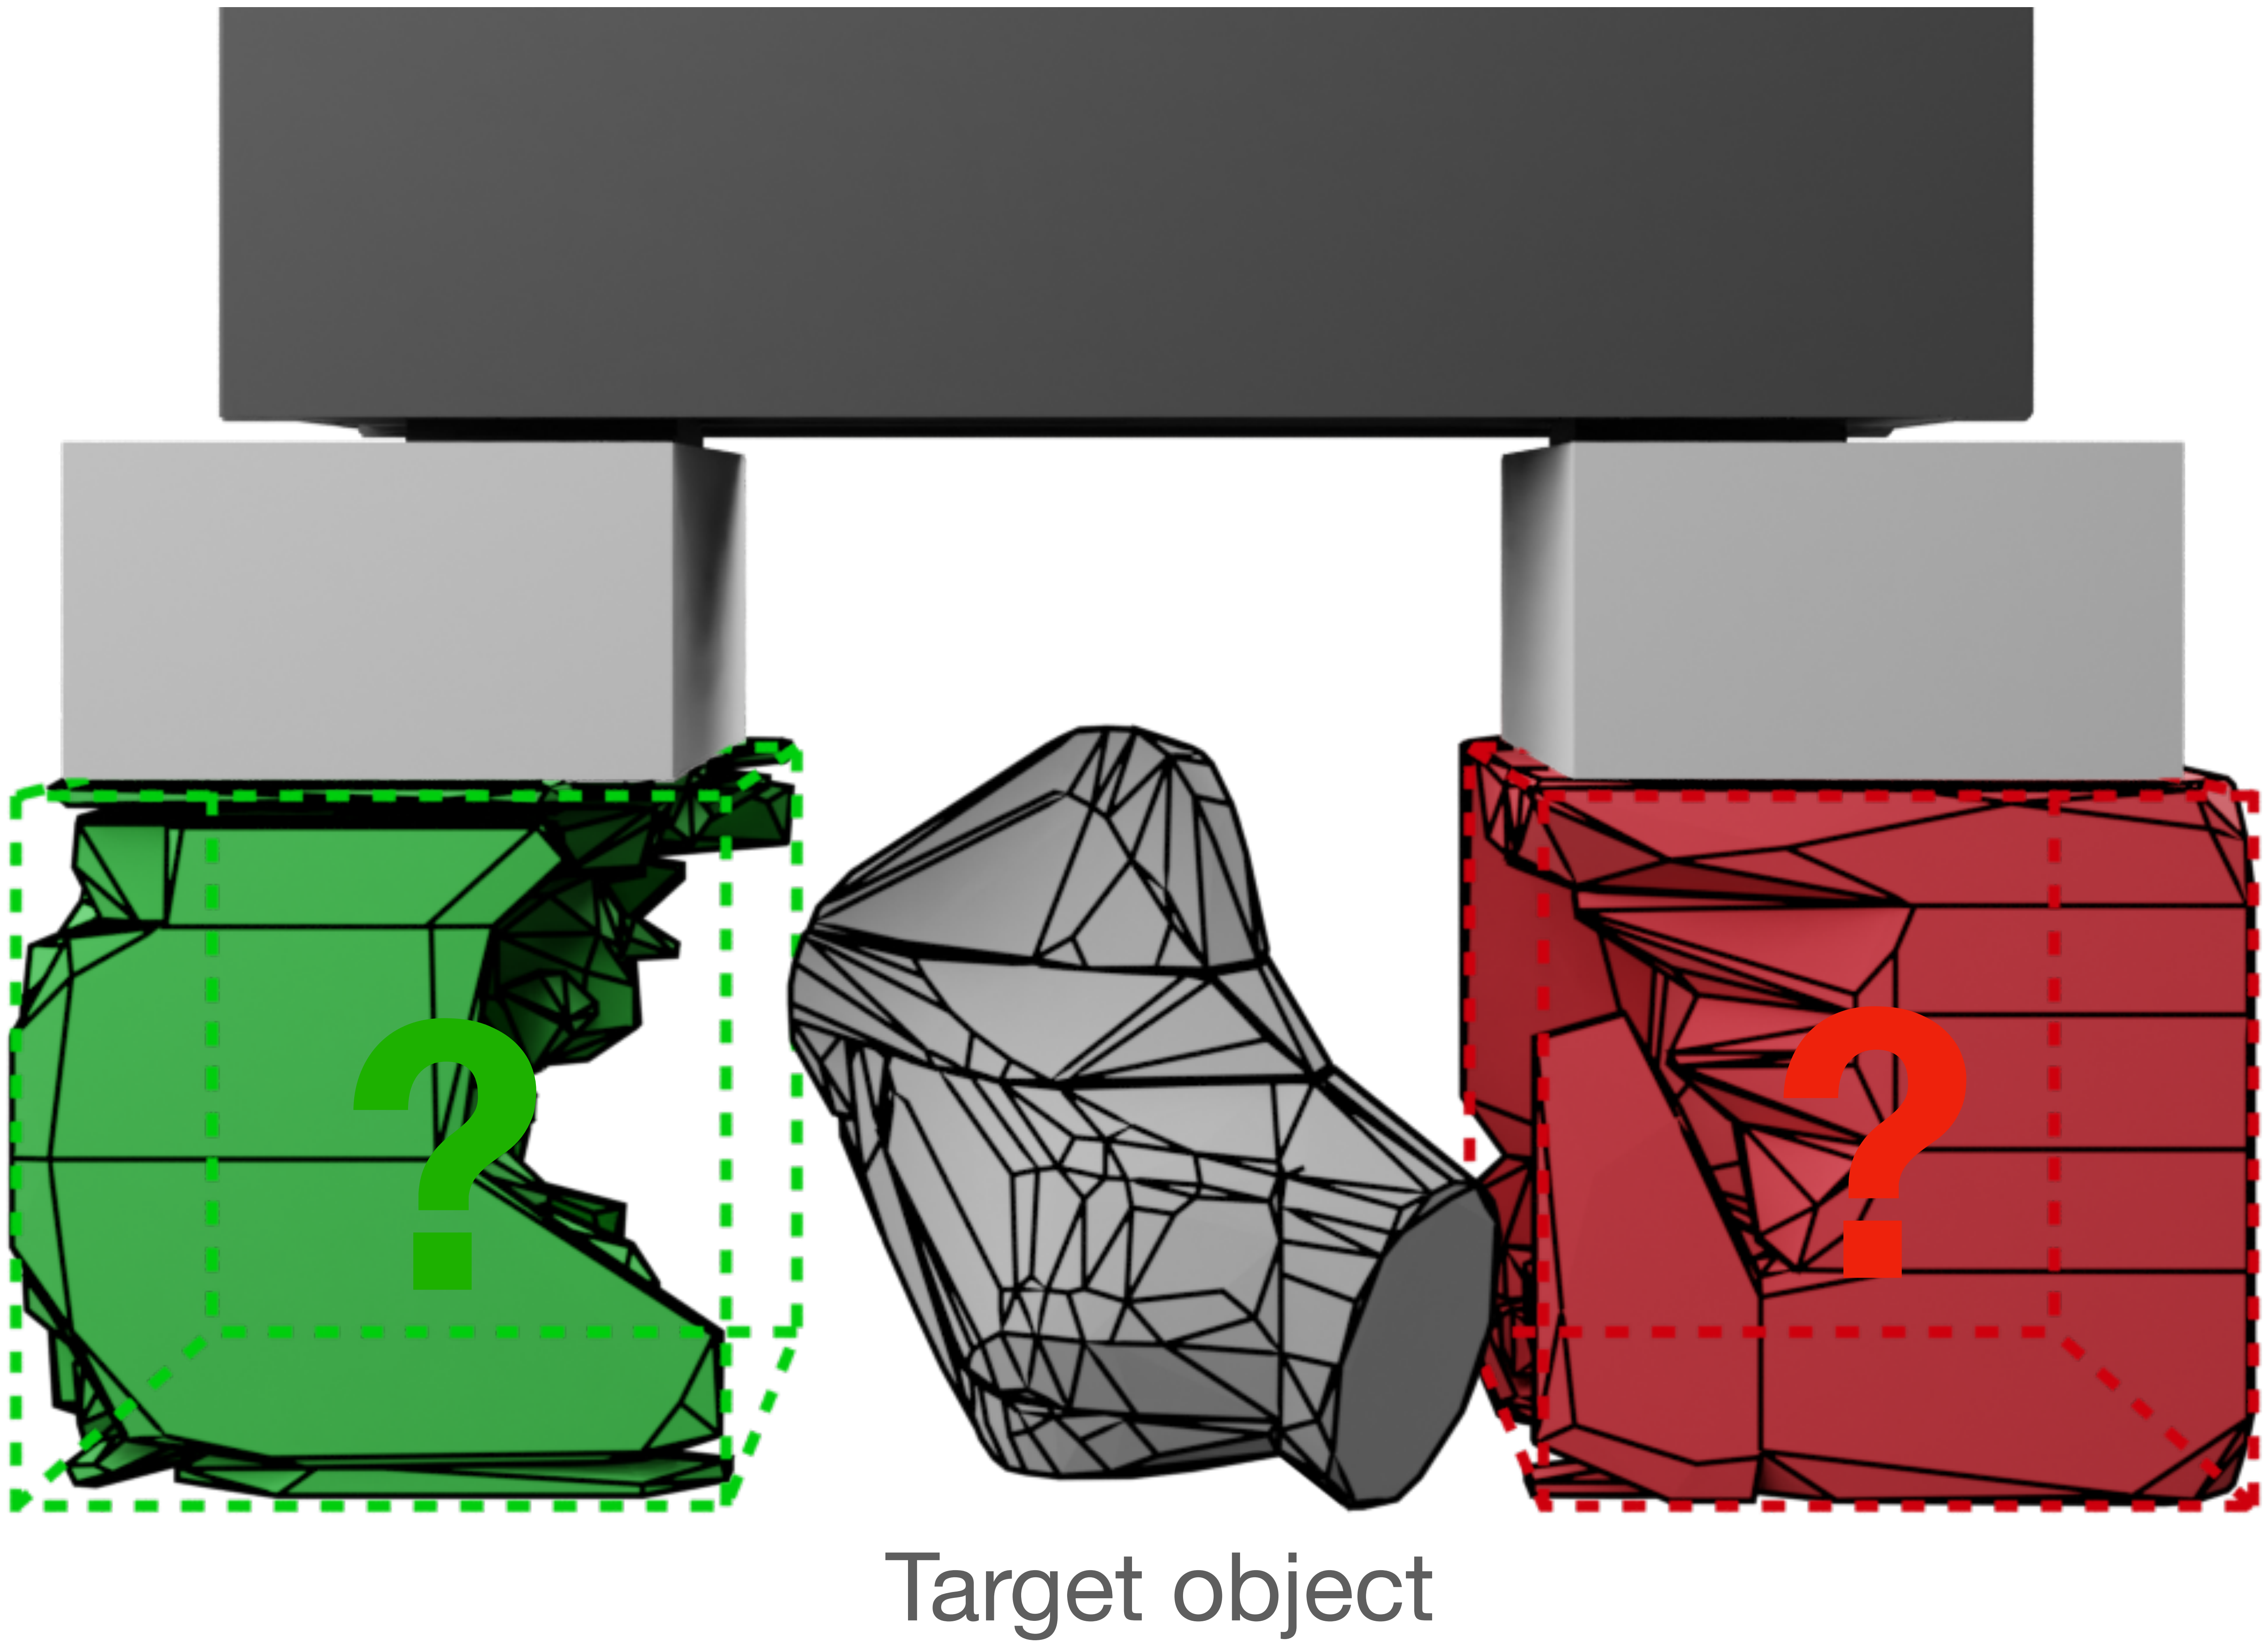

Design Objectives

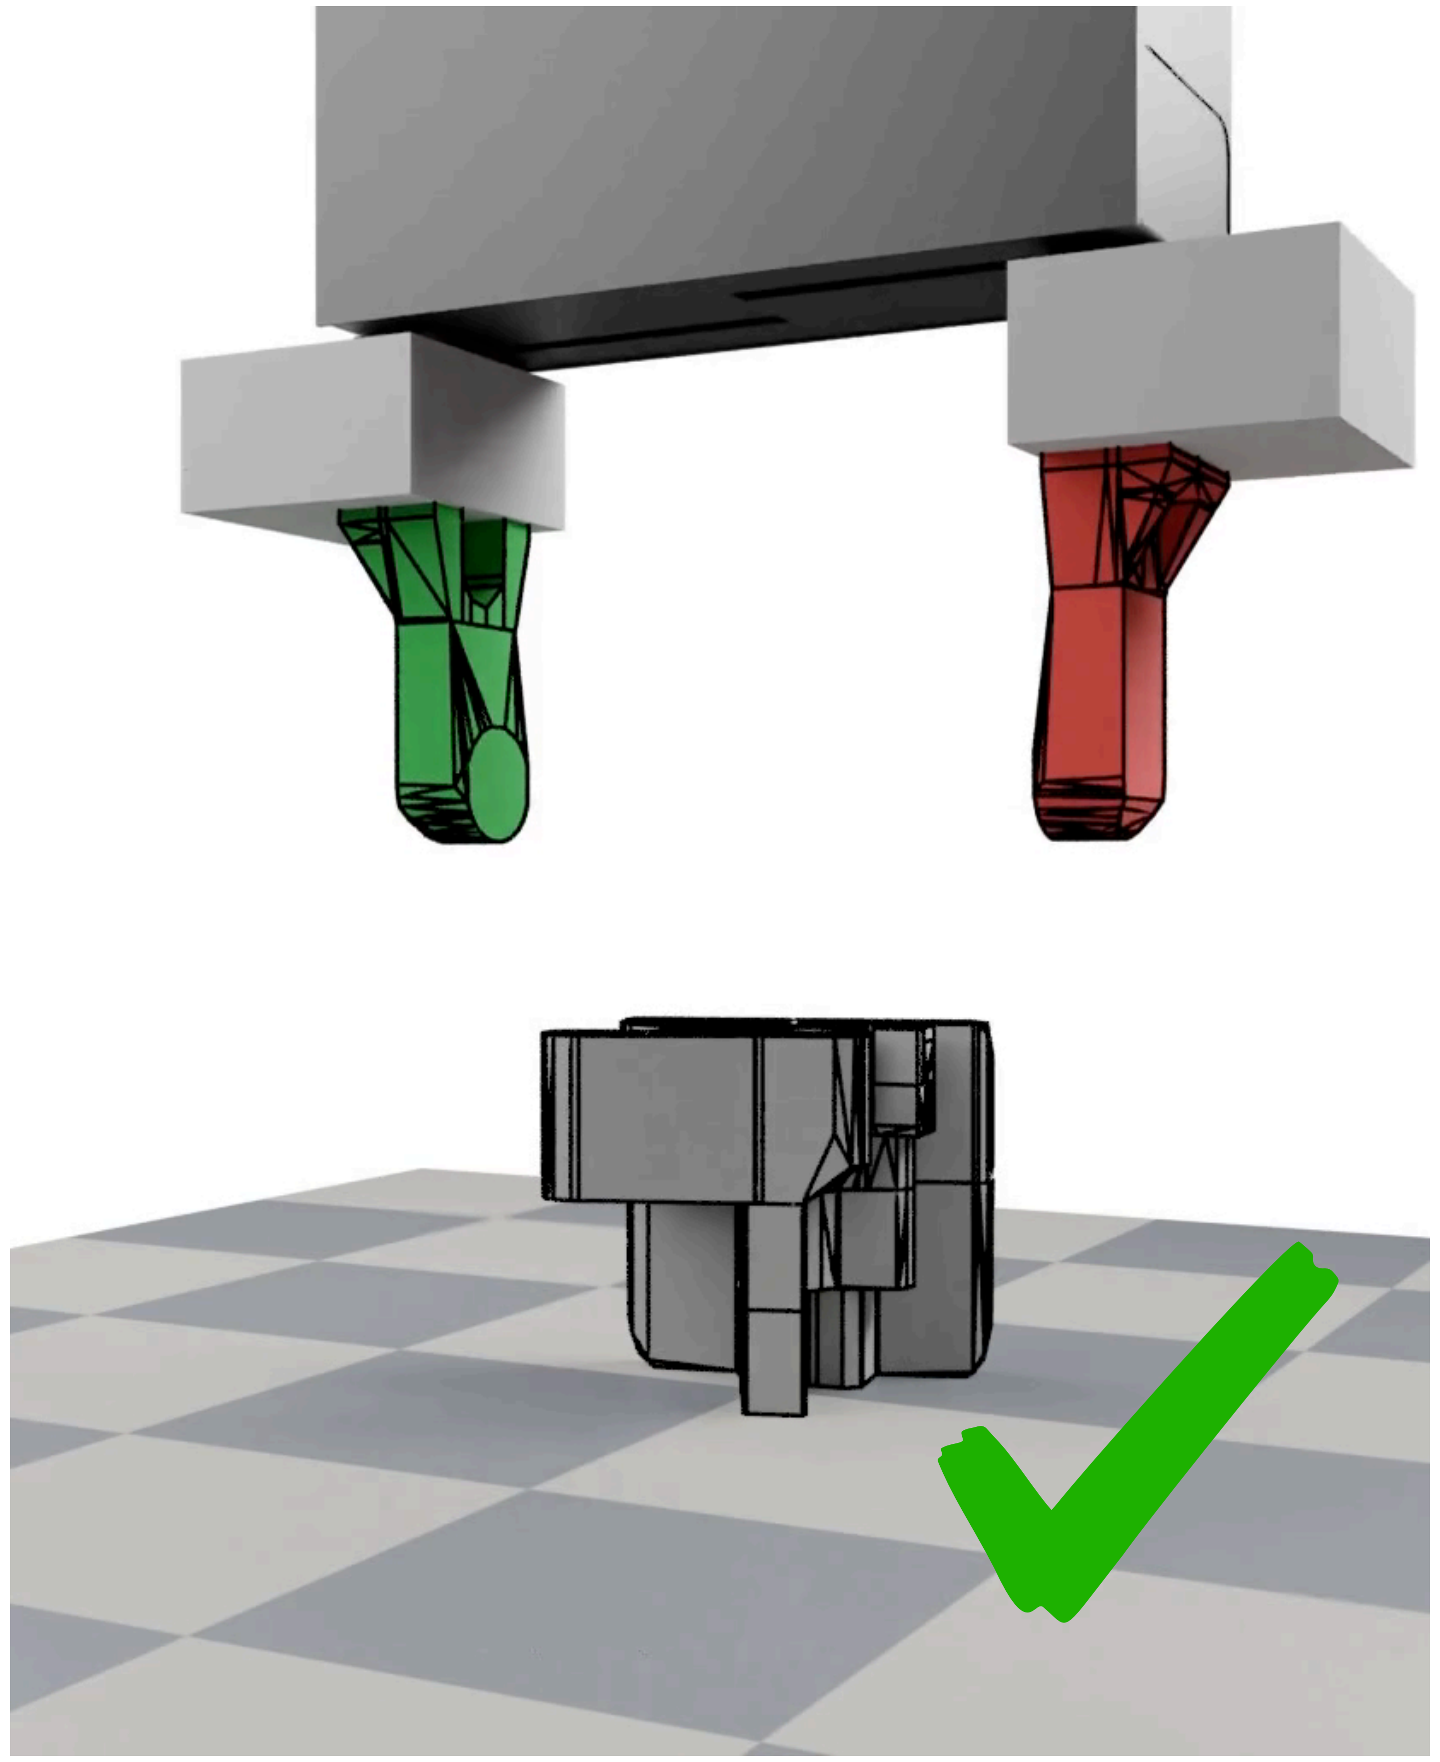

Grasp Success

Stability

Robustness

Feasibility

Infeasible Fingers

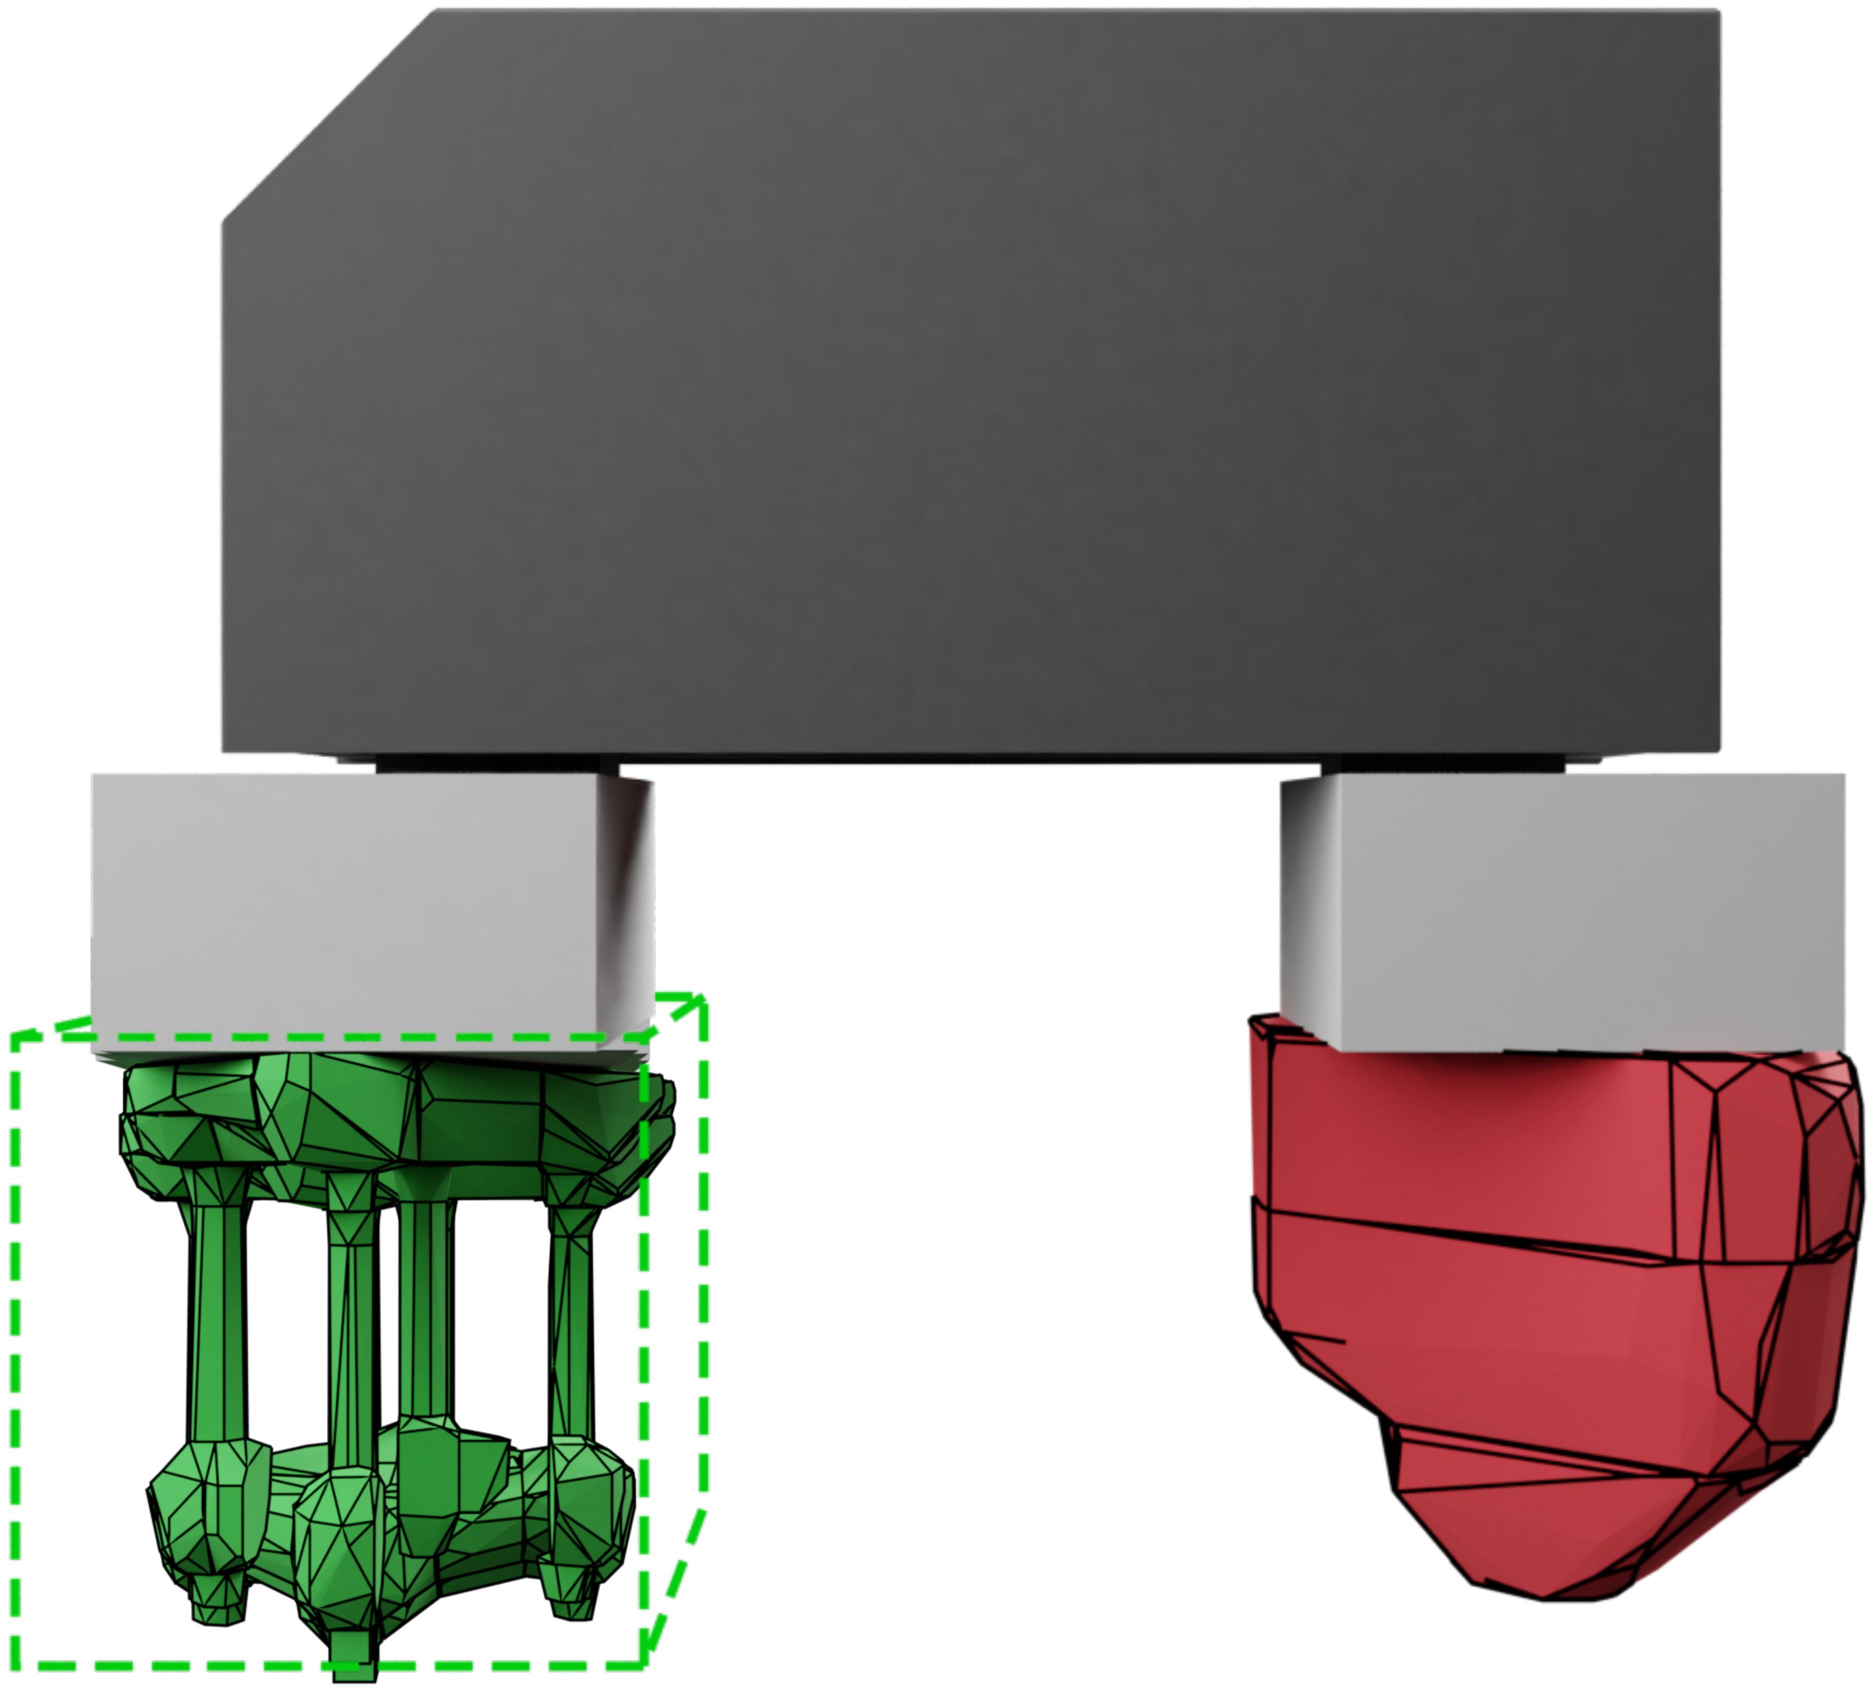

- ✗ Base Disconnected
- ✗ Multiple Connected Component
- ✓ Single Connected Component
- ✓ Base Connected

Approach

Navigate large design space

(all possible shapes in 30^3 Volume)

Represent high-level design objectives

(in a differentiable manner to enable efficient learning)

Approach

Navigate large design space

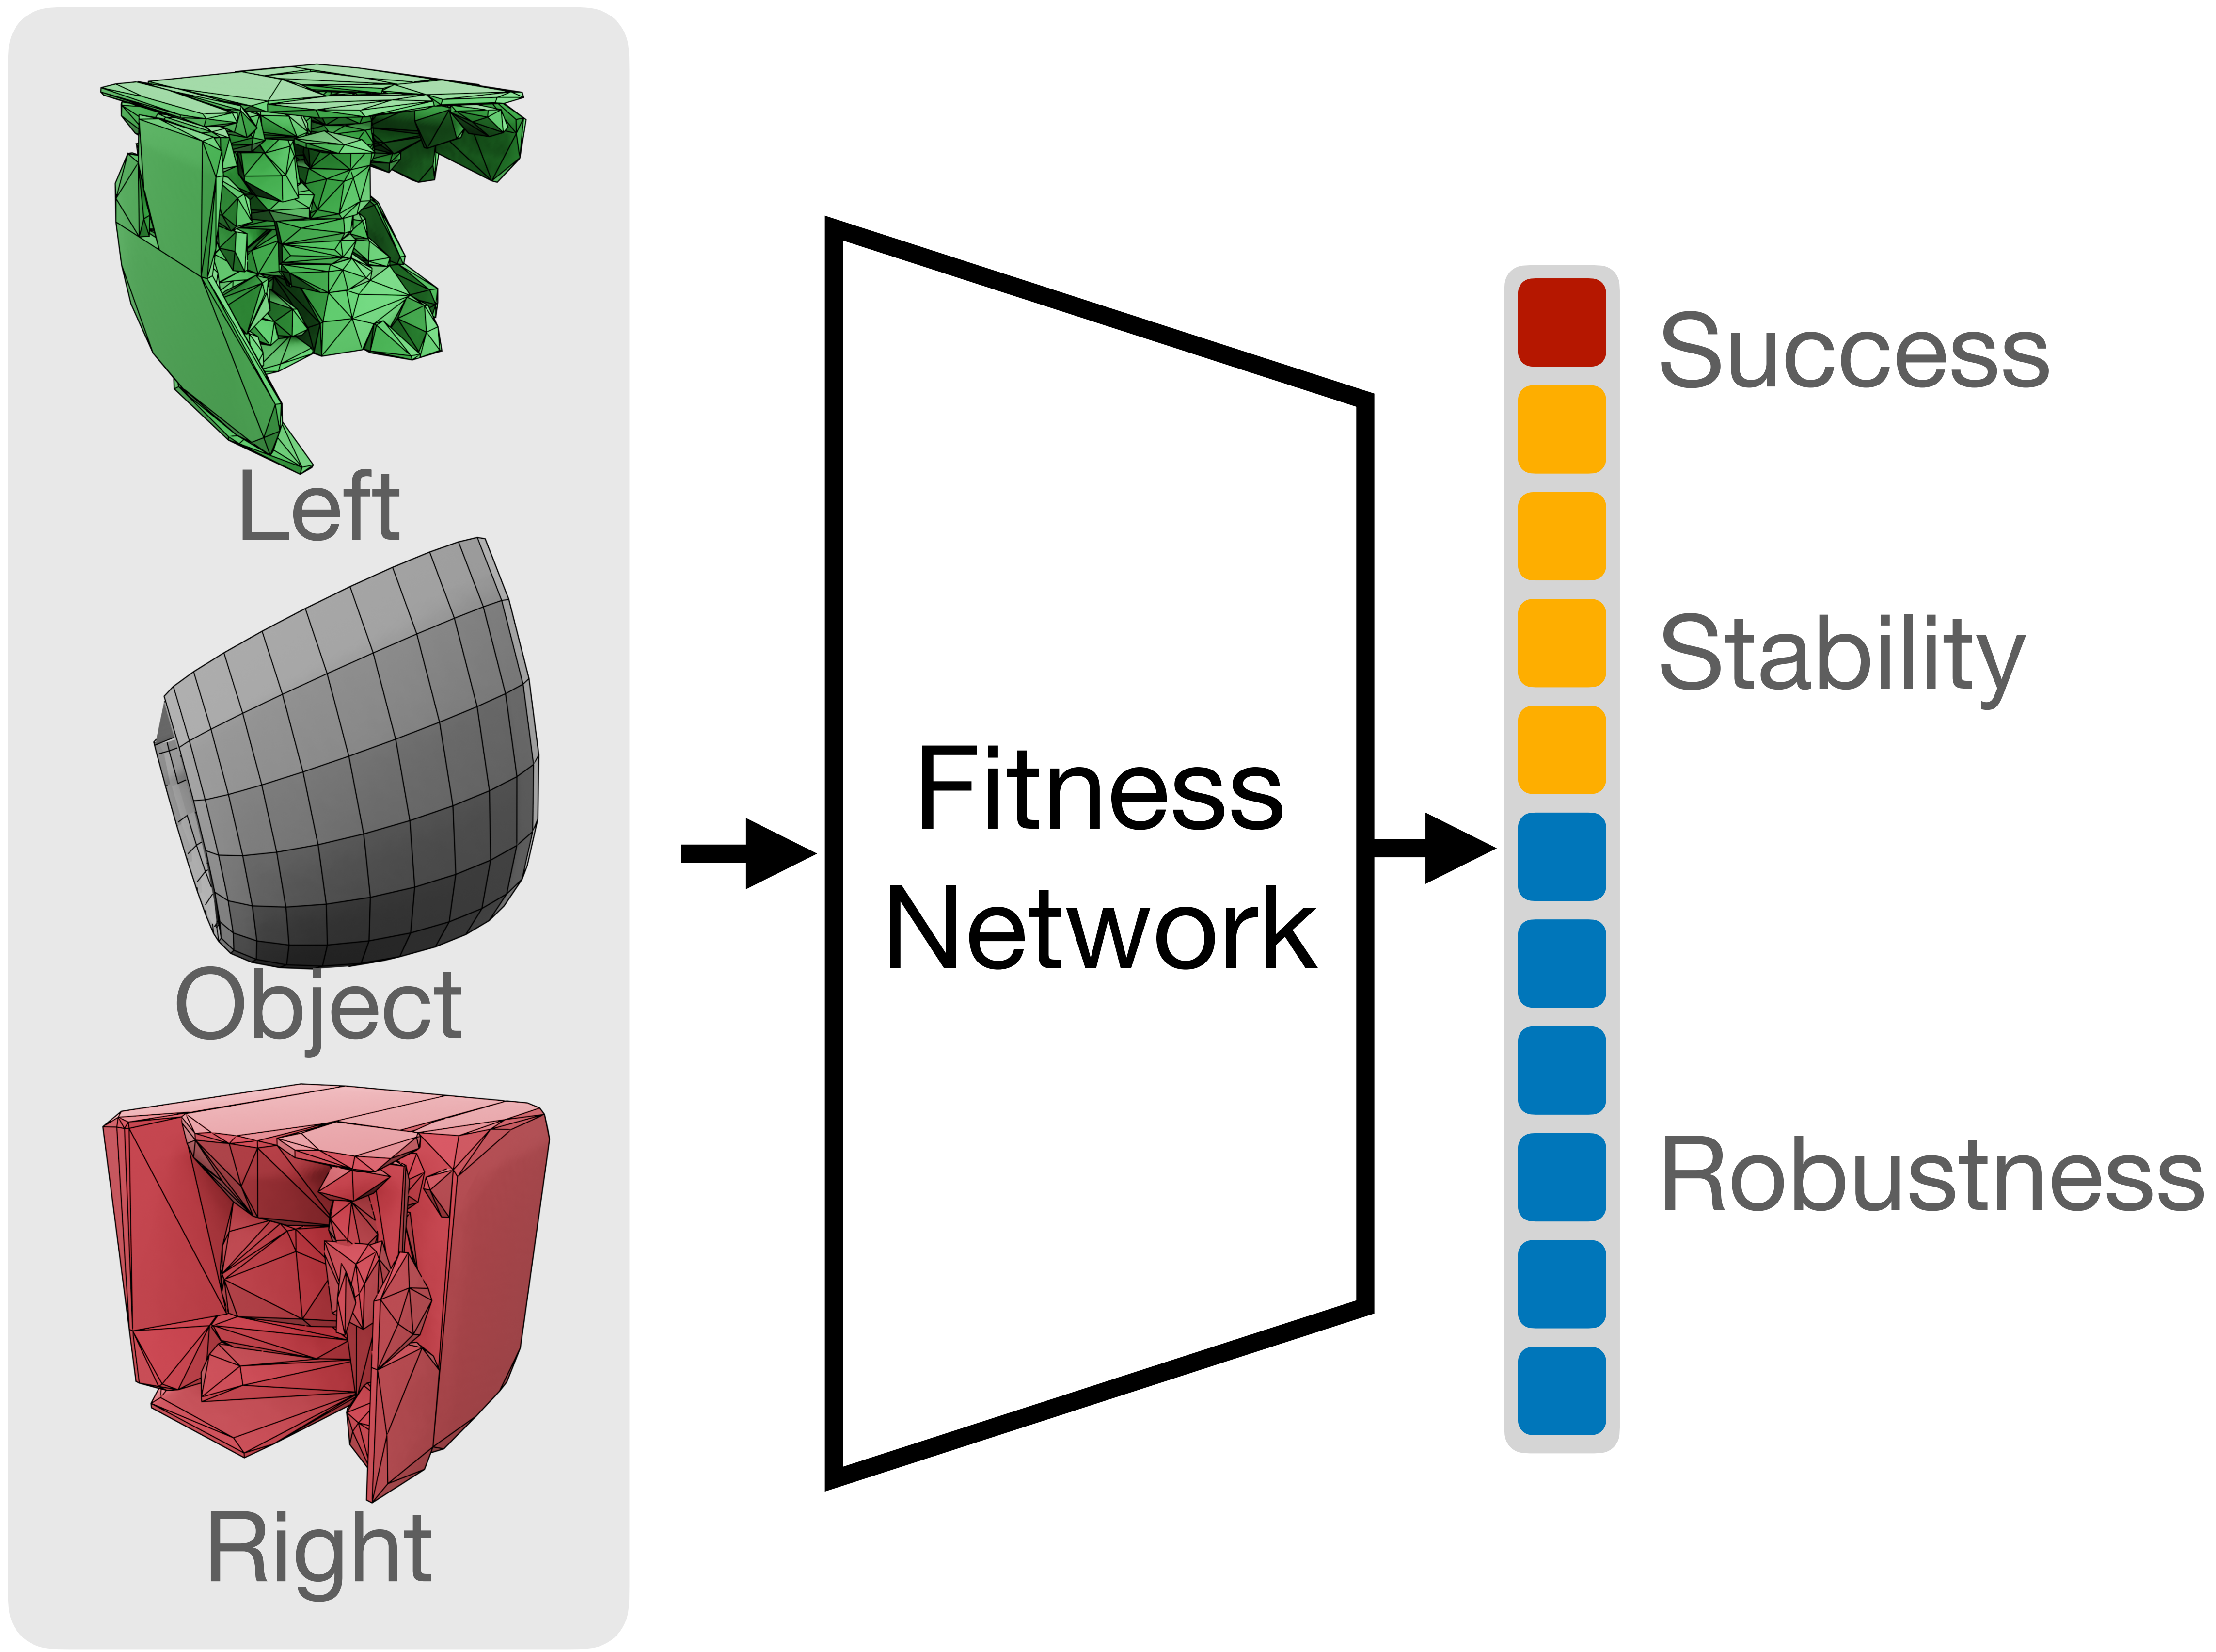

Represent high-level design objectives

Approach

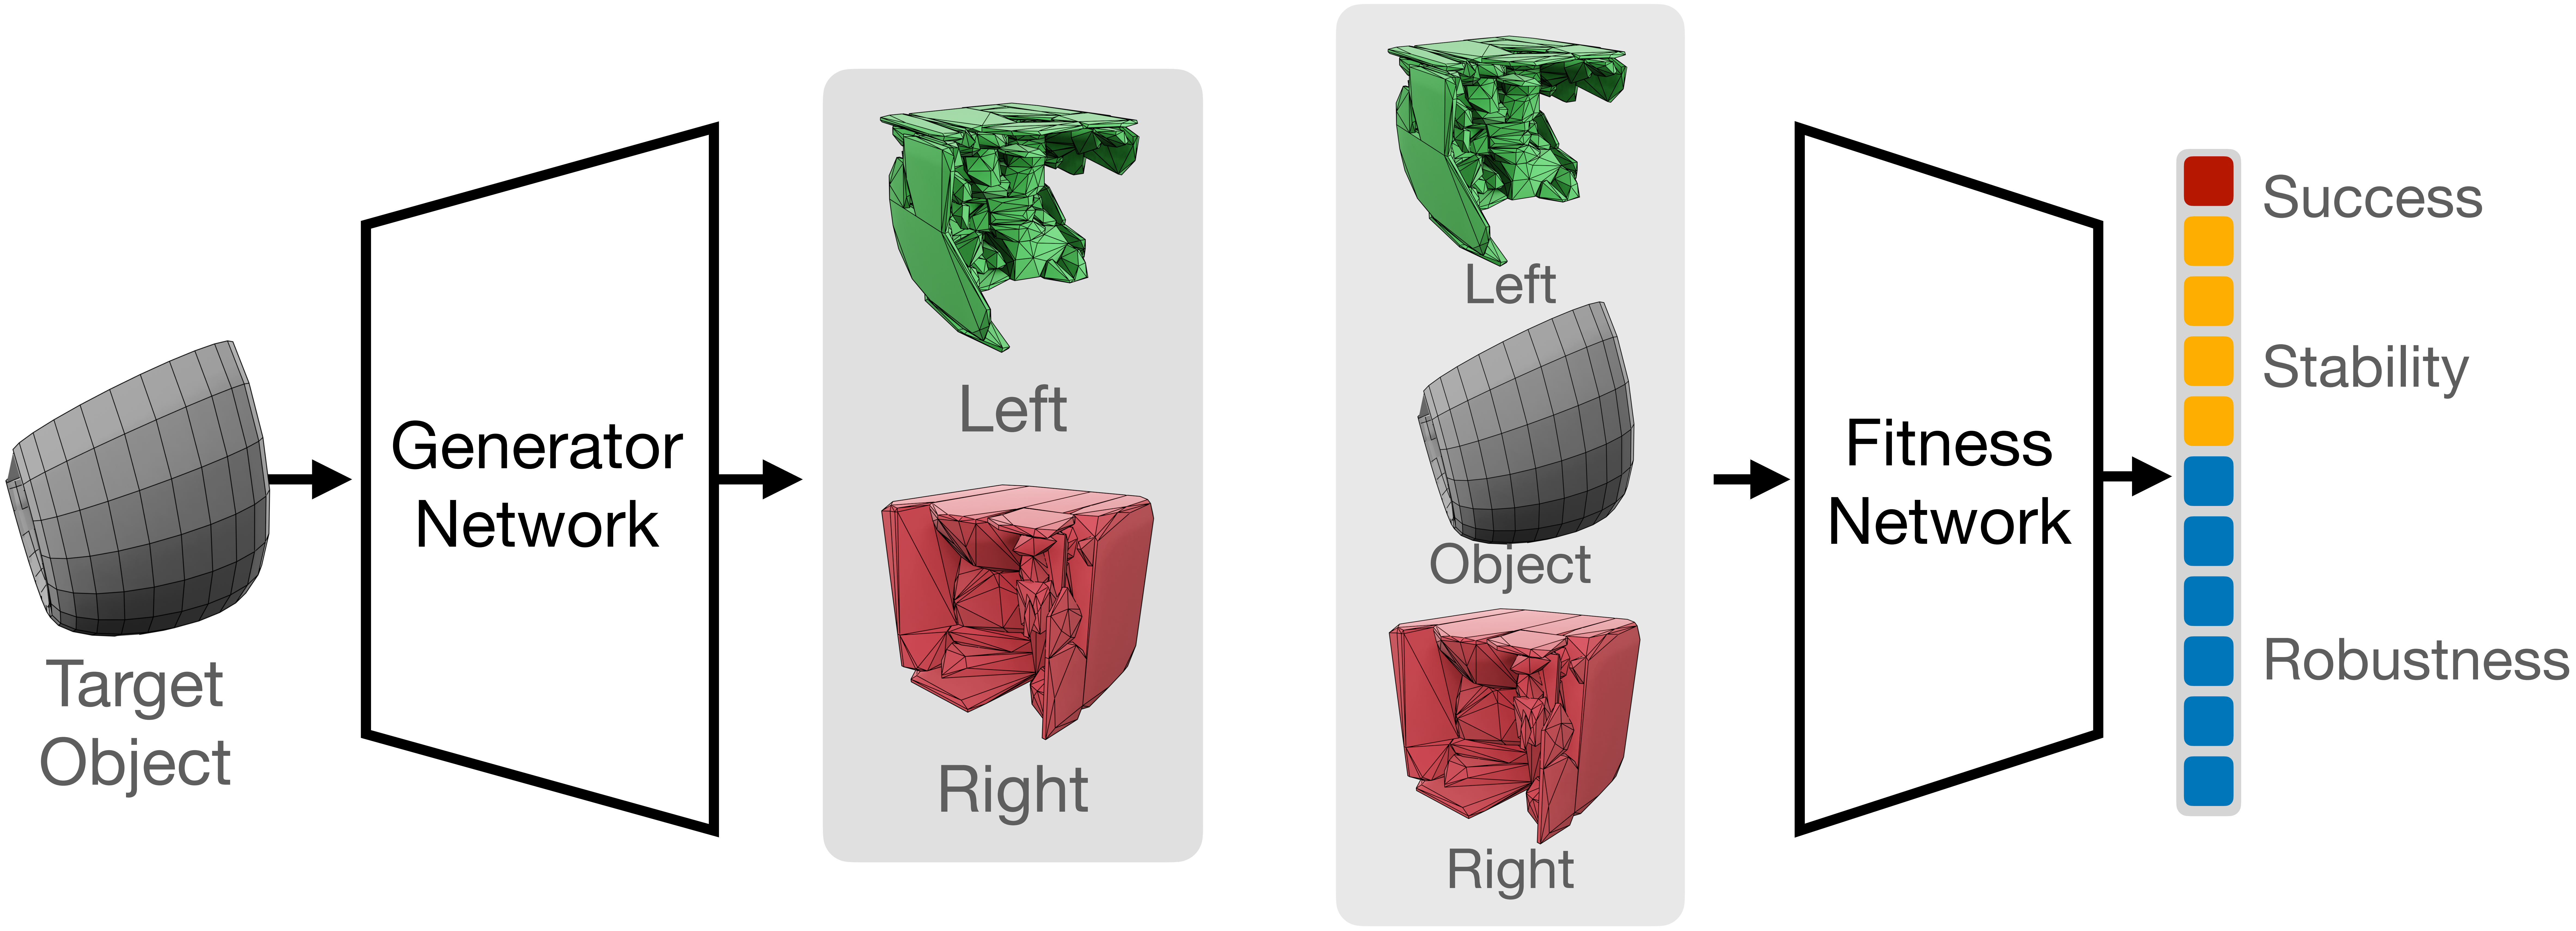

Navigate large design space

Represent high-level design objectives

Approach

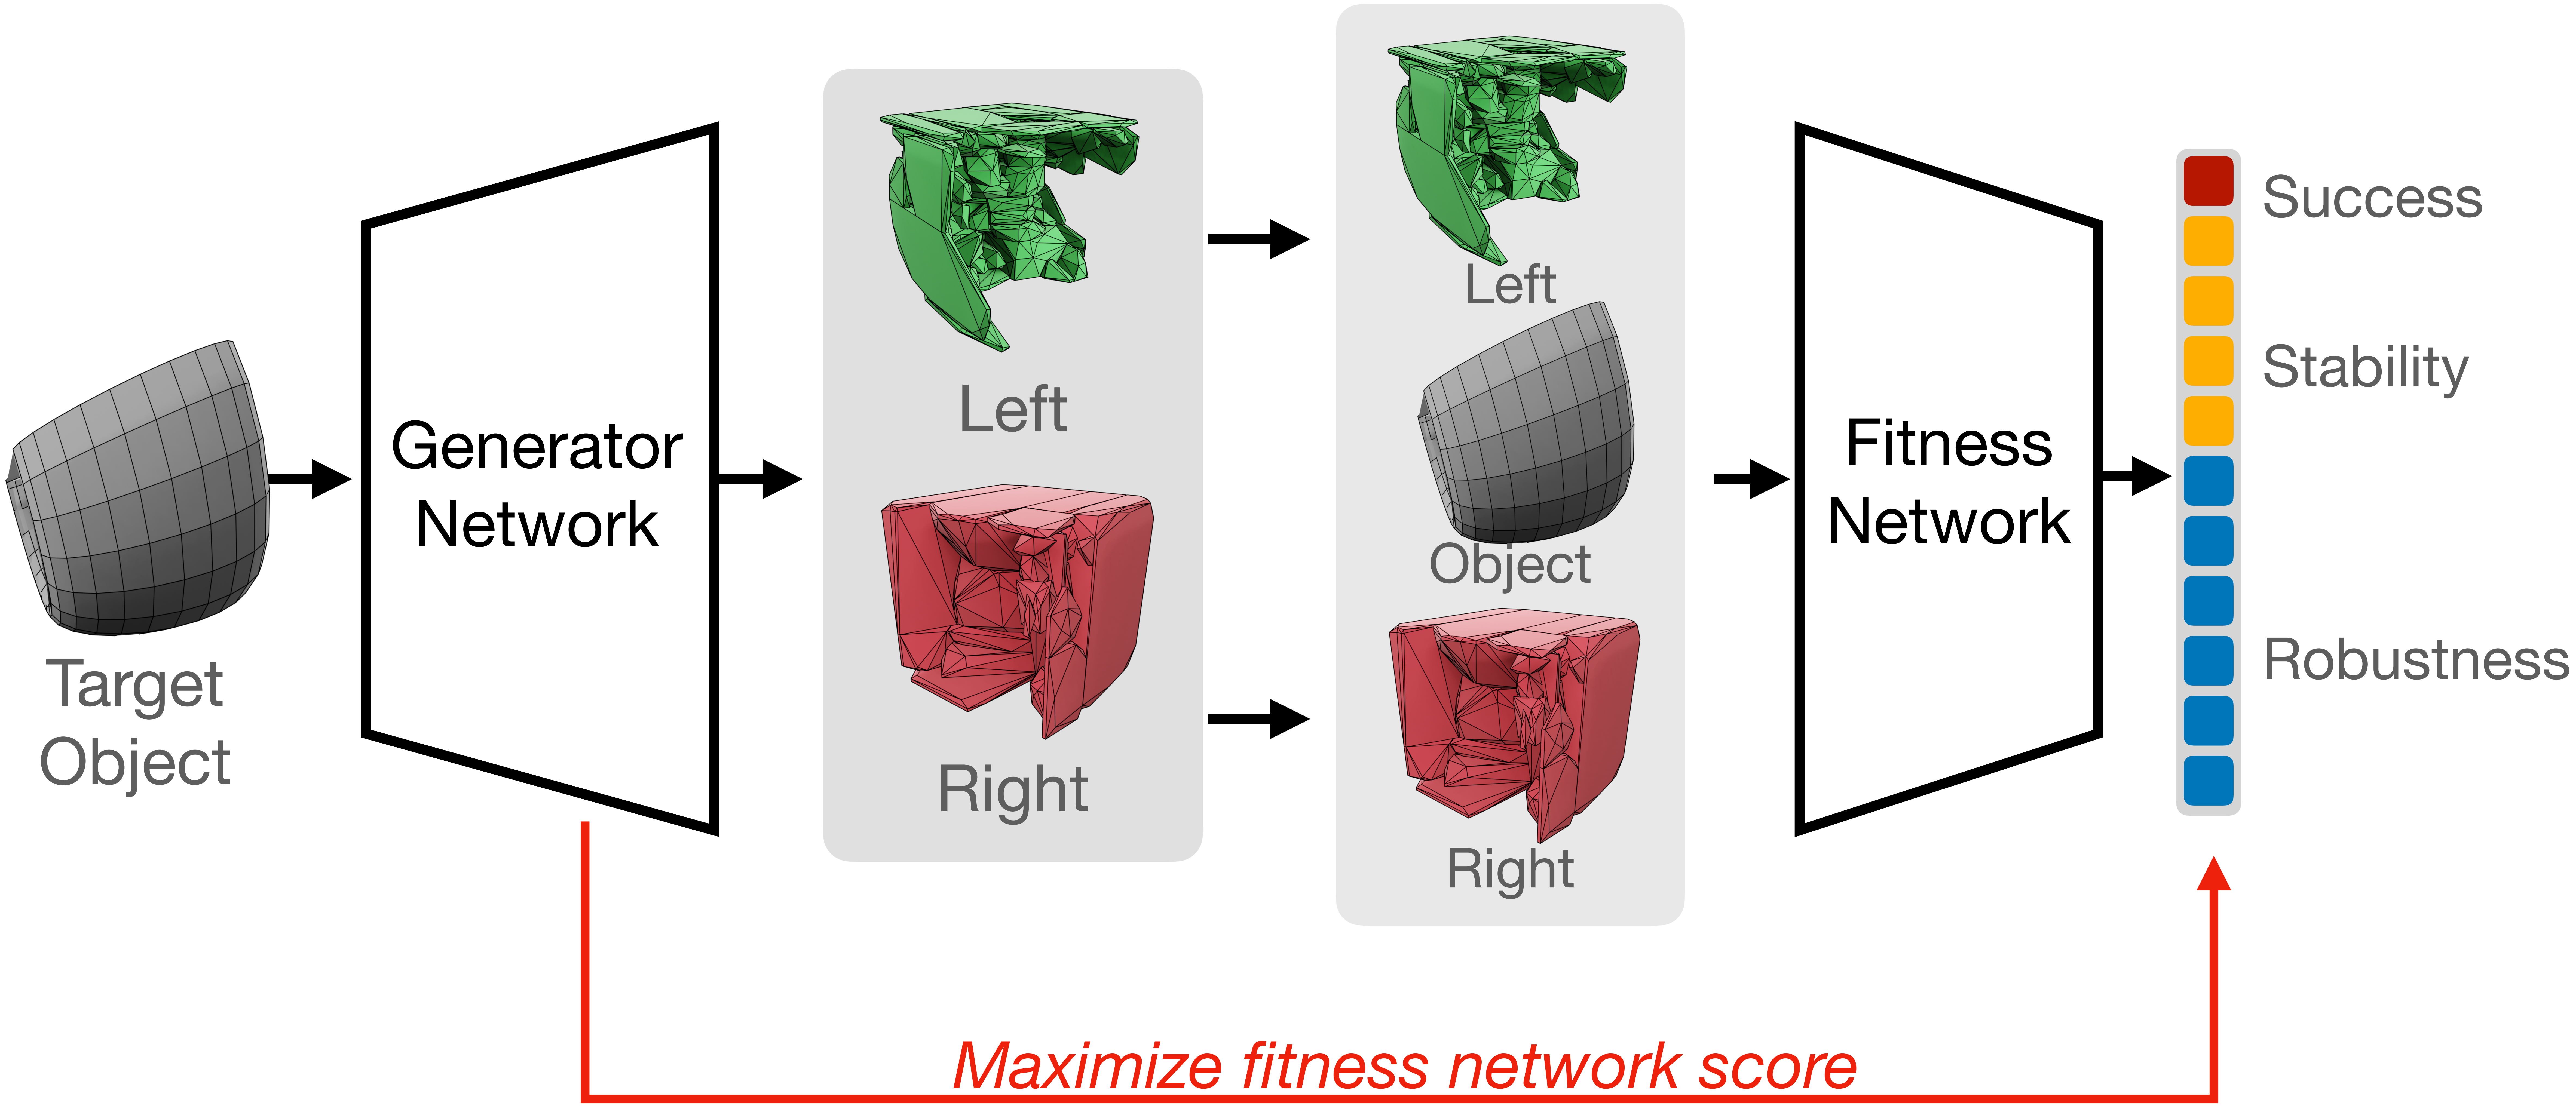

Target
Object
Generator
Network
Left
Right
Left
Object
Right
Fitness
Network
Success
Stability
Robustness
Maximize fitness network score

Initiate Training with Random Object Shapes

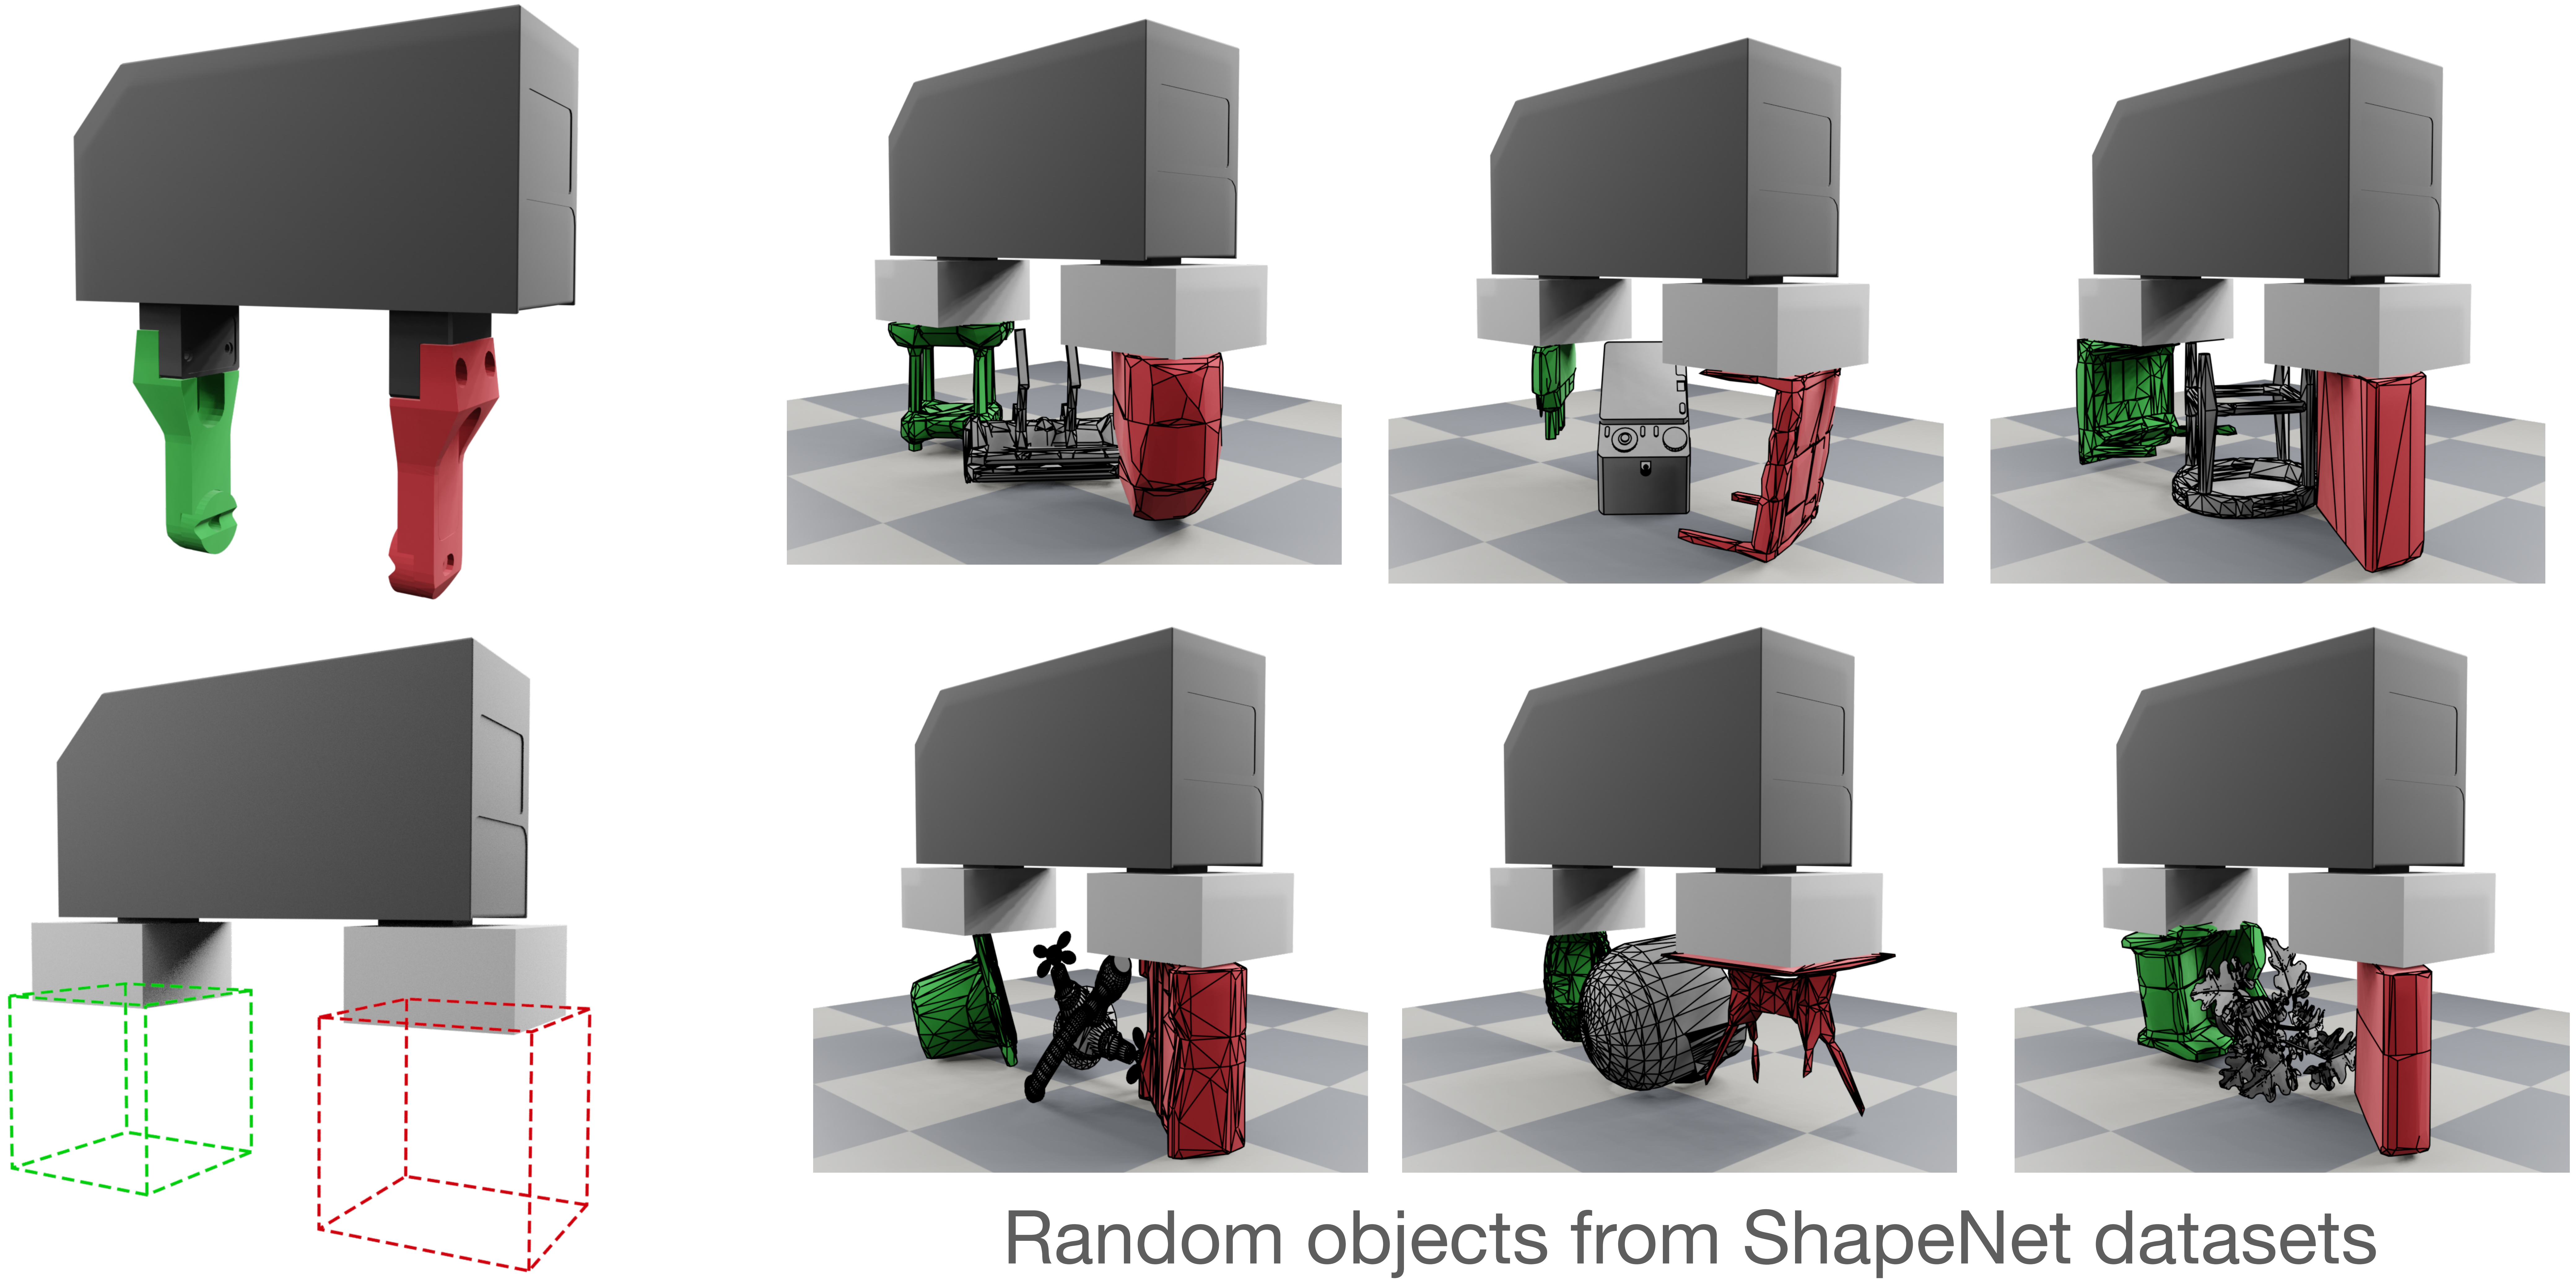

Random objects from ShapeNet datasets

Fit2Form

On Adversarial Grasp Objects [Wang *et al.* CASE'19]

Real world experiments

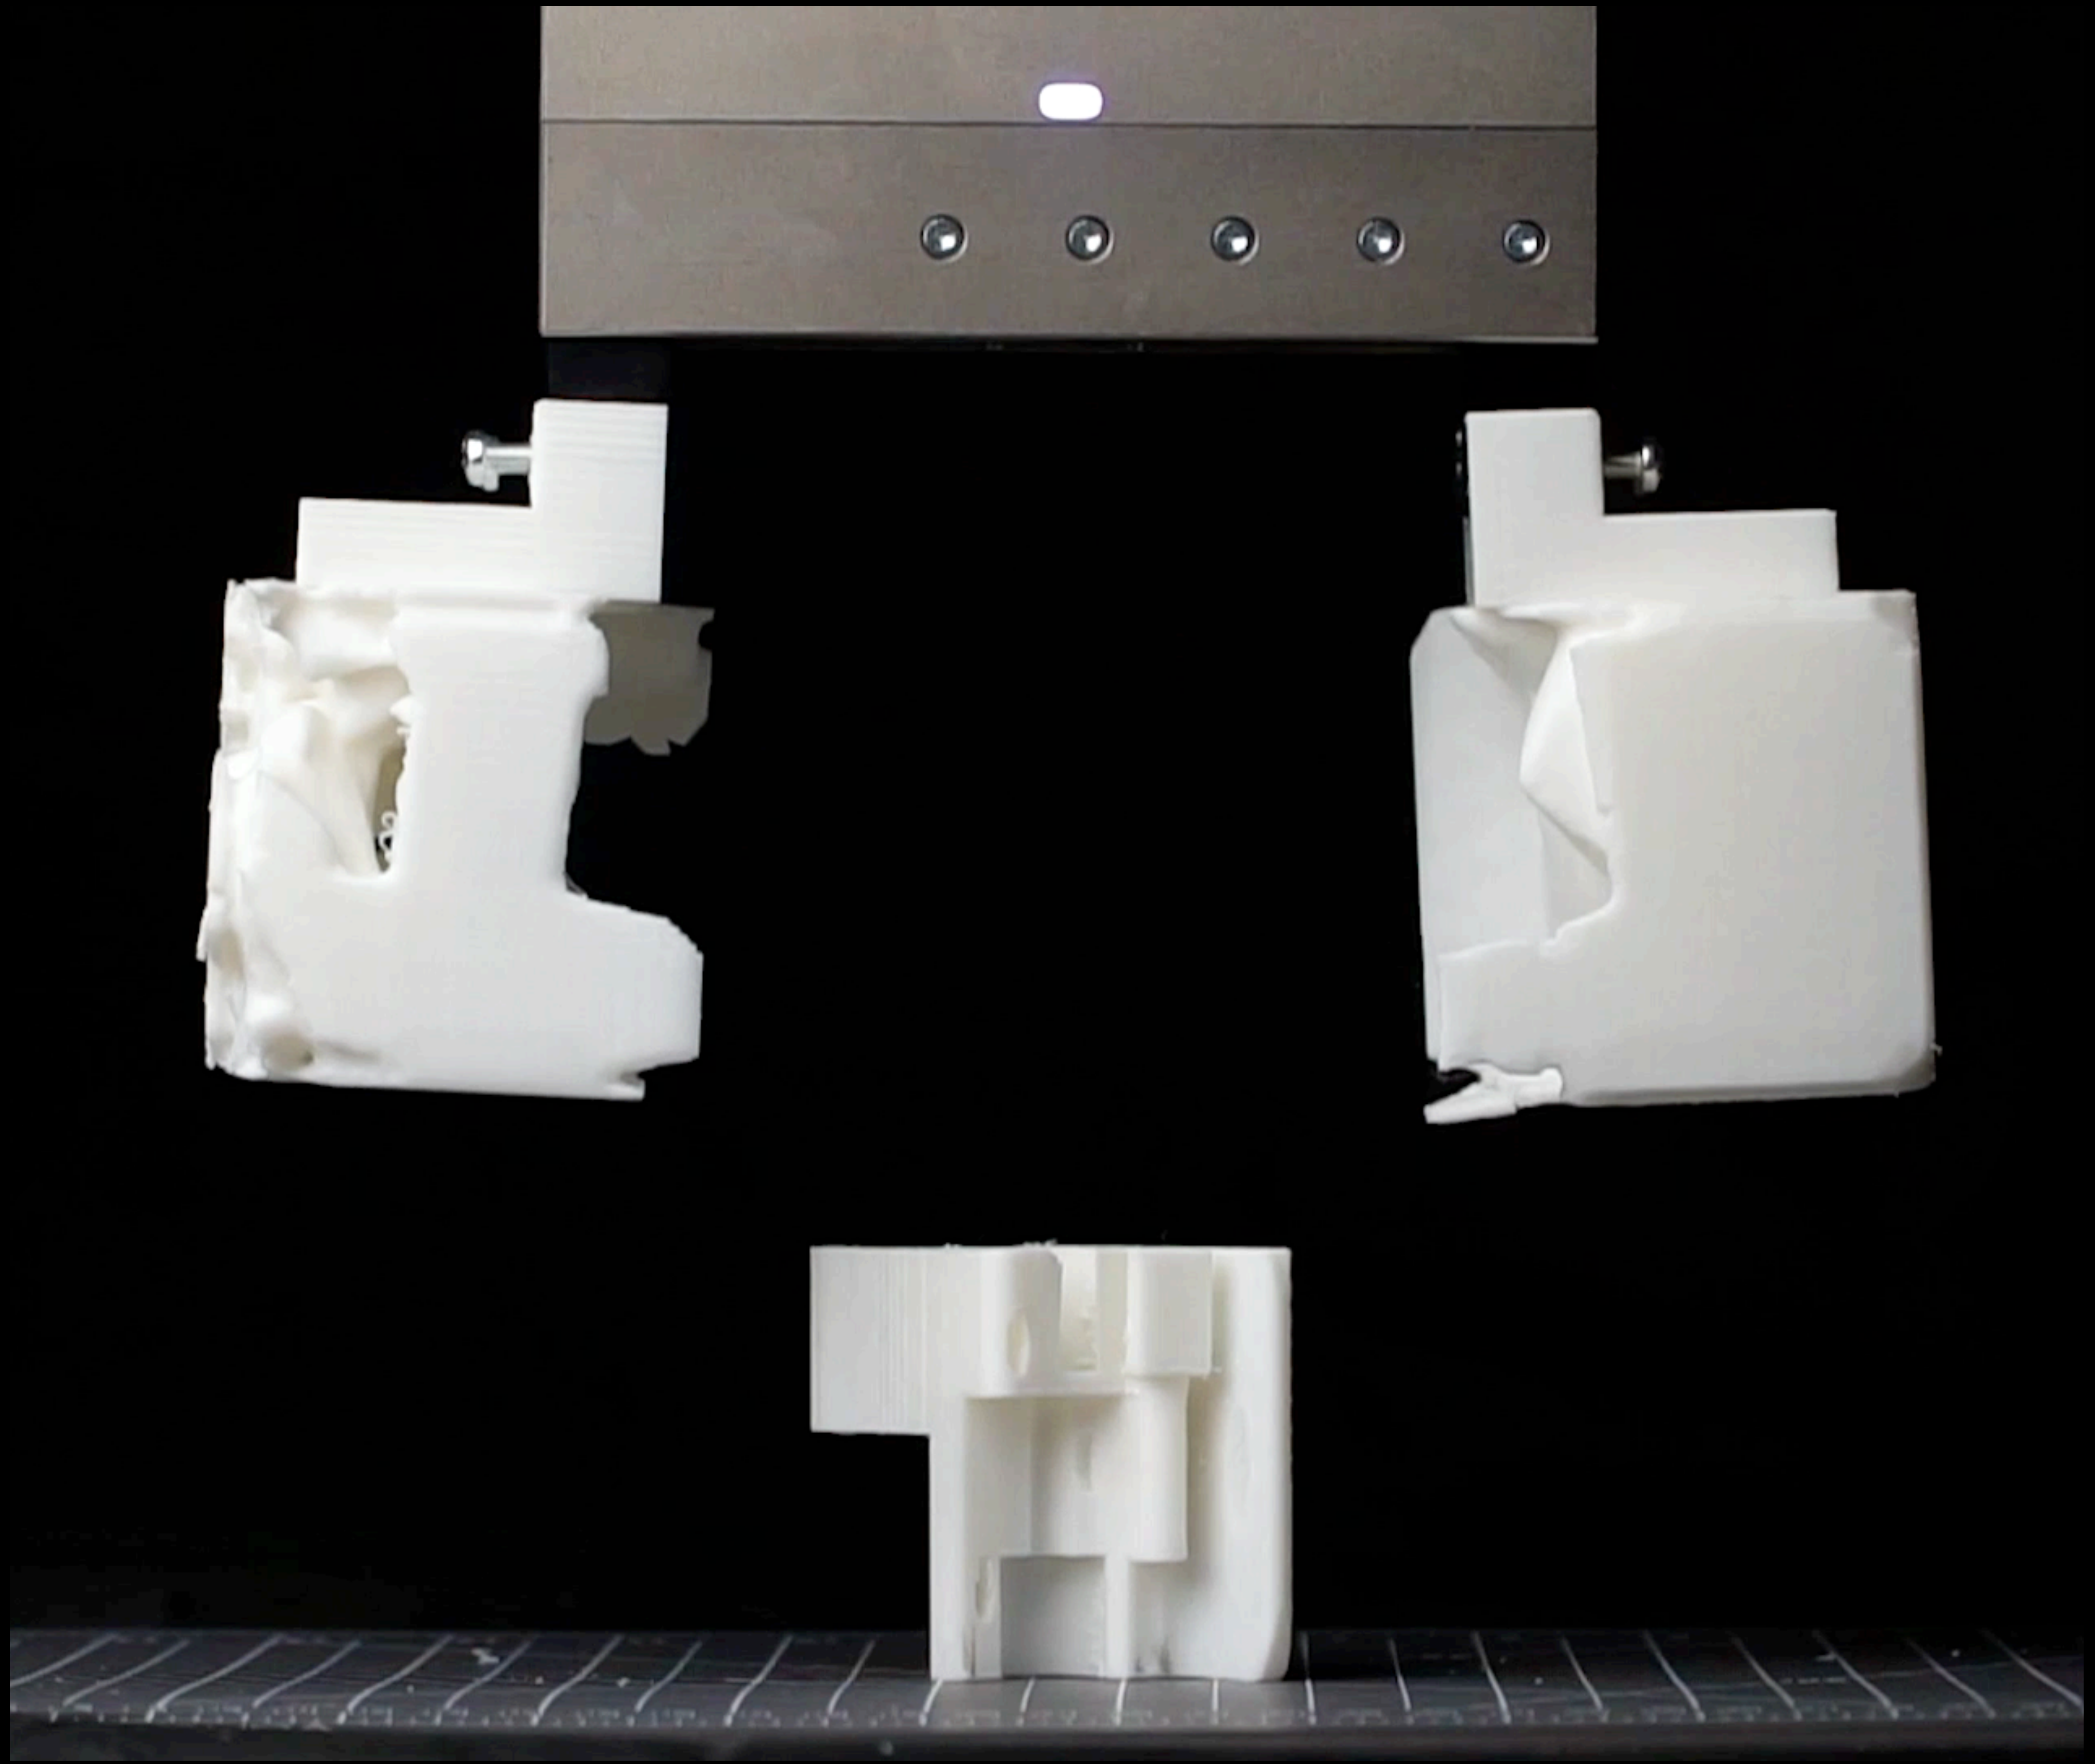

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Fit2Form

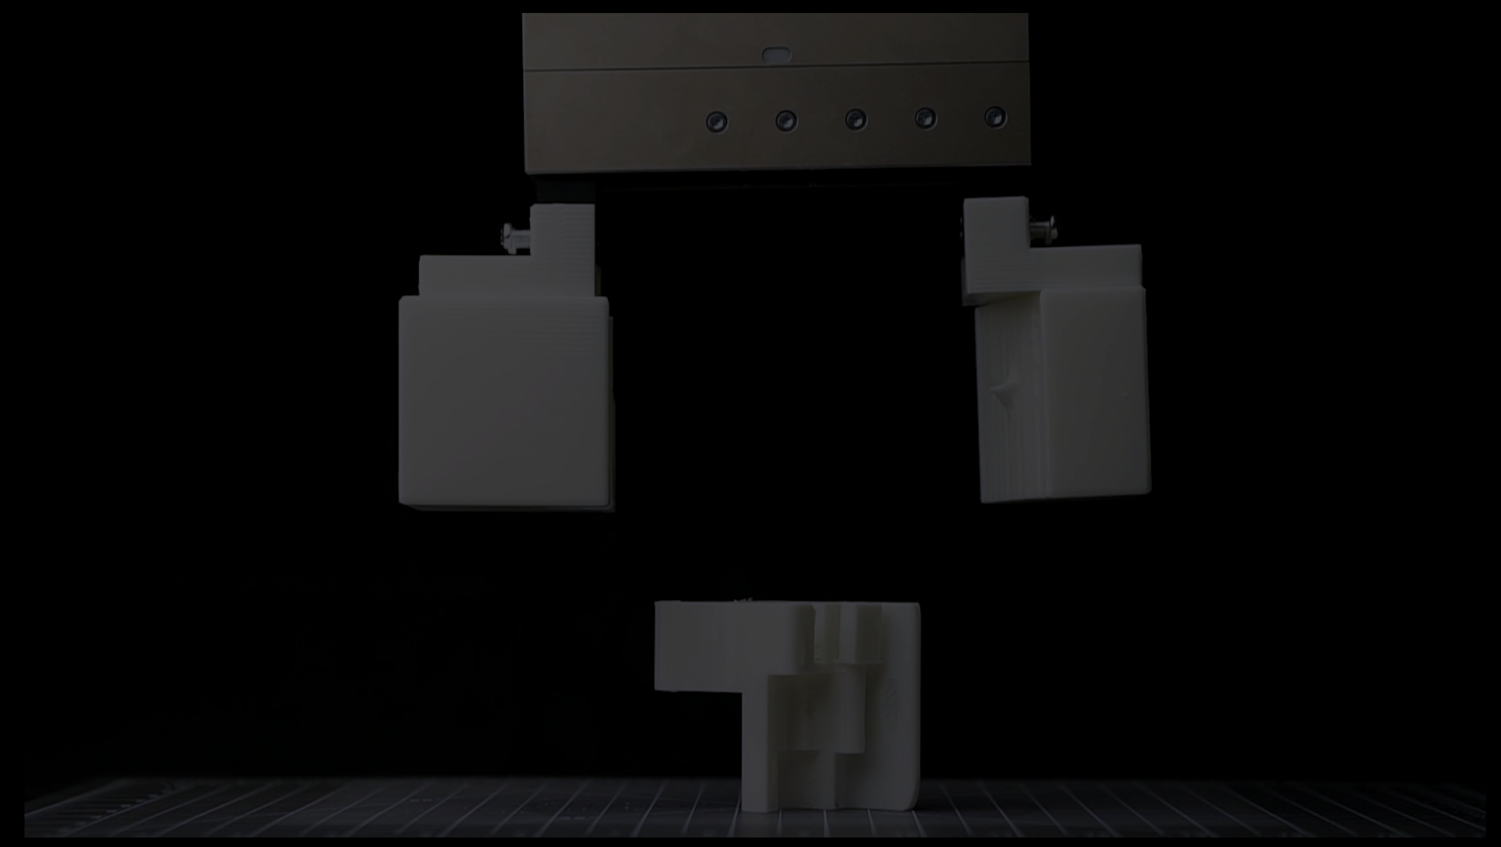

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Imprint

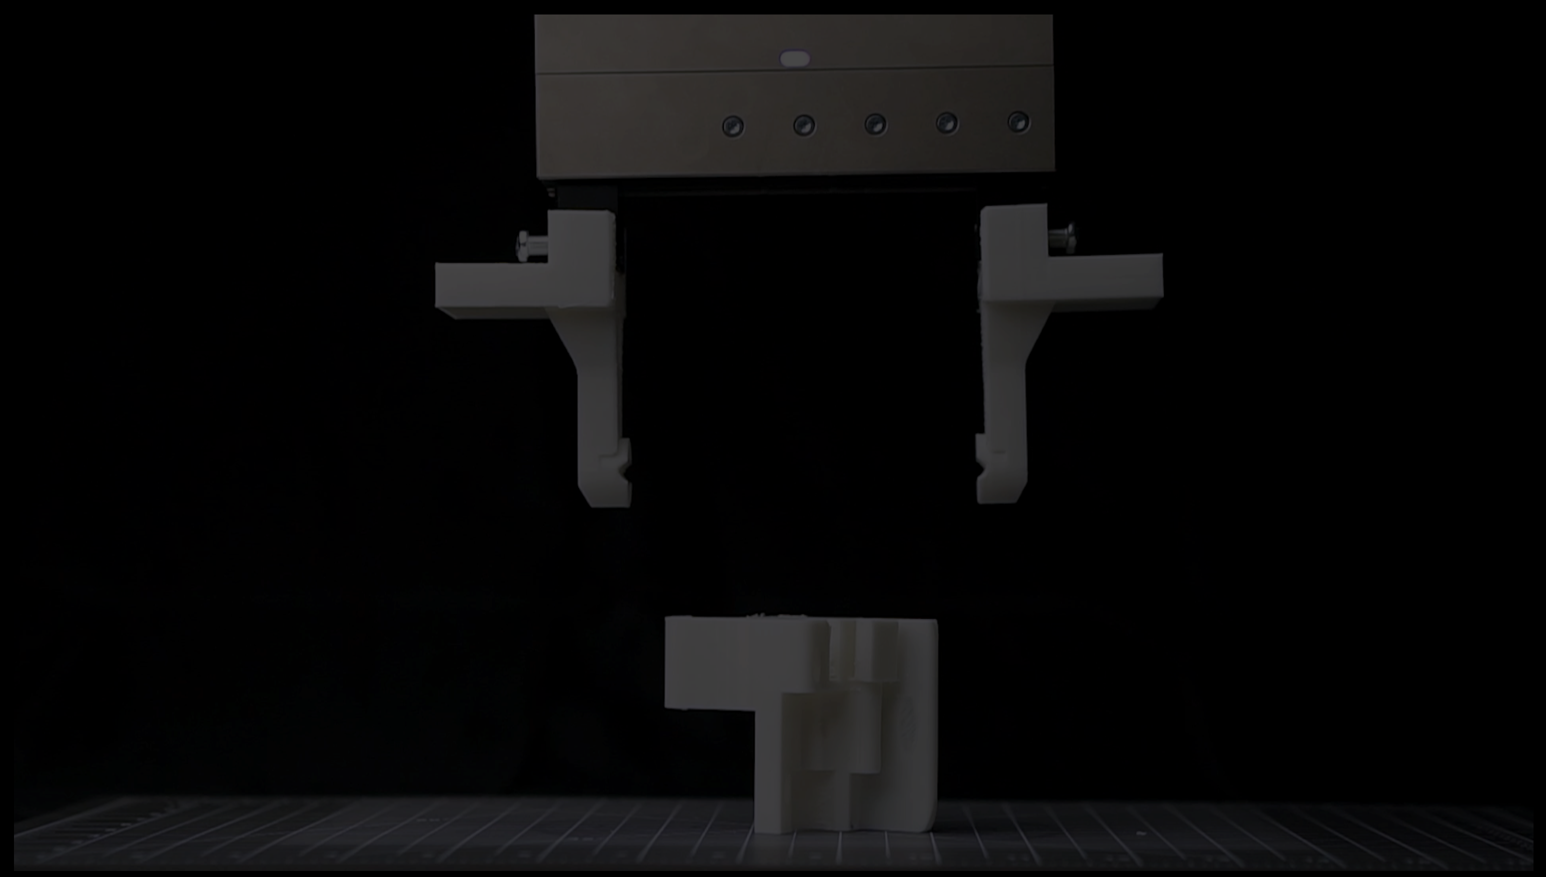

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

WSG 50

Real world experiments

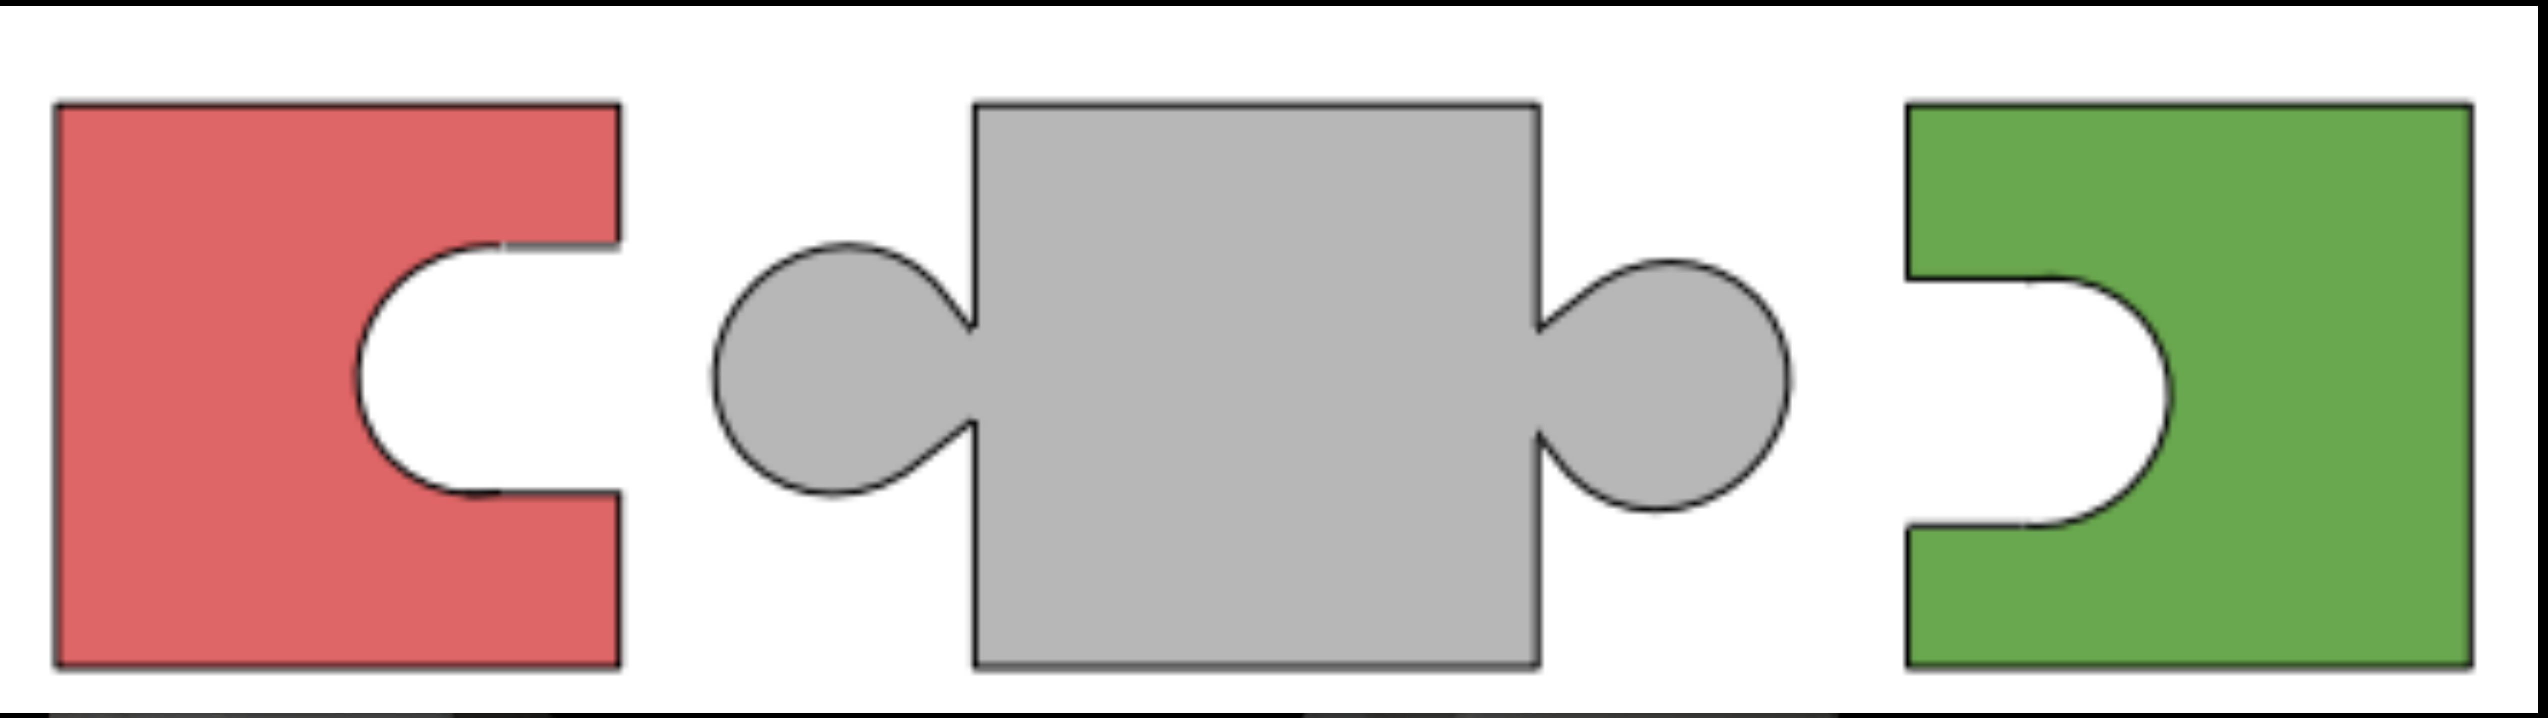

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Fit2Form

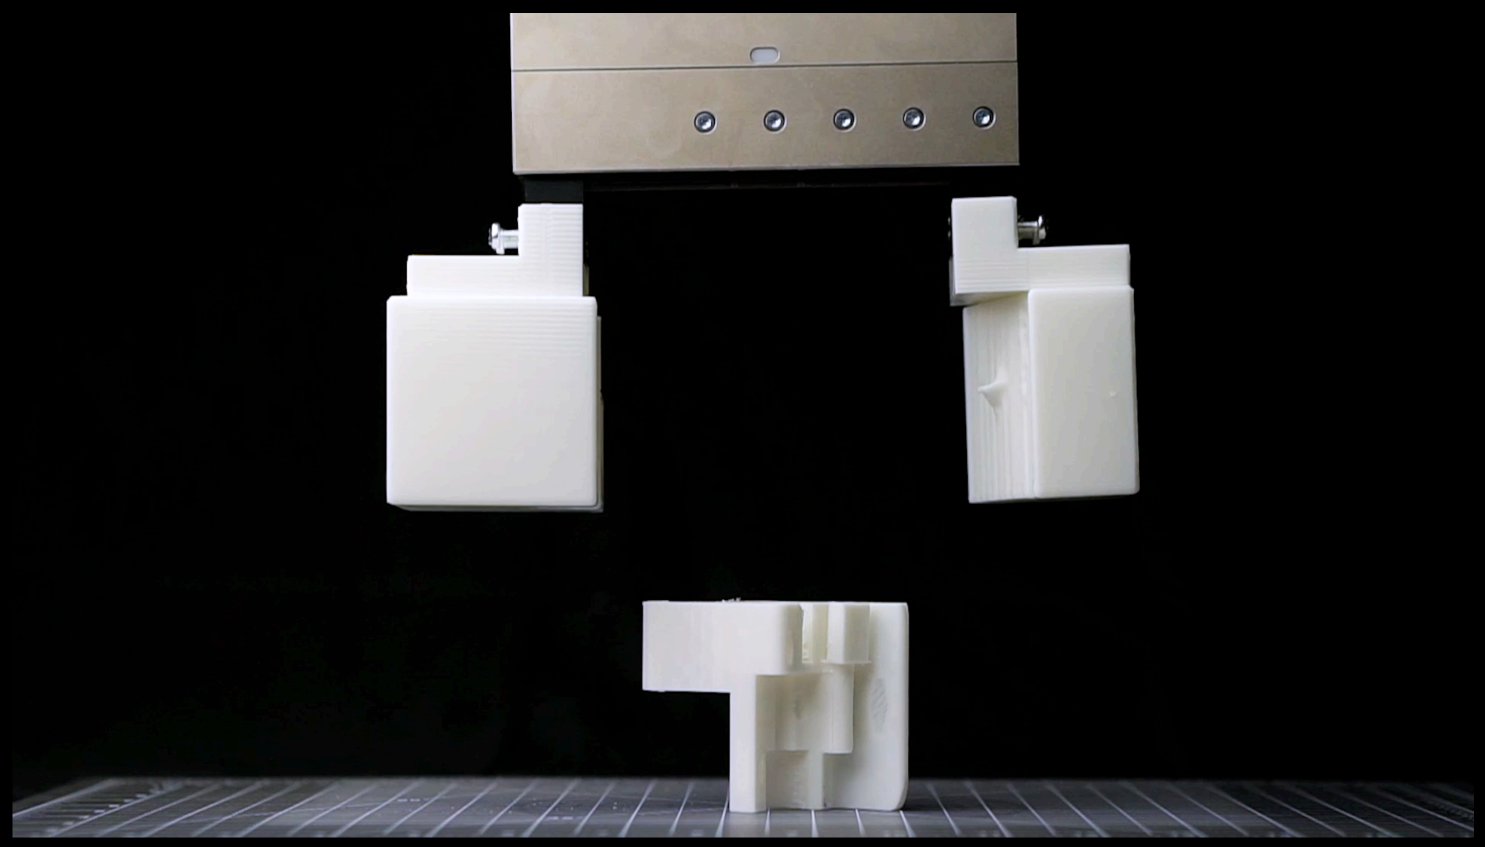

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Imprint

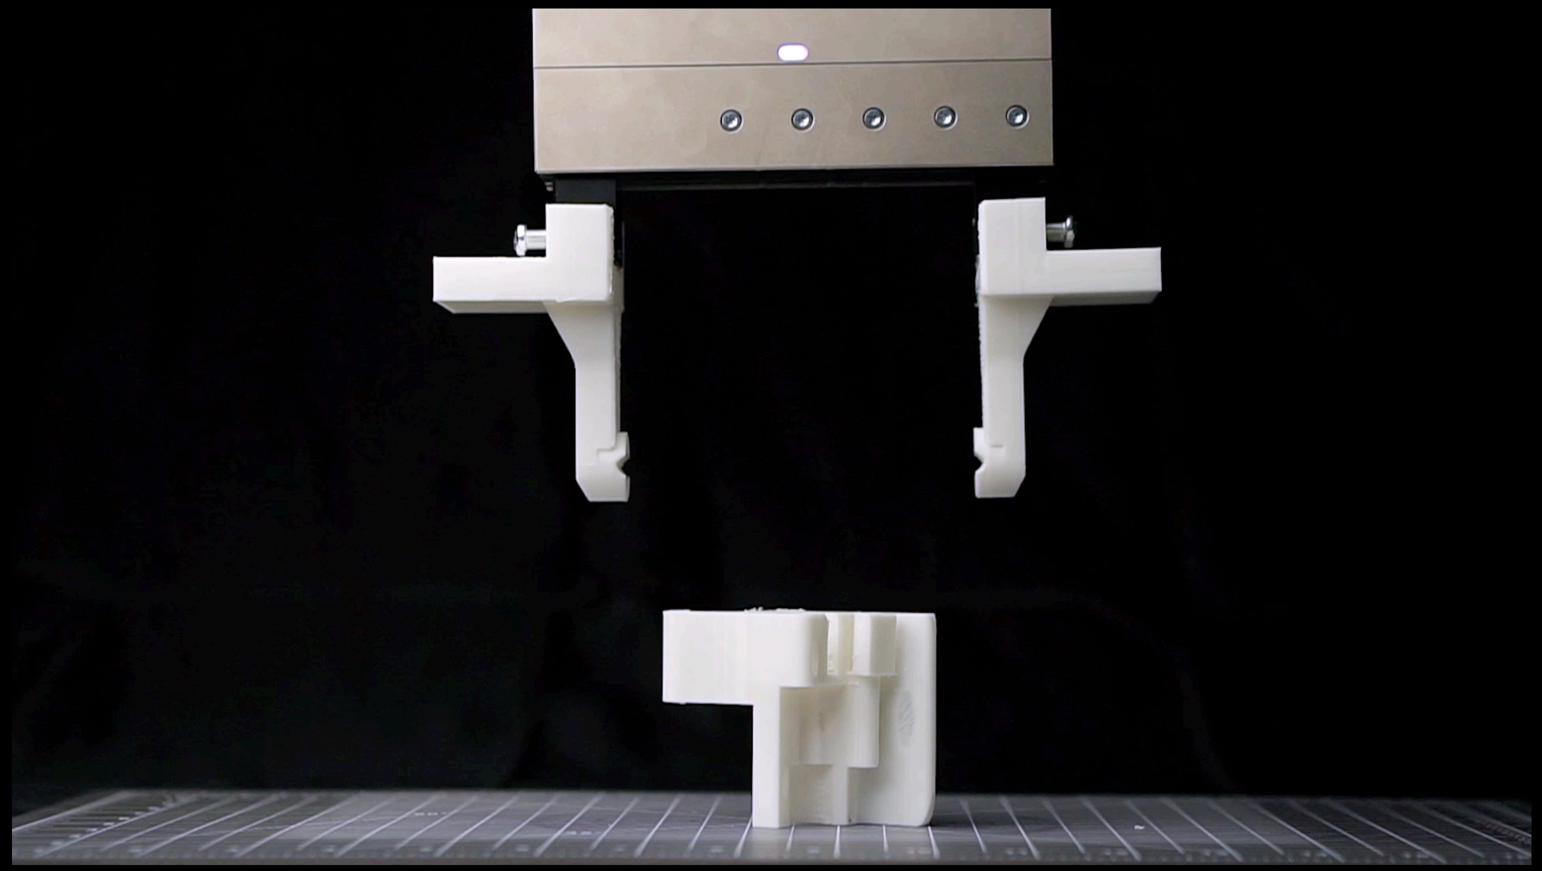

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

WSG 50

Real world experiments

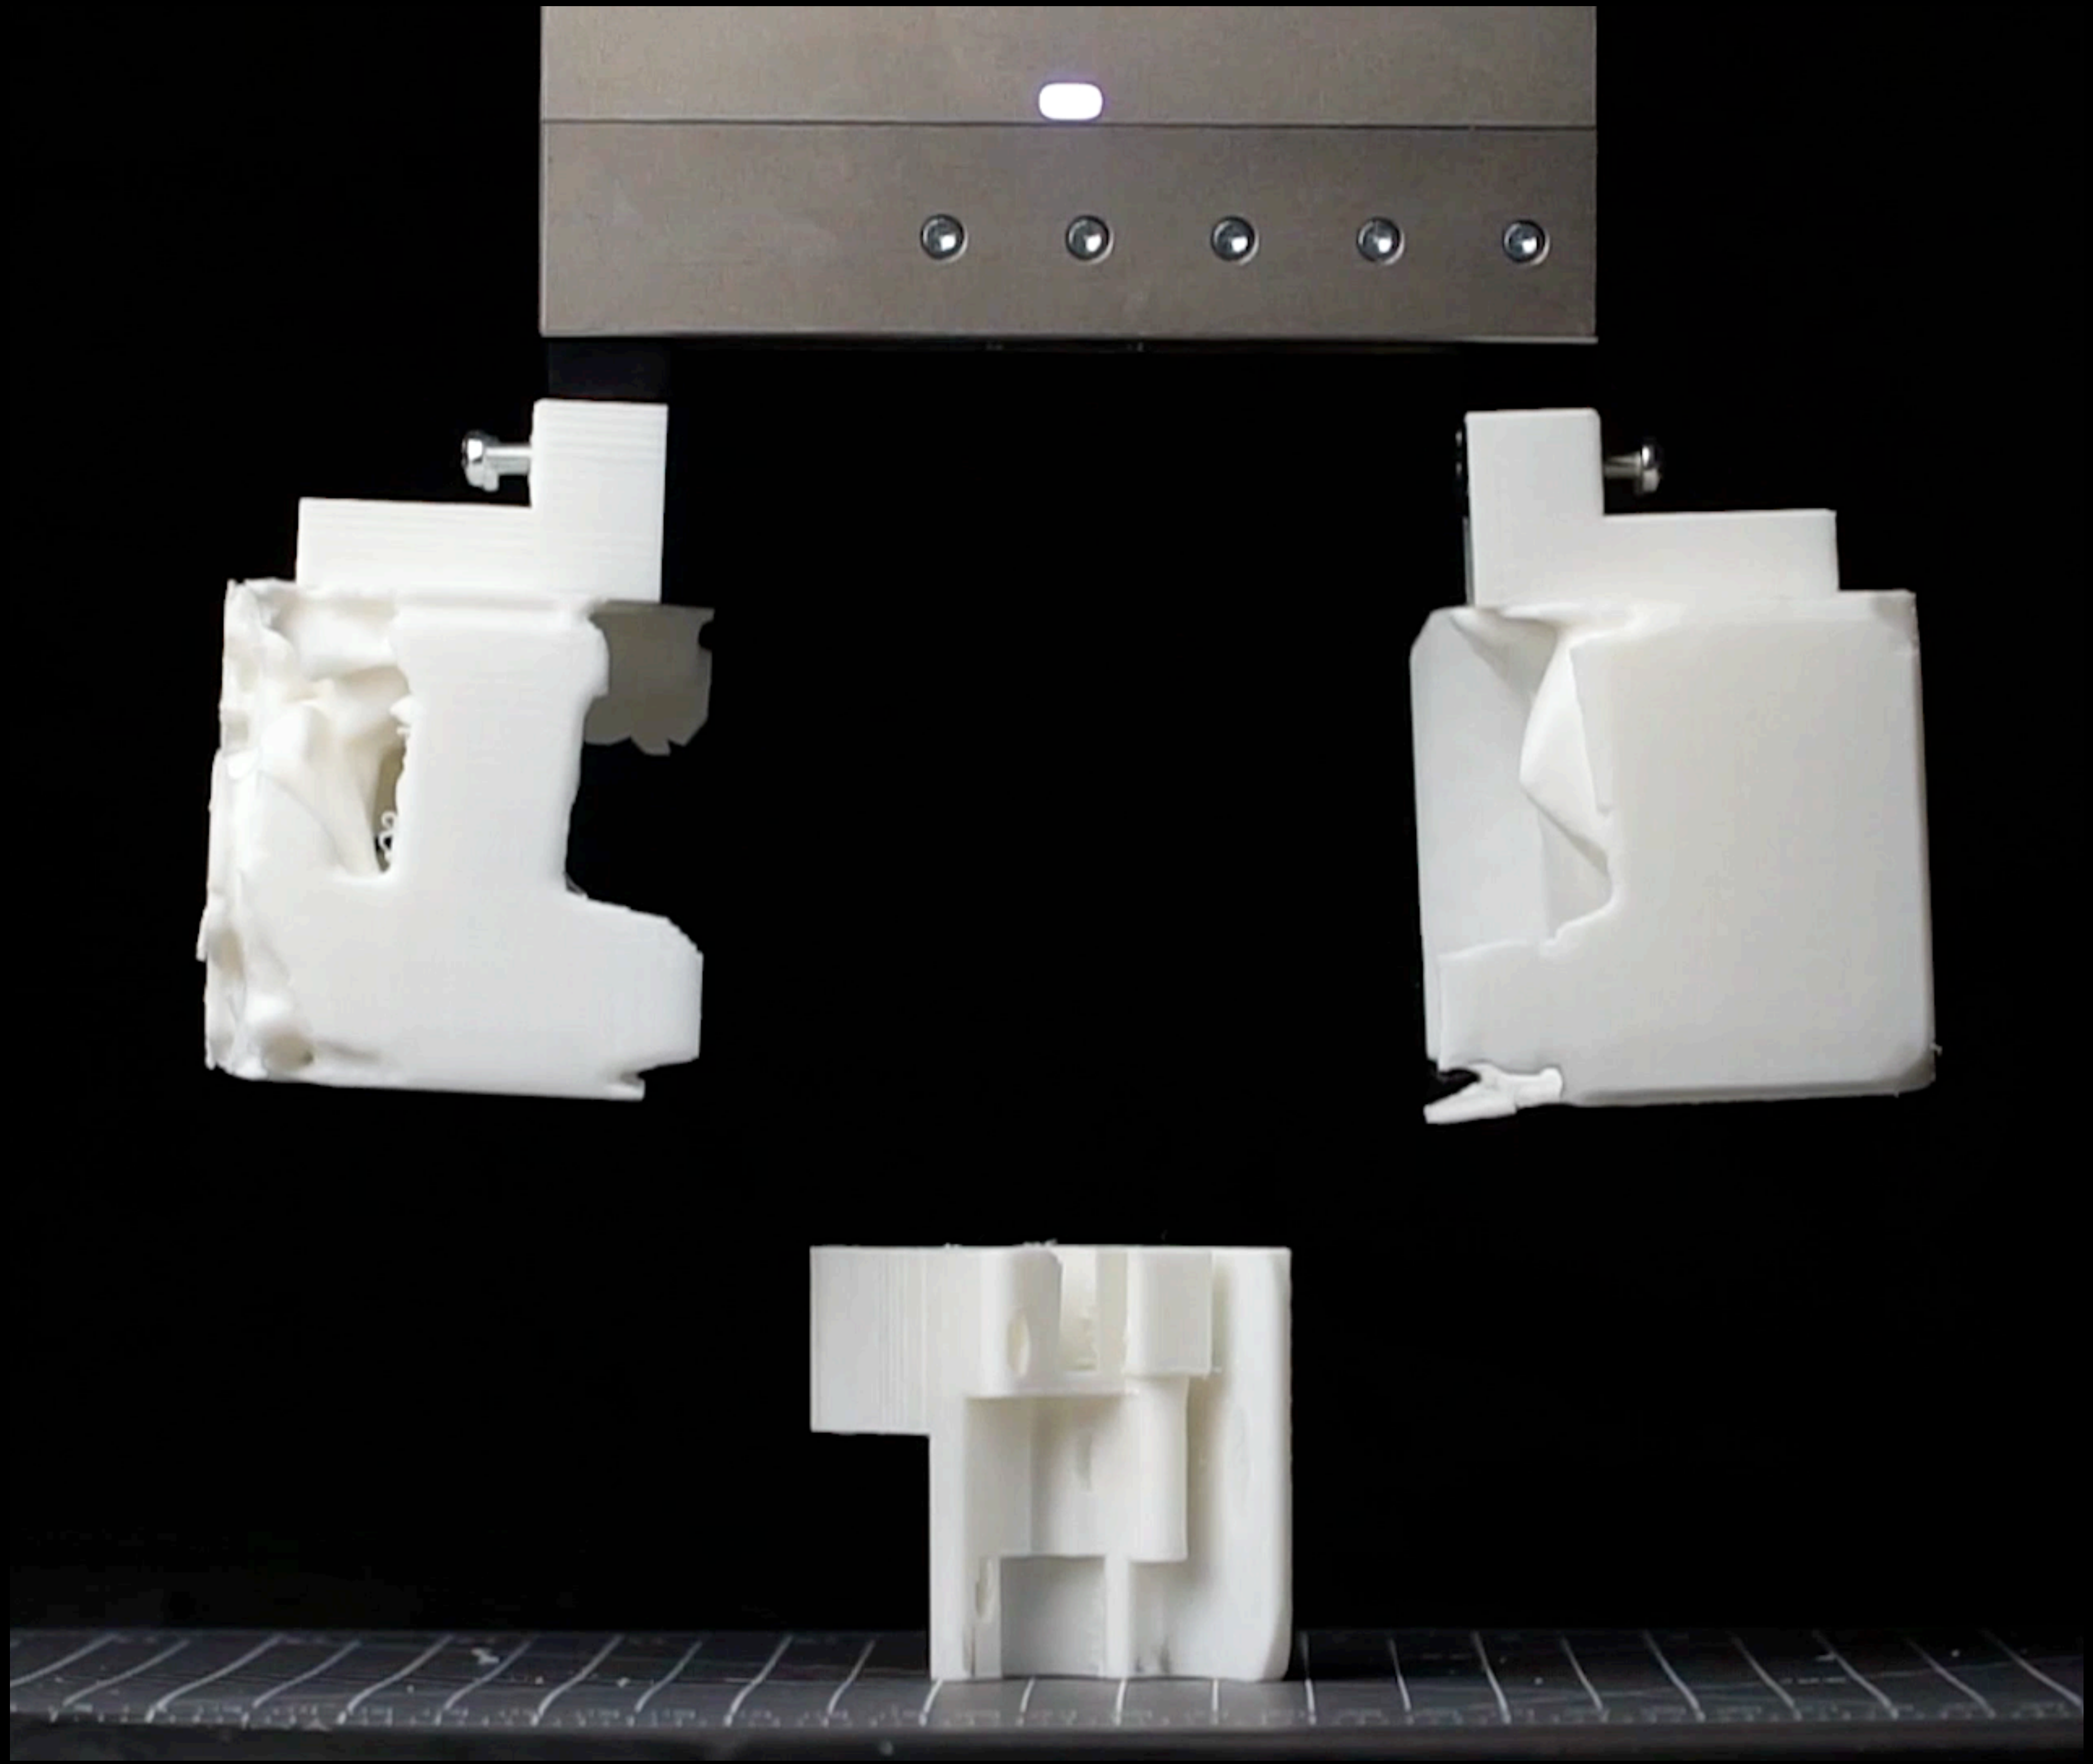

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Fit2Form

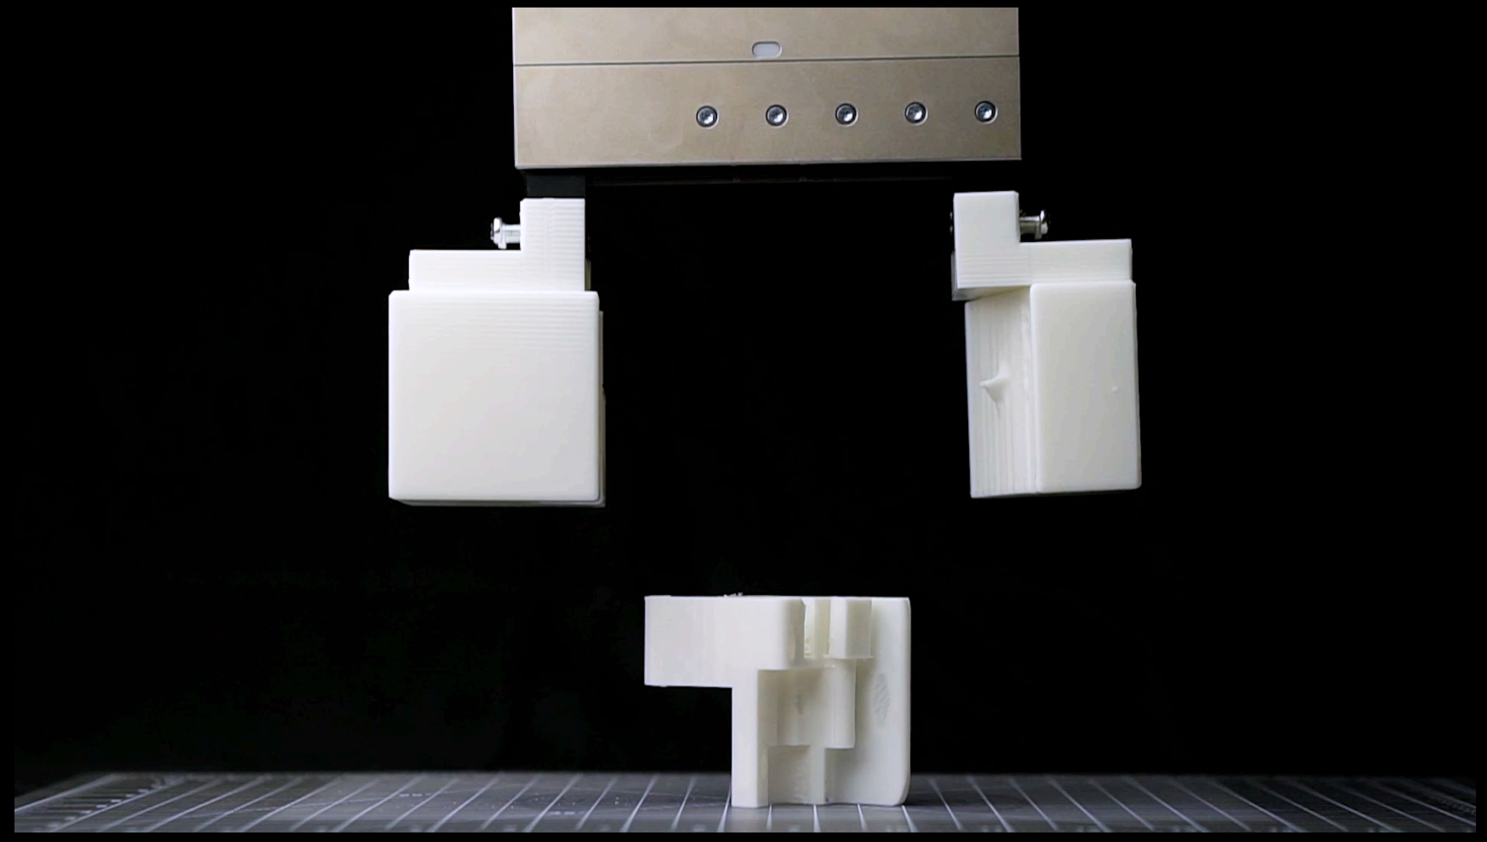

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

Imprint

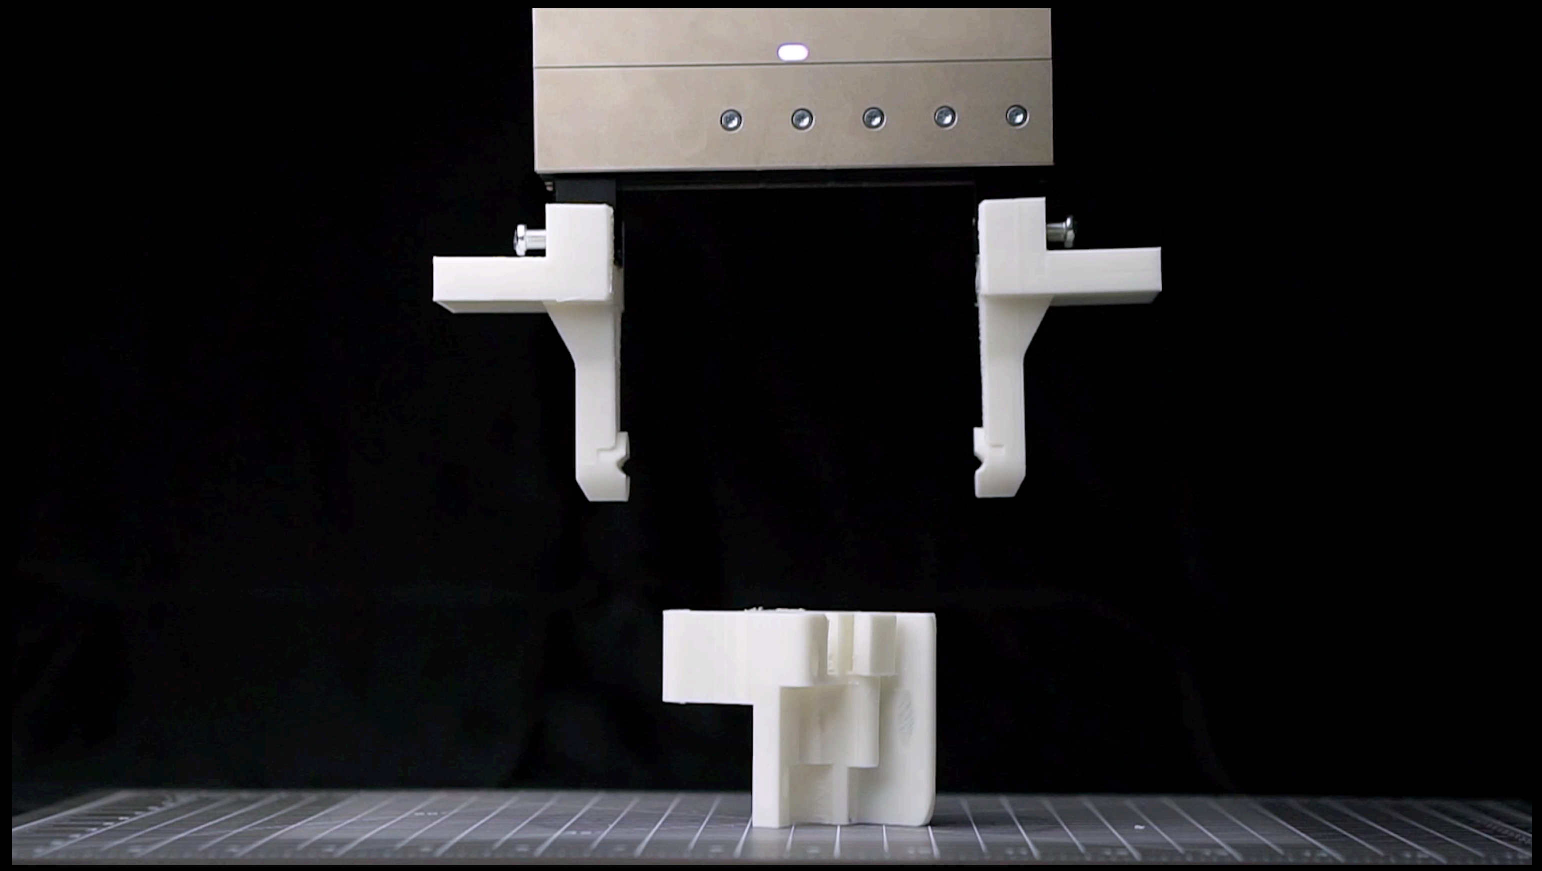

G S_0 S_1 S_2 S_3 S_4 R_1 R_2 R_3 R_4 R_5 R_6

WSG 50

Ablation Study

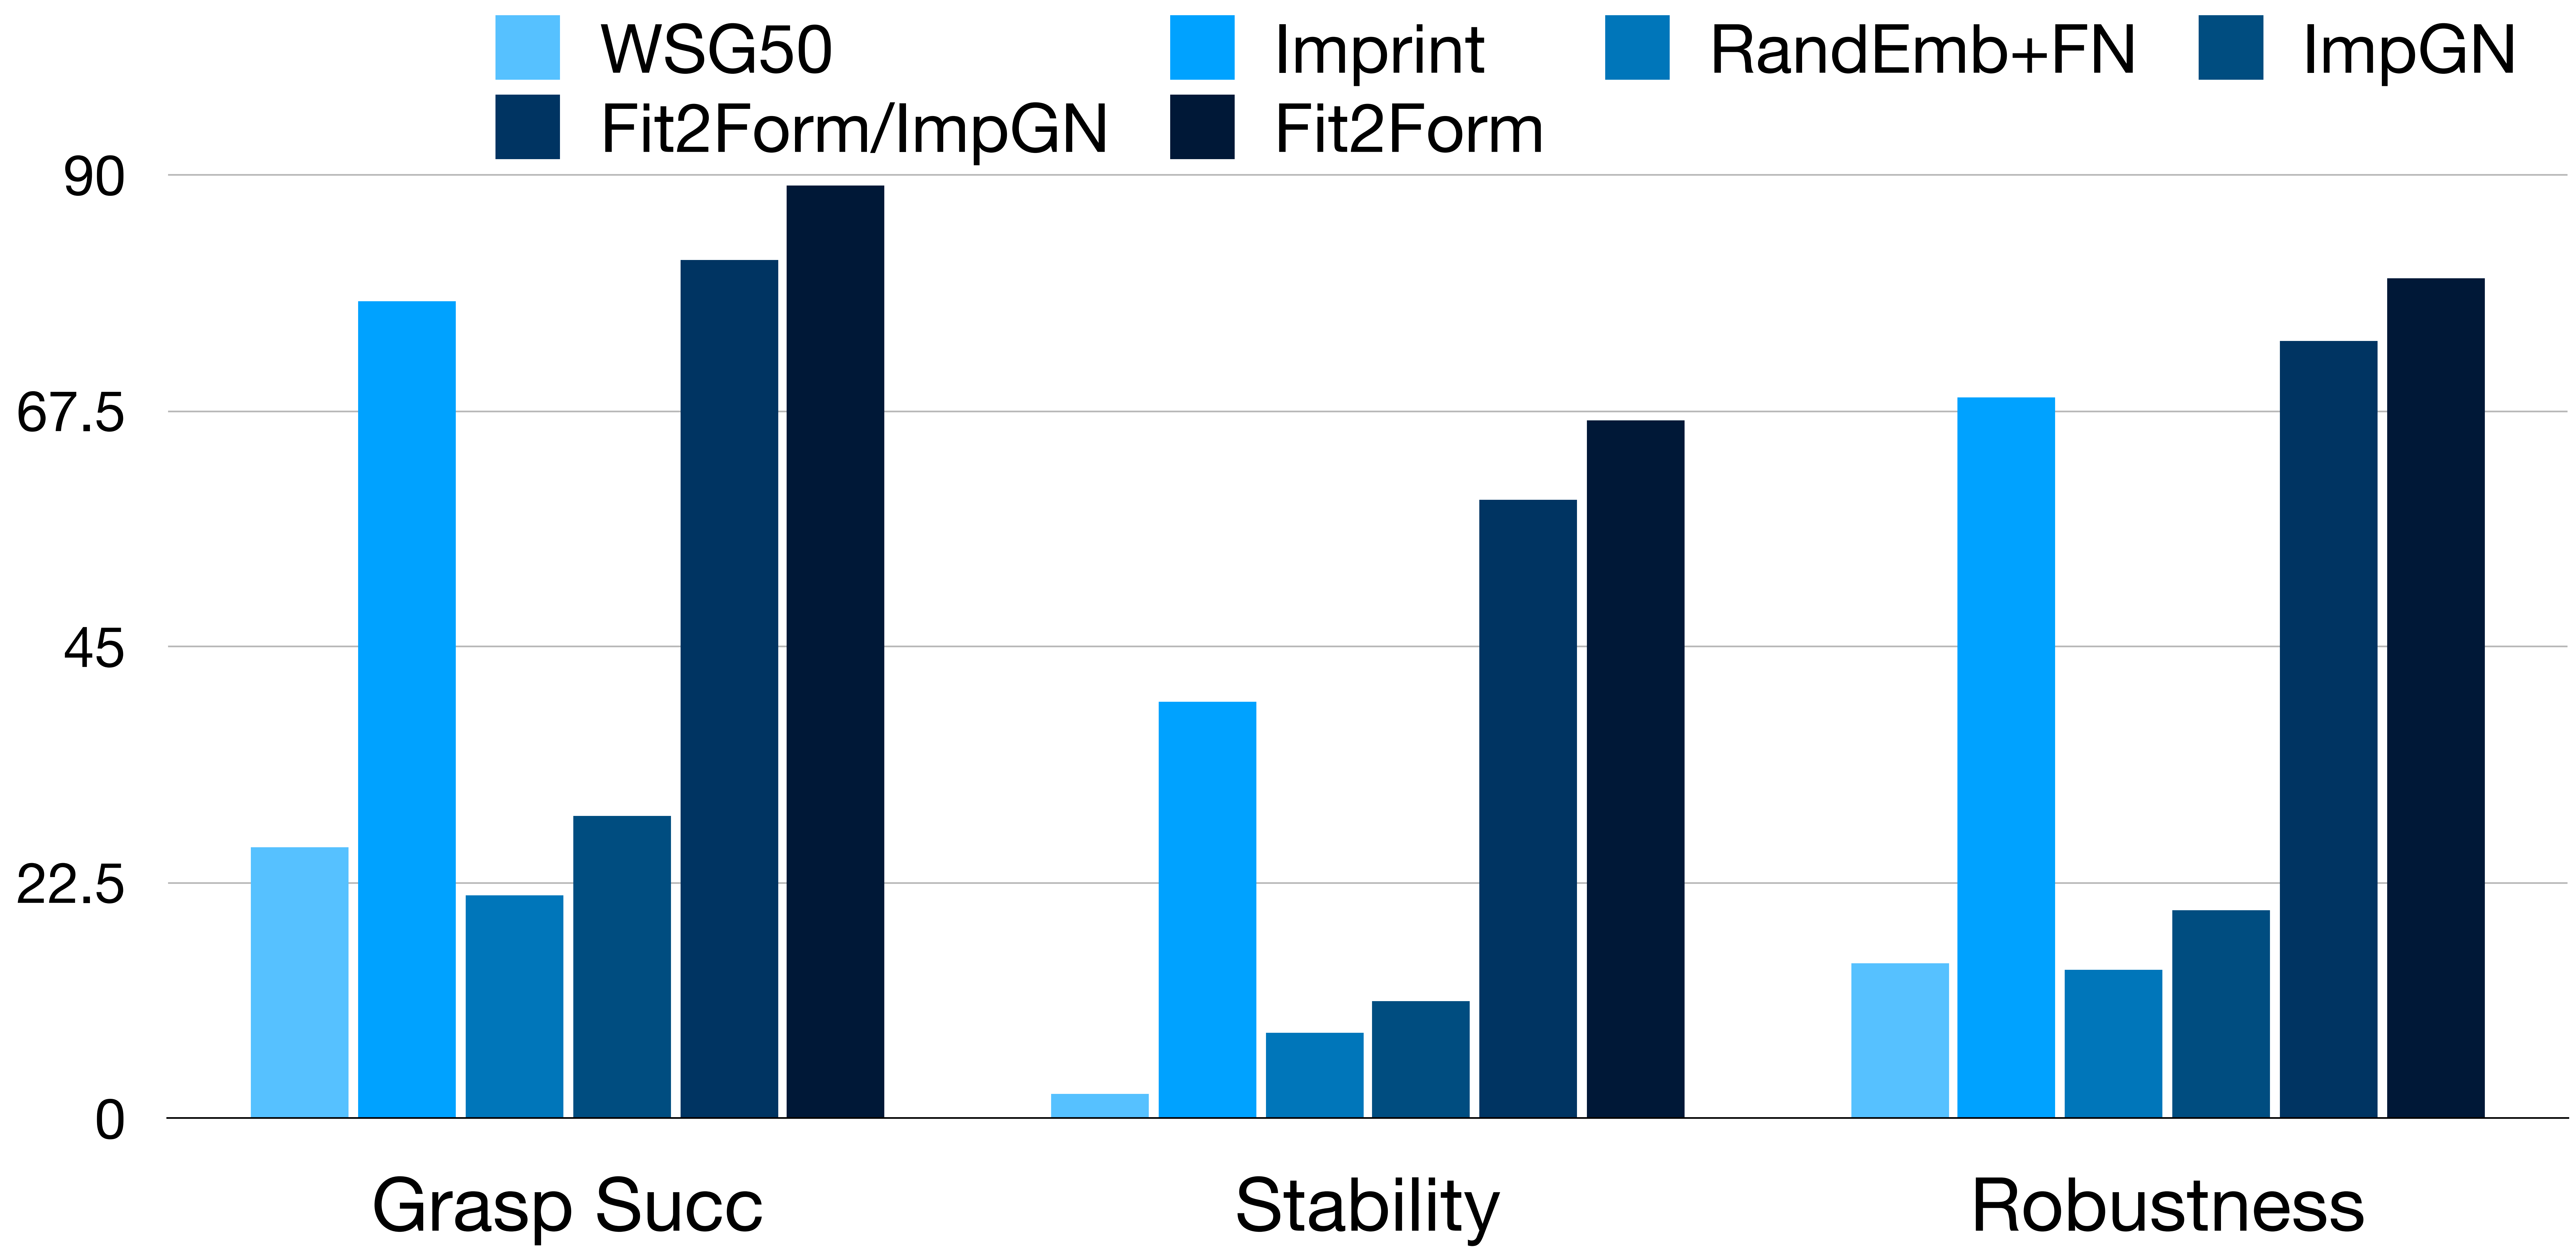

What does Fit2Form Learn?

Single-sided Scoop

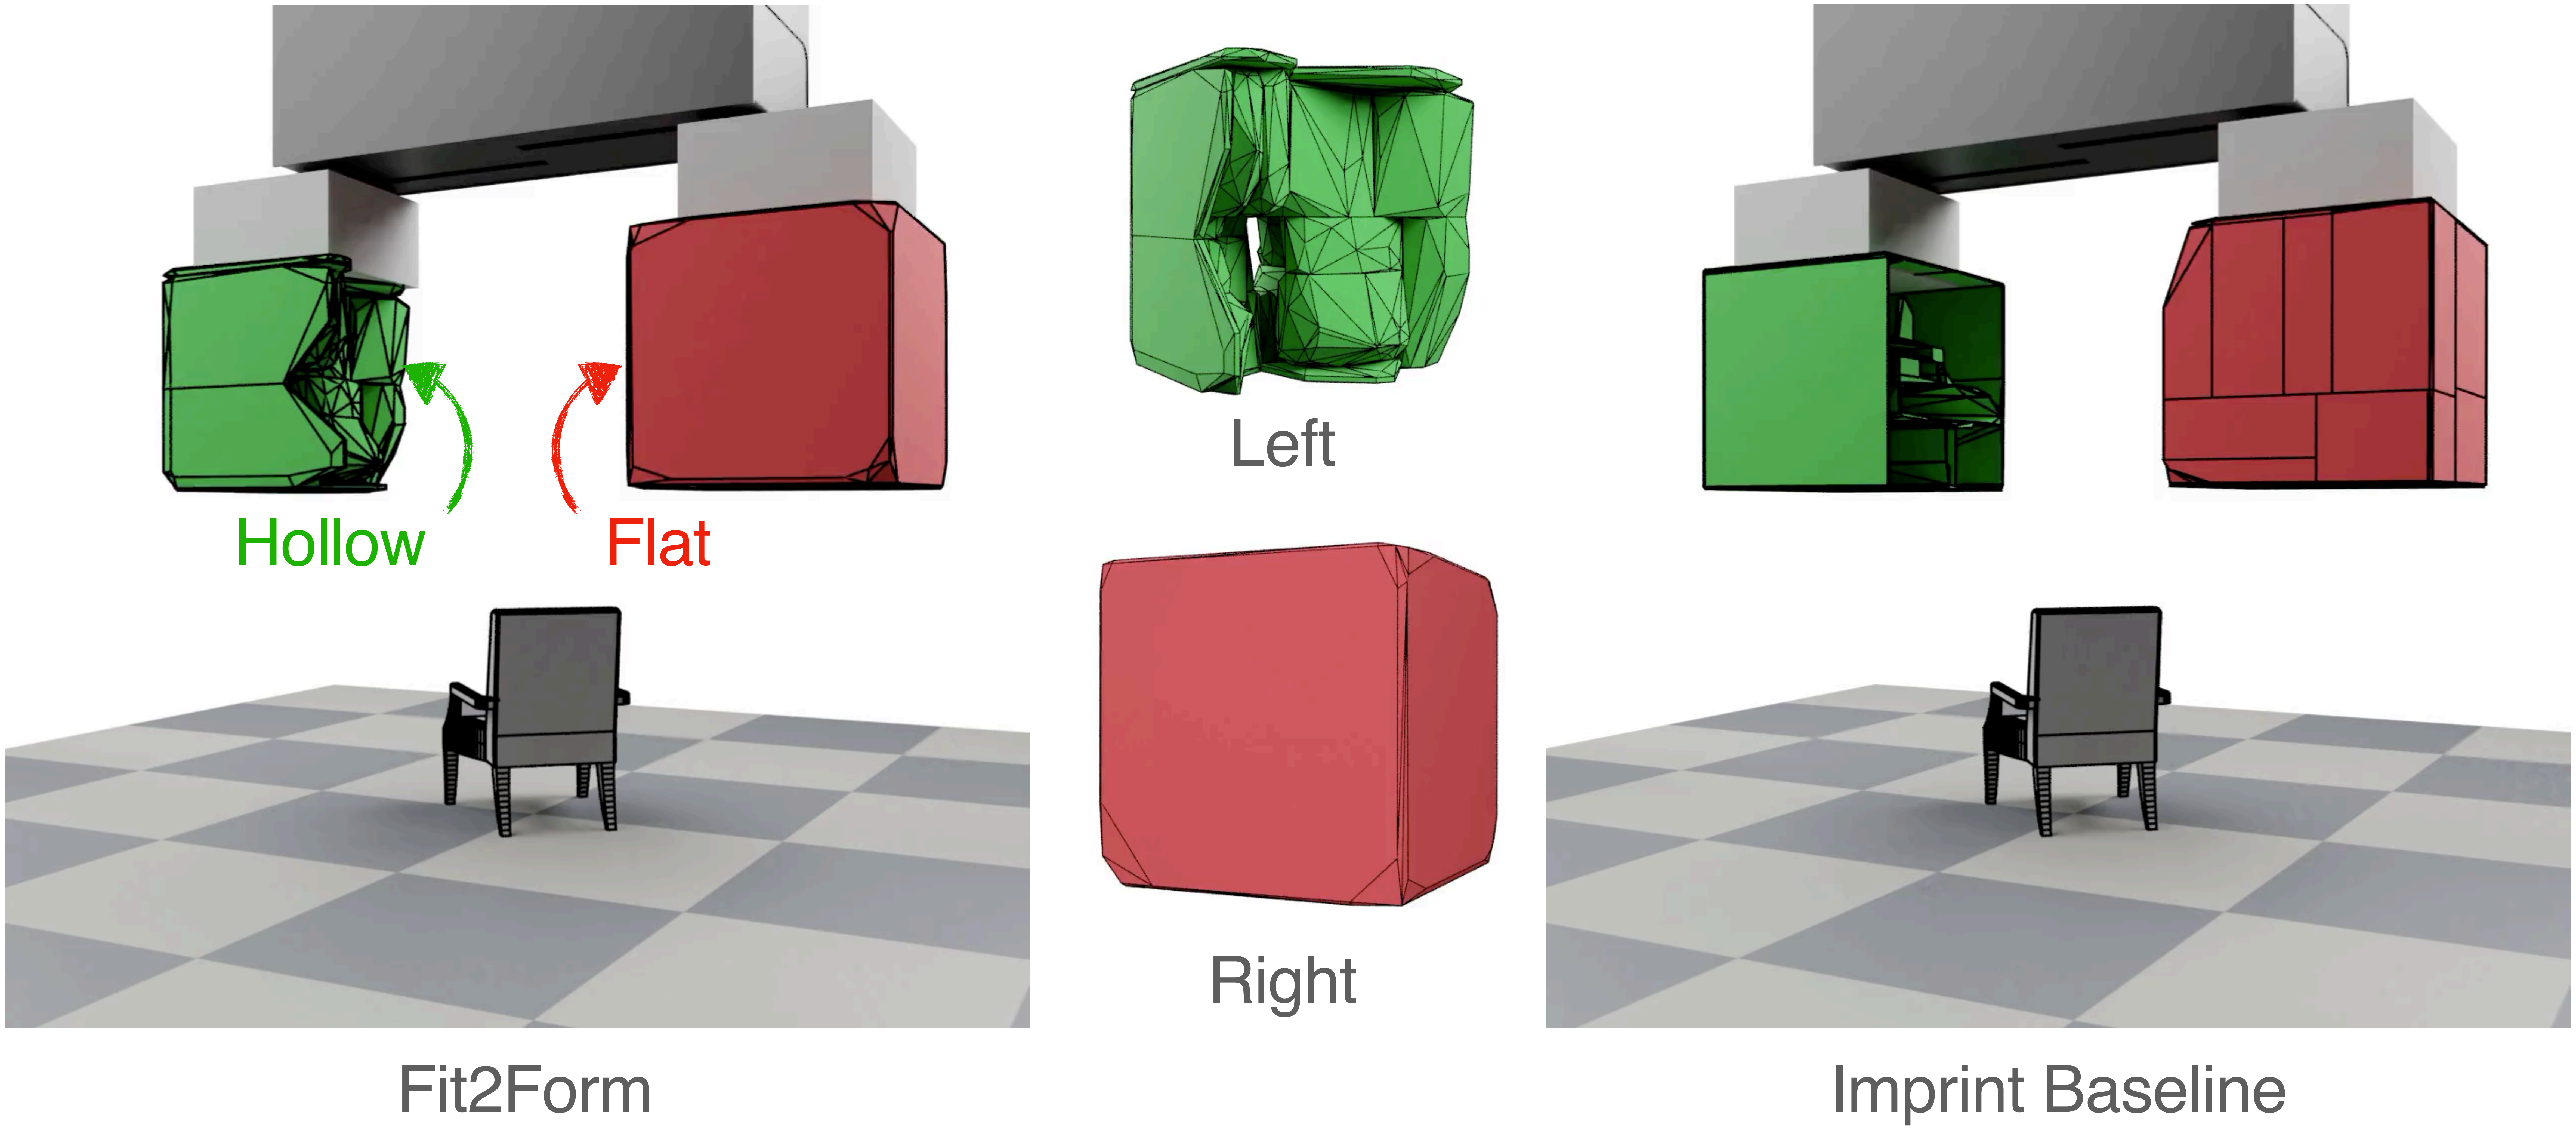

Double-sided Scoop

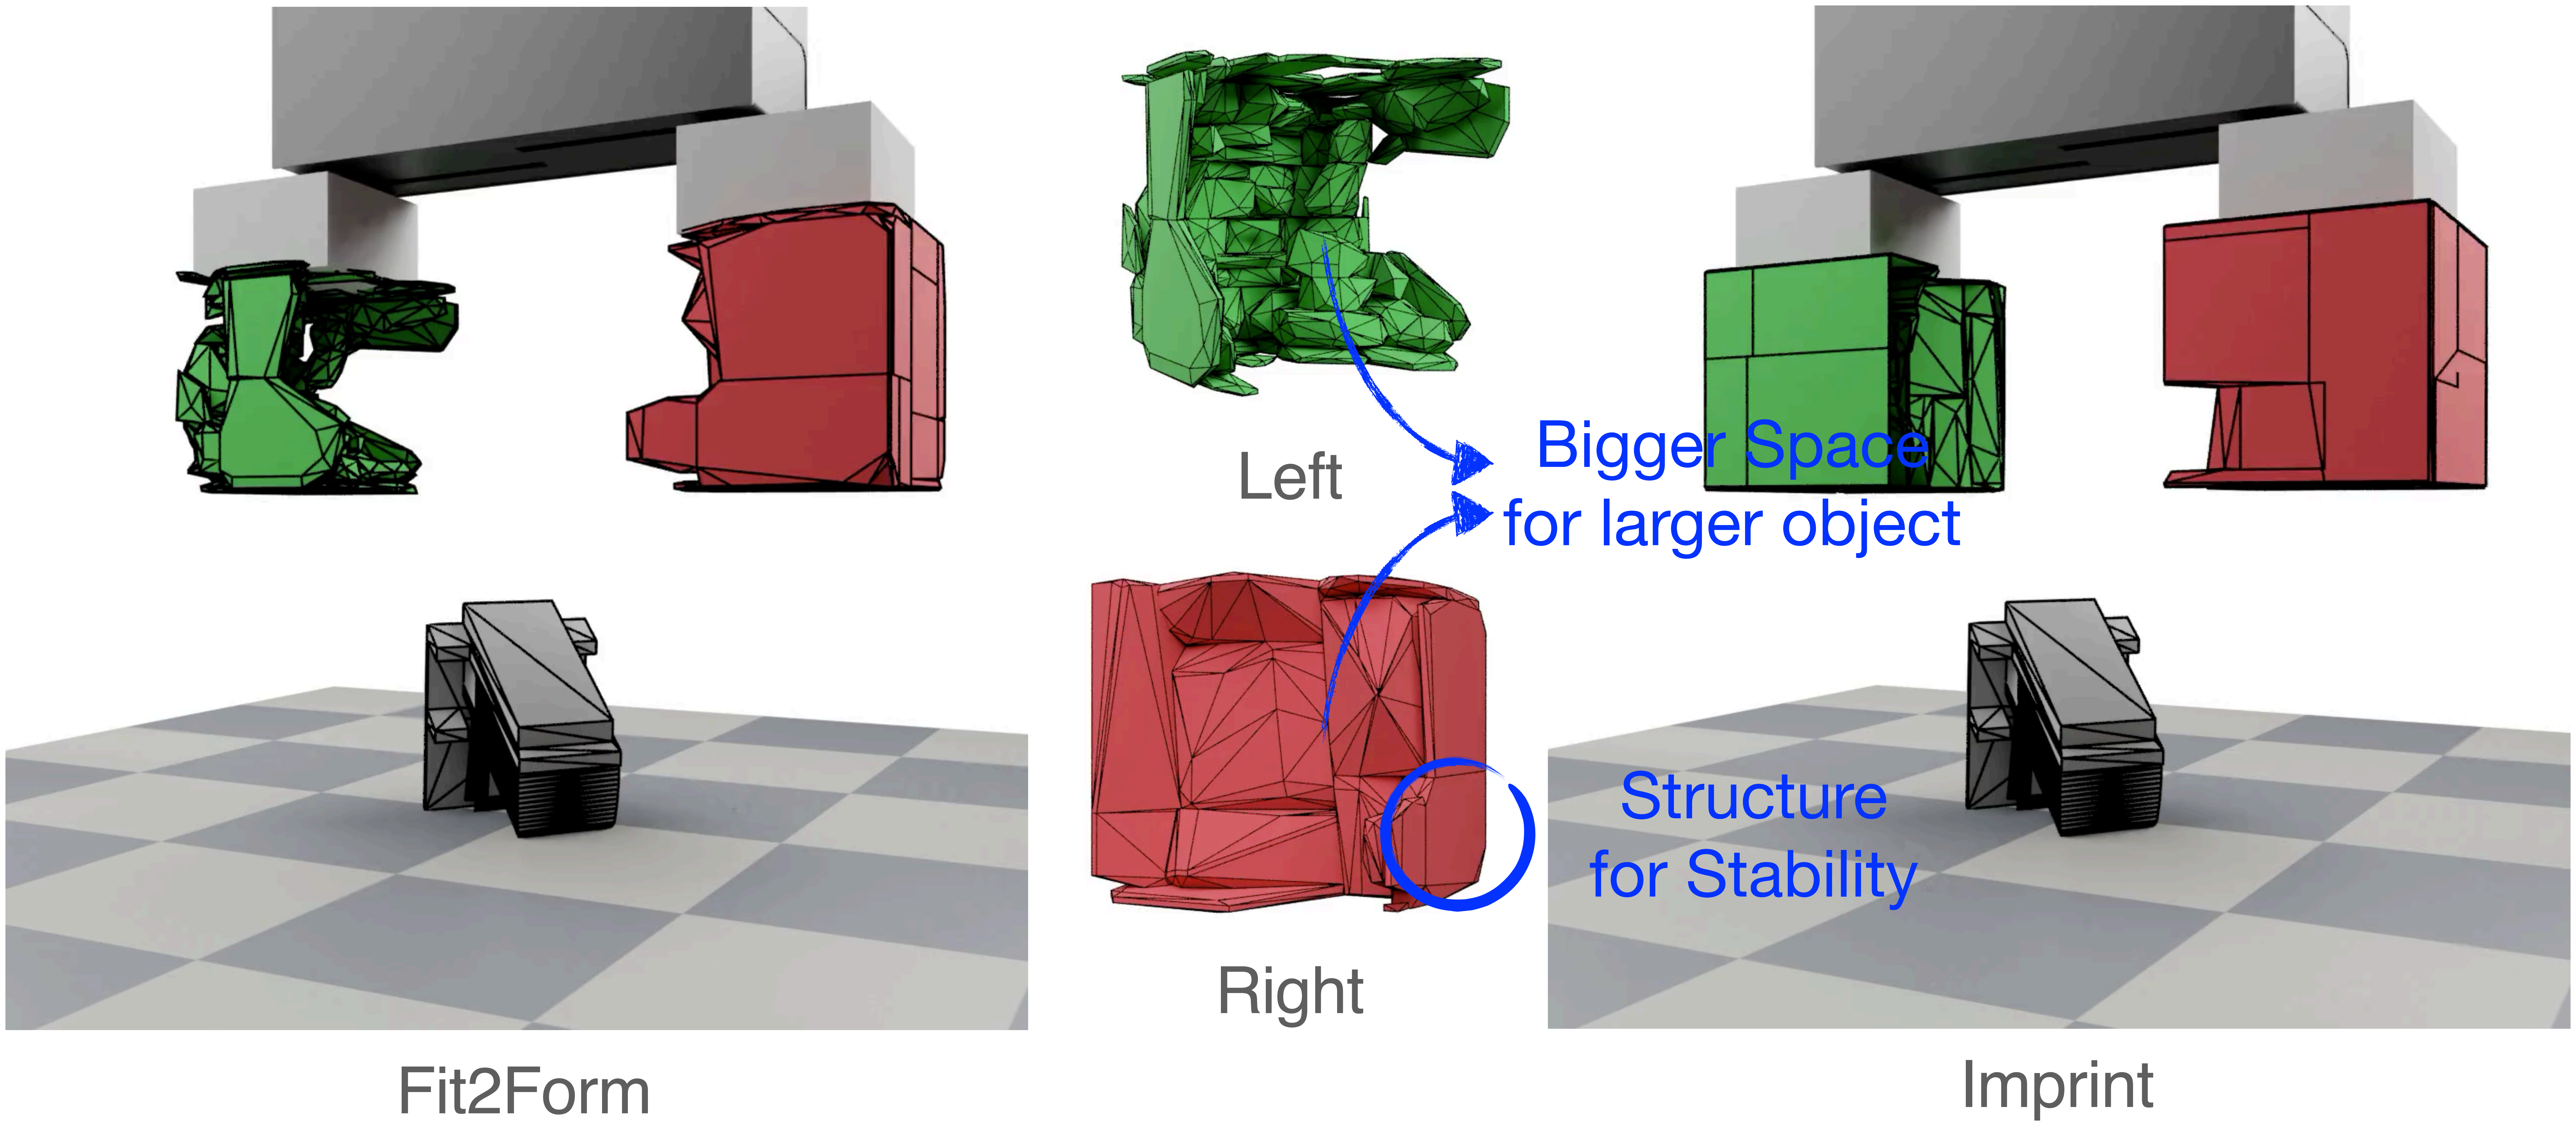

Dynamic Hook

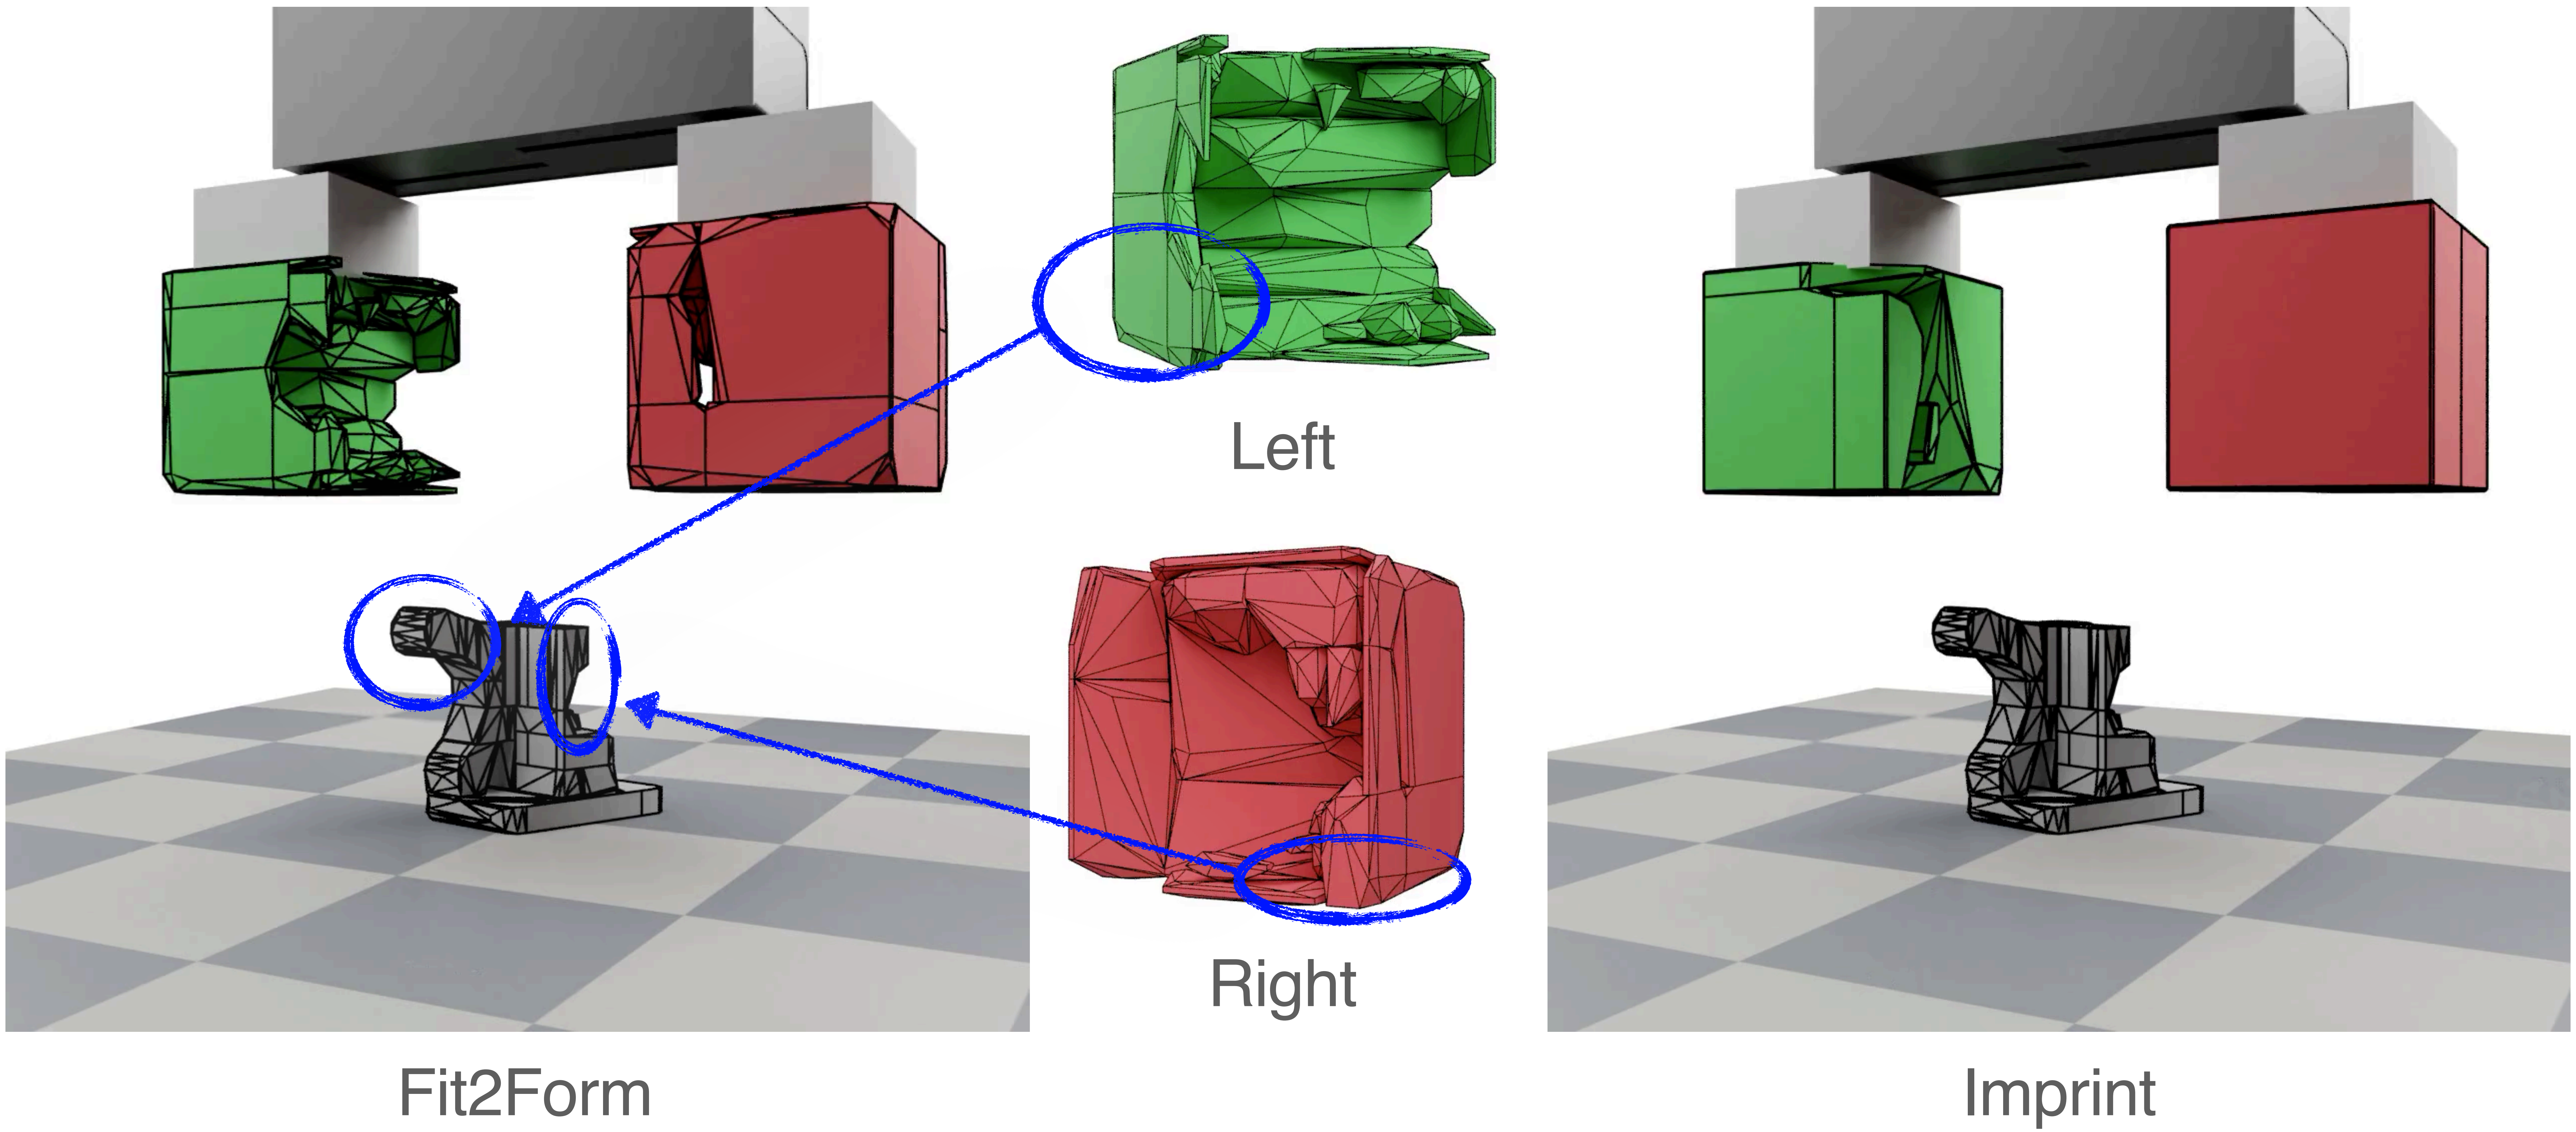

Relaxed Imprint

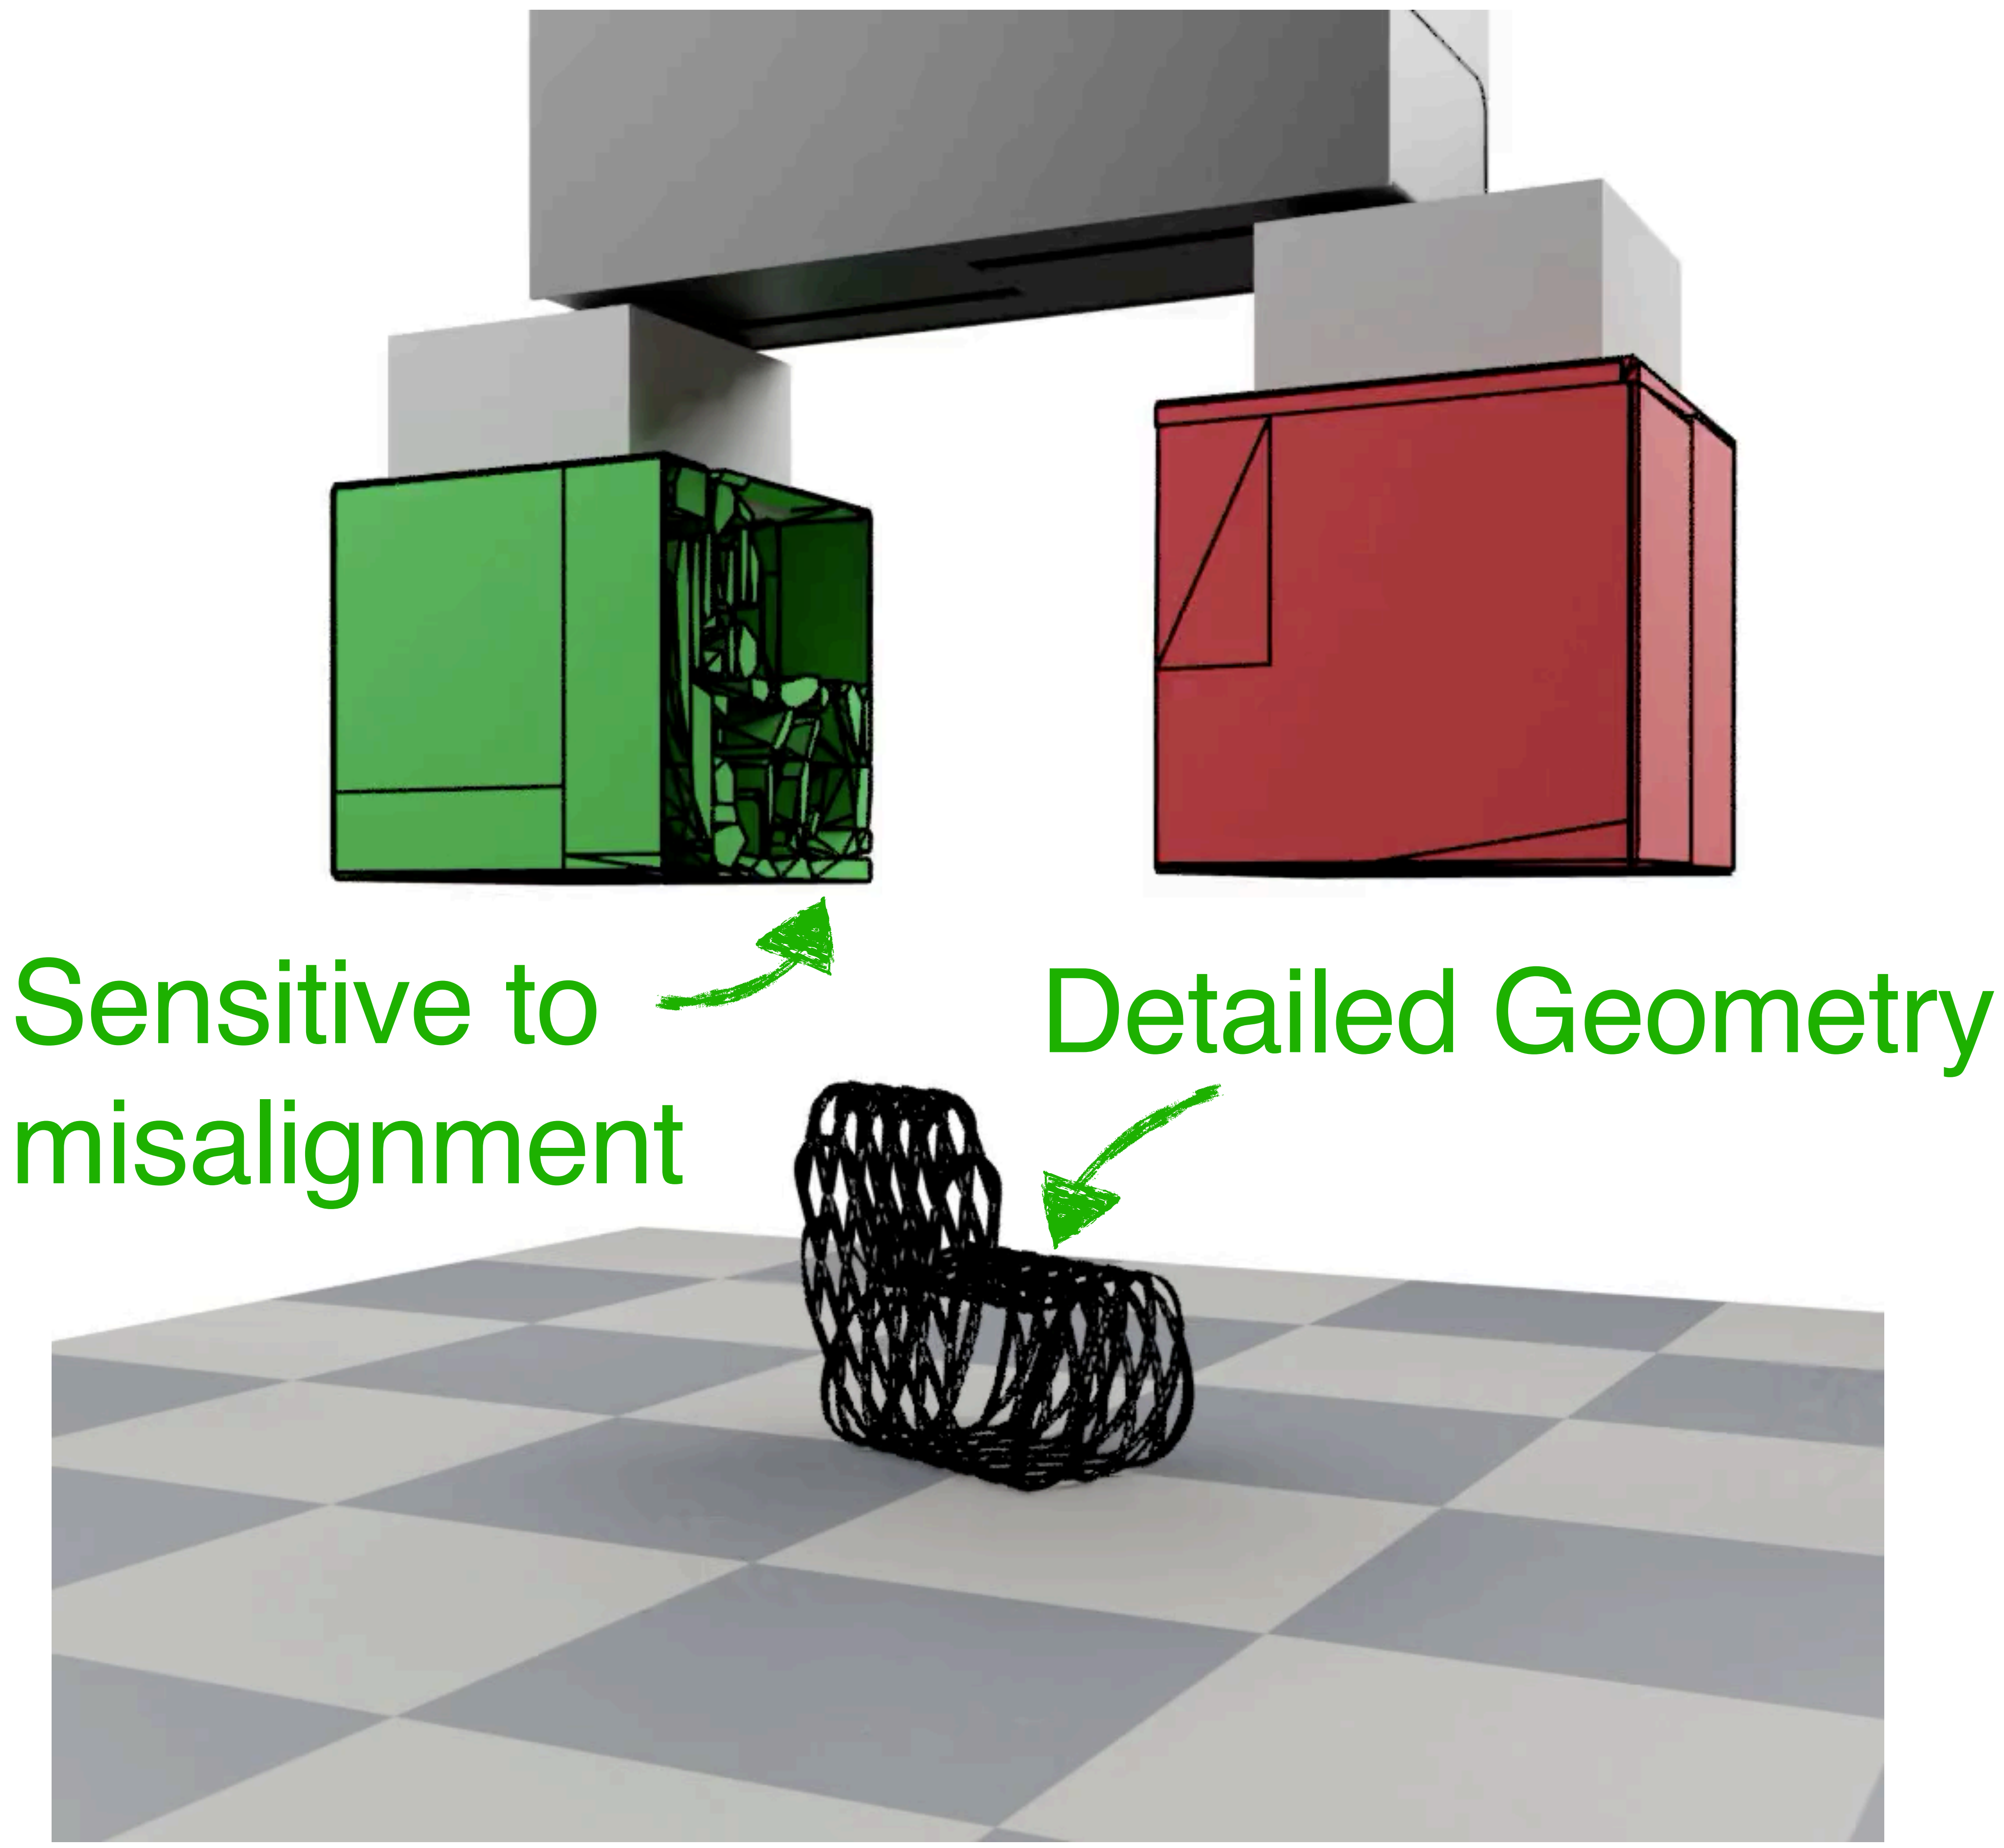

Imprint Baseline

Self-Adaptive Manipulation

Treating a robot's hardware as an integral part of robot behavior —

the learned manipulation policies should be conditioned on their hardware

also inform how hardware should be improved.

A <u>3D generative design</u> framework that uses machine learning algorithms to automate the process of robot hardware design.

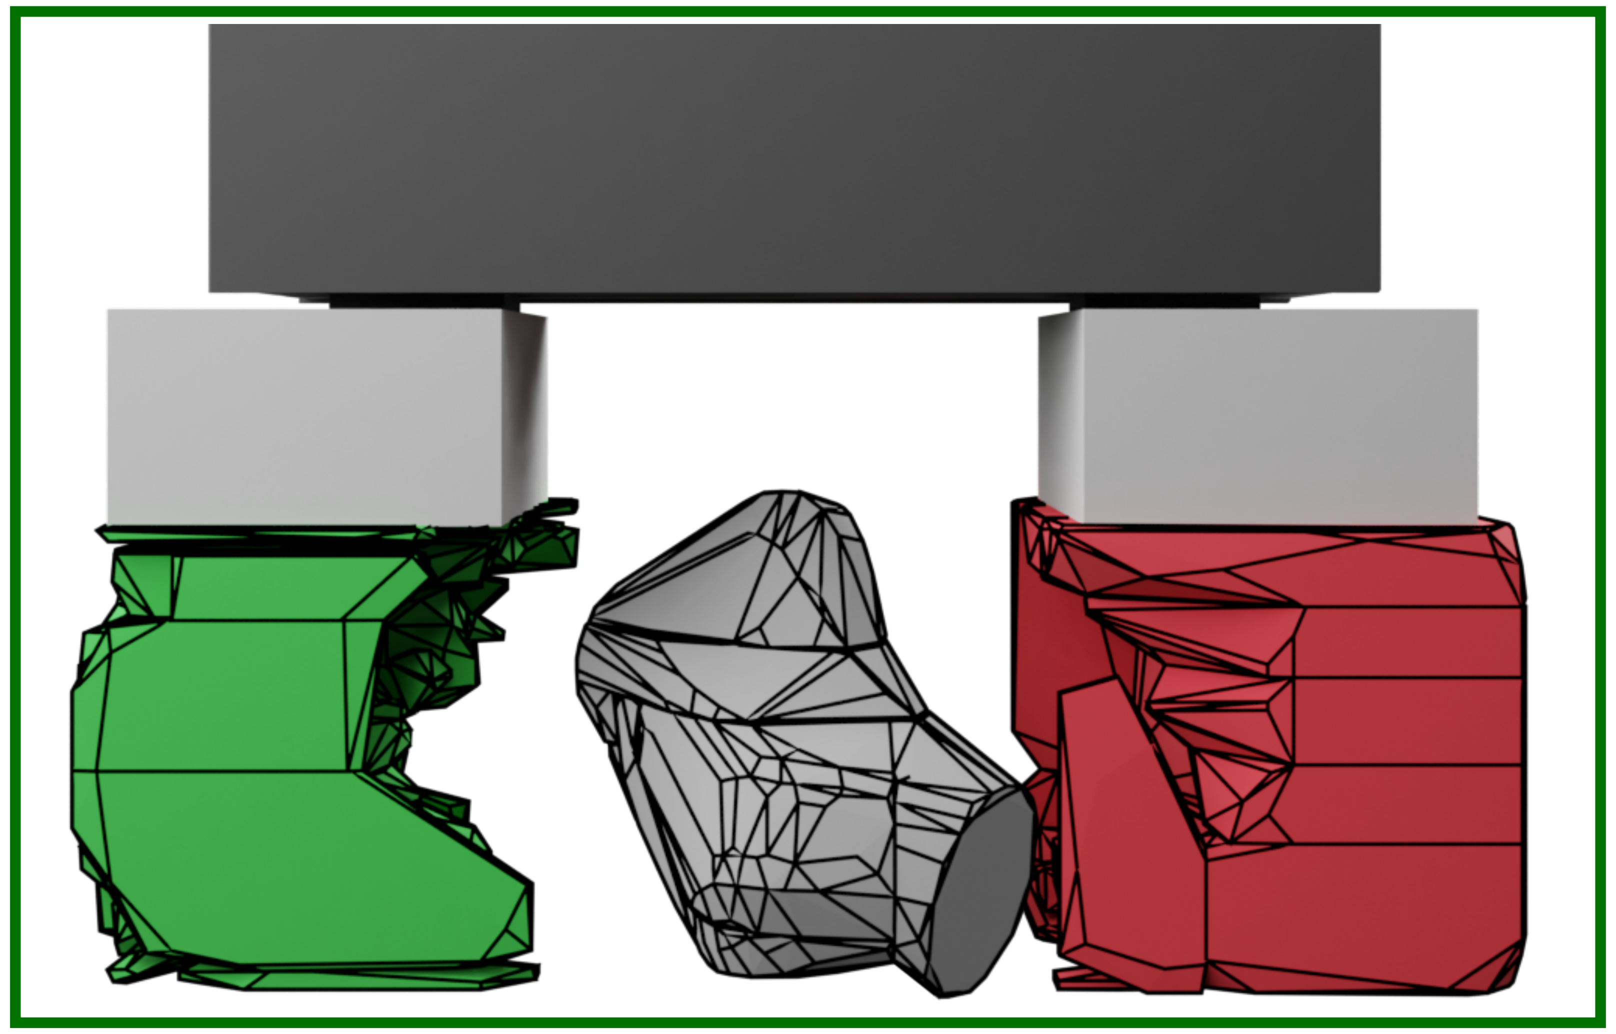

Fit2Form *(CoRL'20)*

Self-Adaptive Manipulation

Treating a robot's hardware as an integral part of robot behavior —

the learned manipulation policies should be conditioned on their hardware

also inform how hardware should be improved.

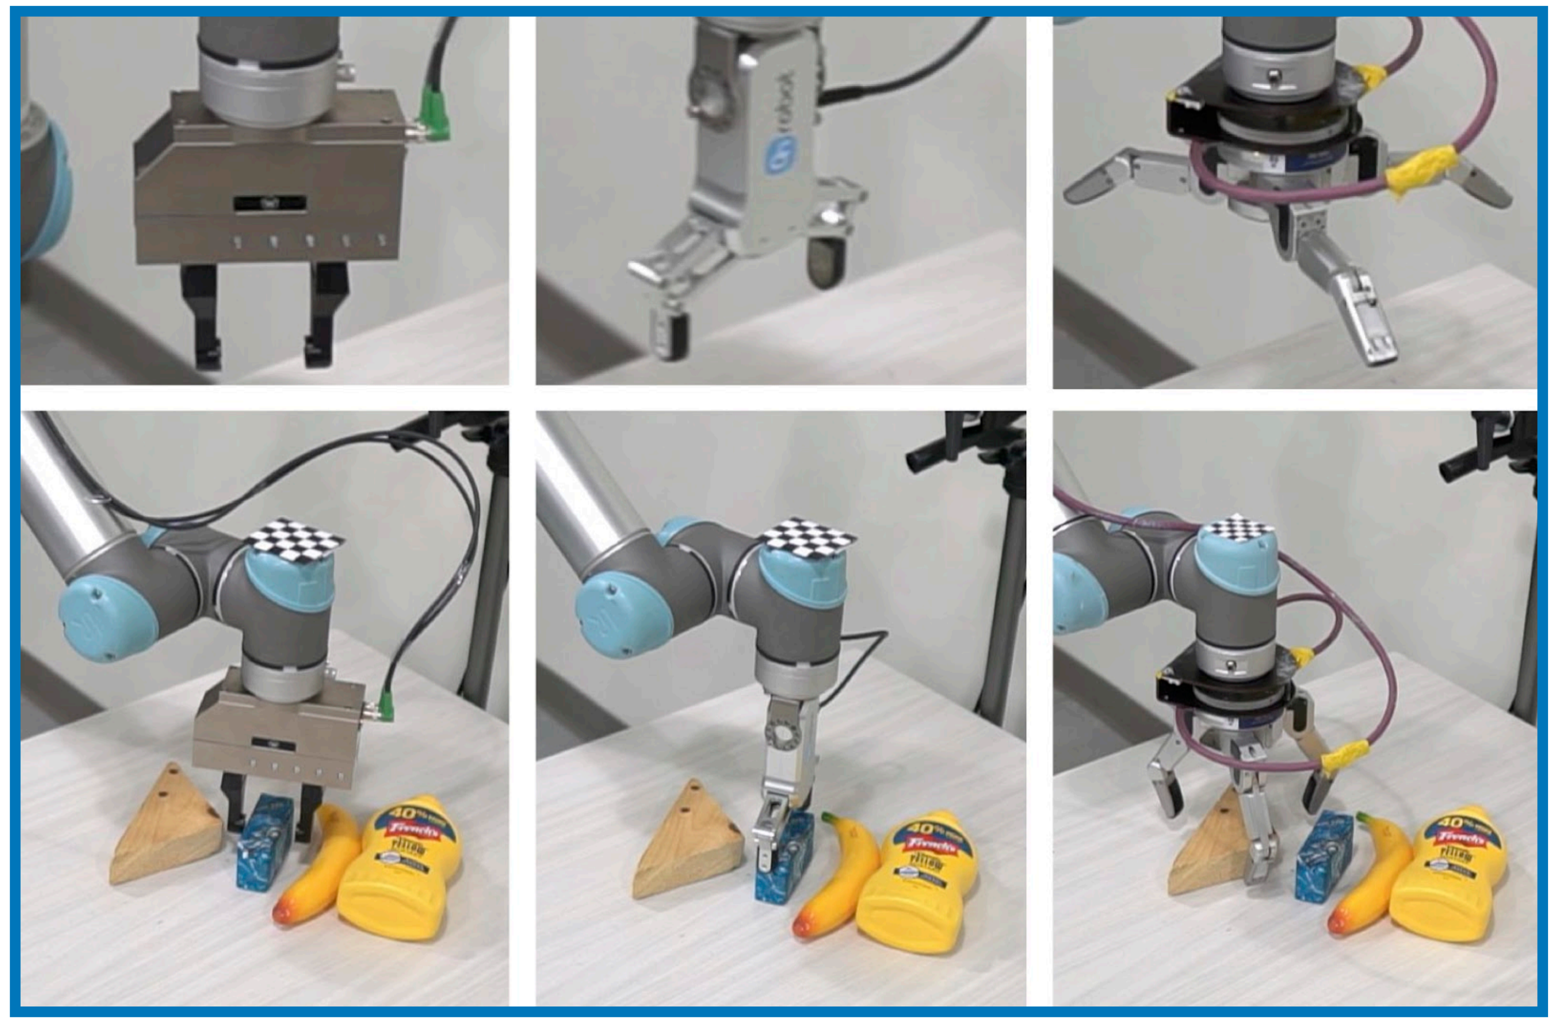

Adagrasp *(ICRA'21)*

https://adagrasp.cs.columbia.edu/

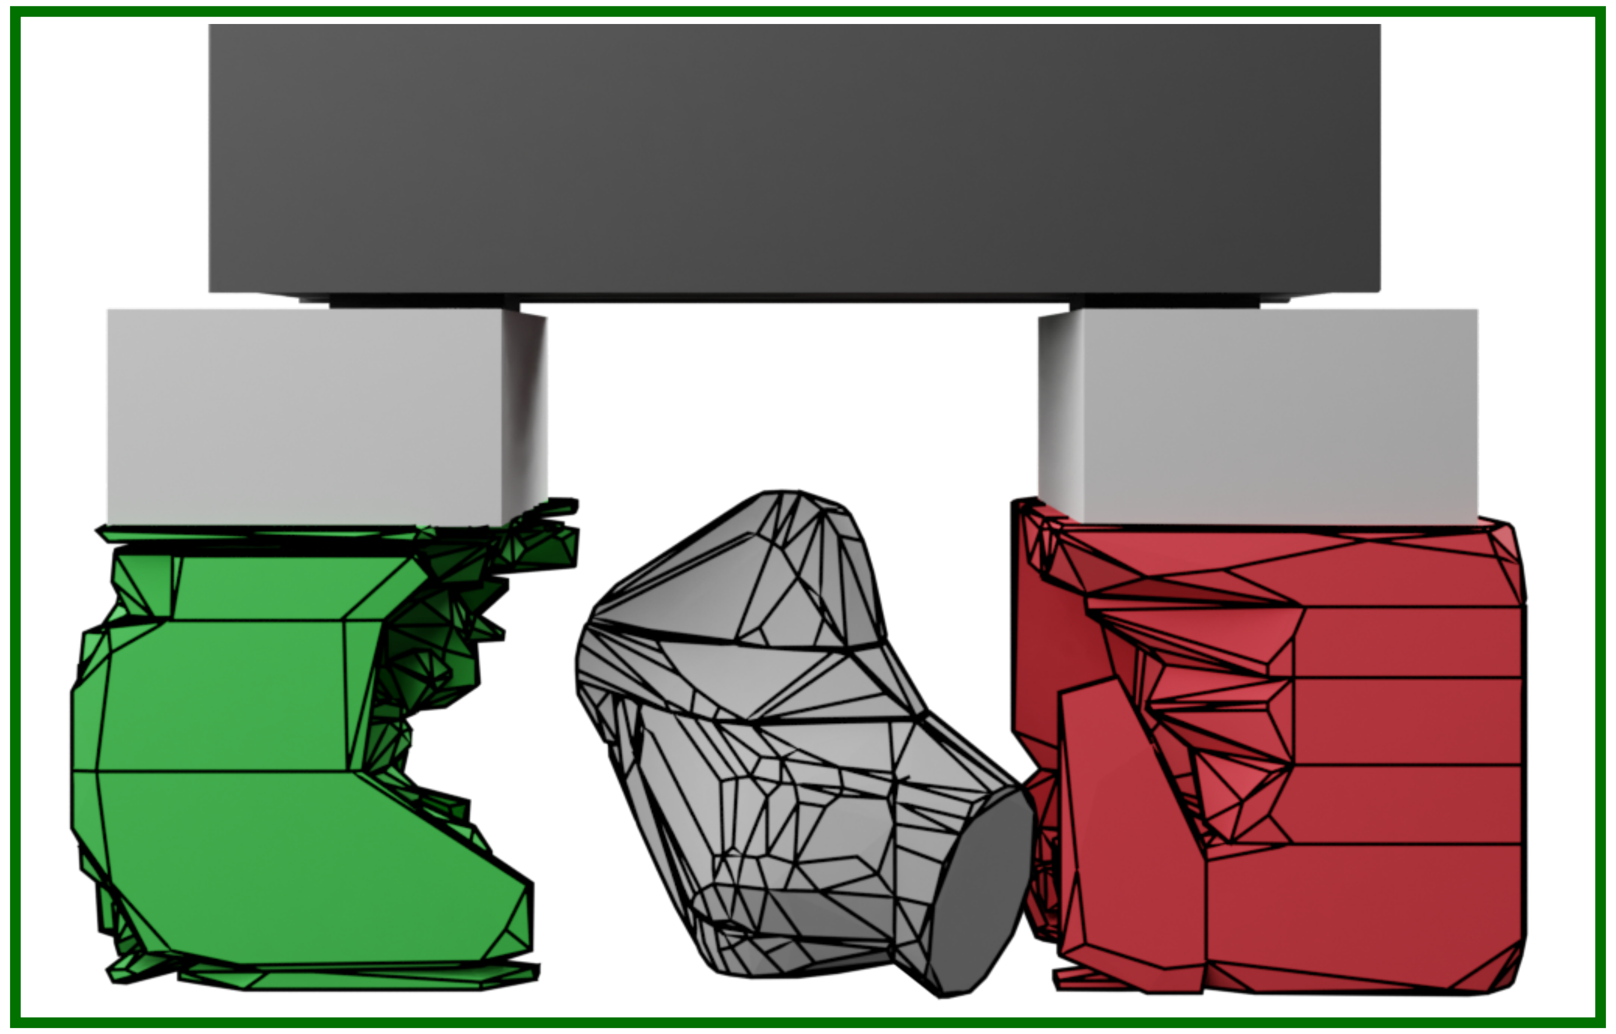

Fit2Form *(CoRL'20)*

https://fit2form.cs.columbia.edu/

Self-Adaptive Manipulation

What's Missing?

Self-Adaptive Manipulation

What's Missing?

Whole-body Motion Planning

Self-Adaptive Manipulation

What's Missing?

Type

Tactile*

Force

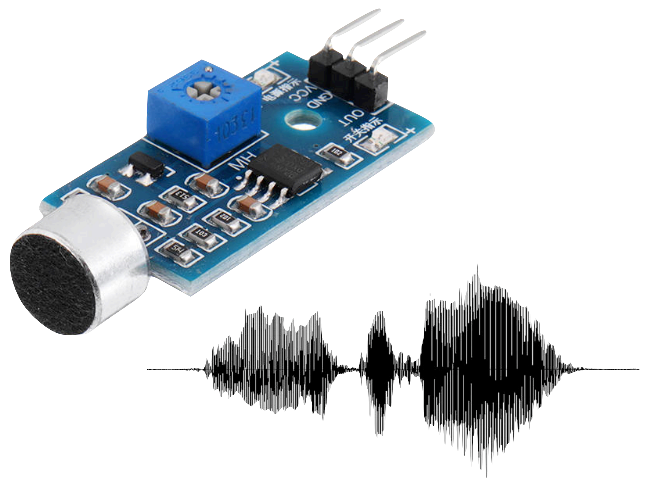

Sound

Sensor Input and Placement

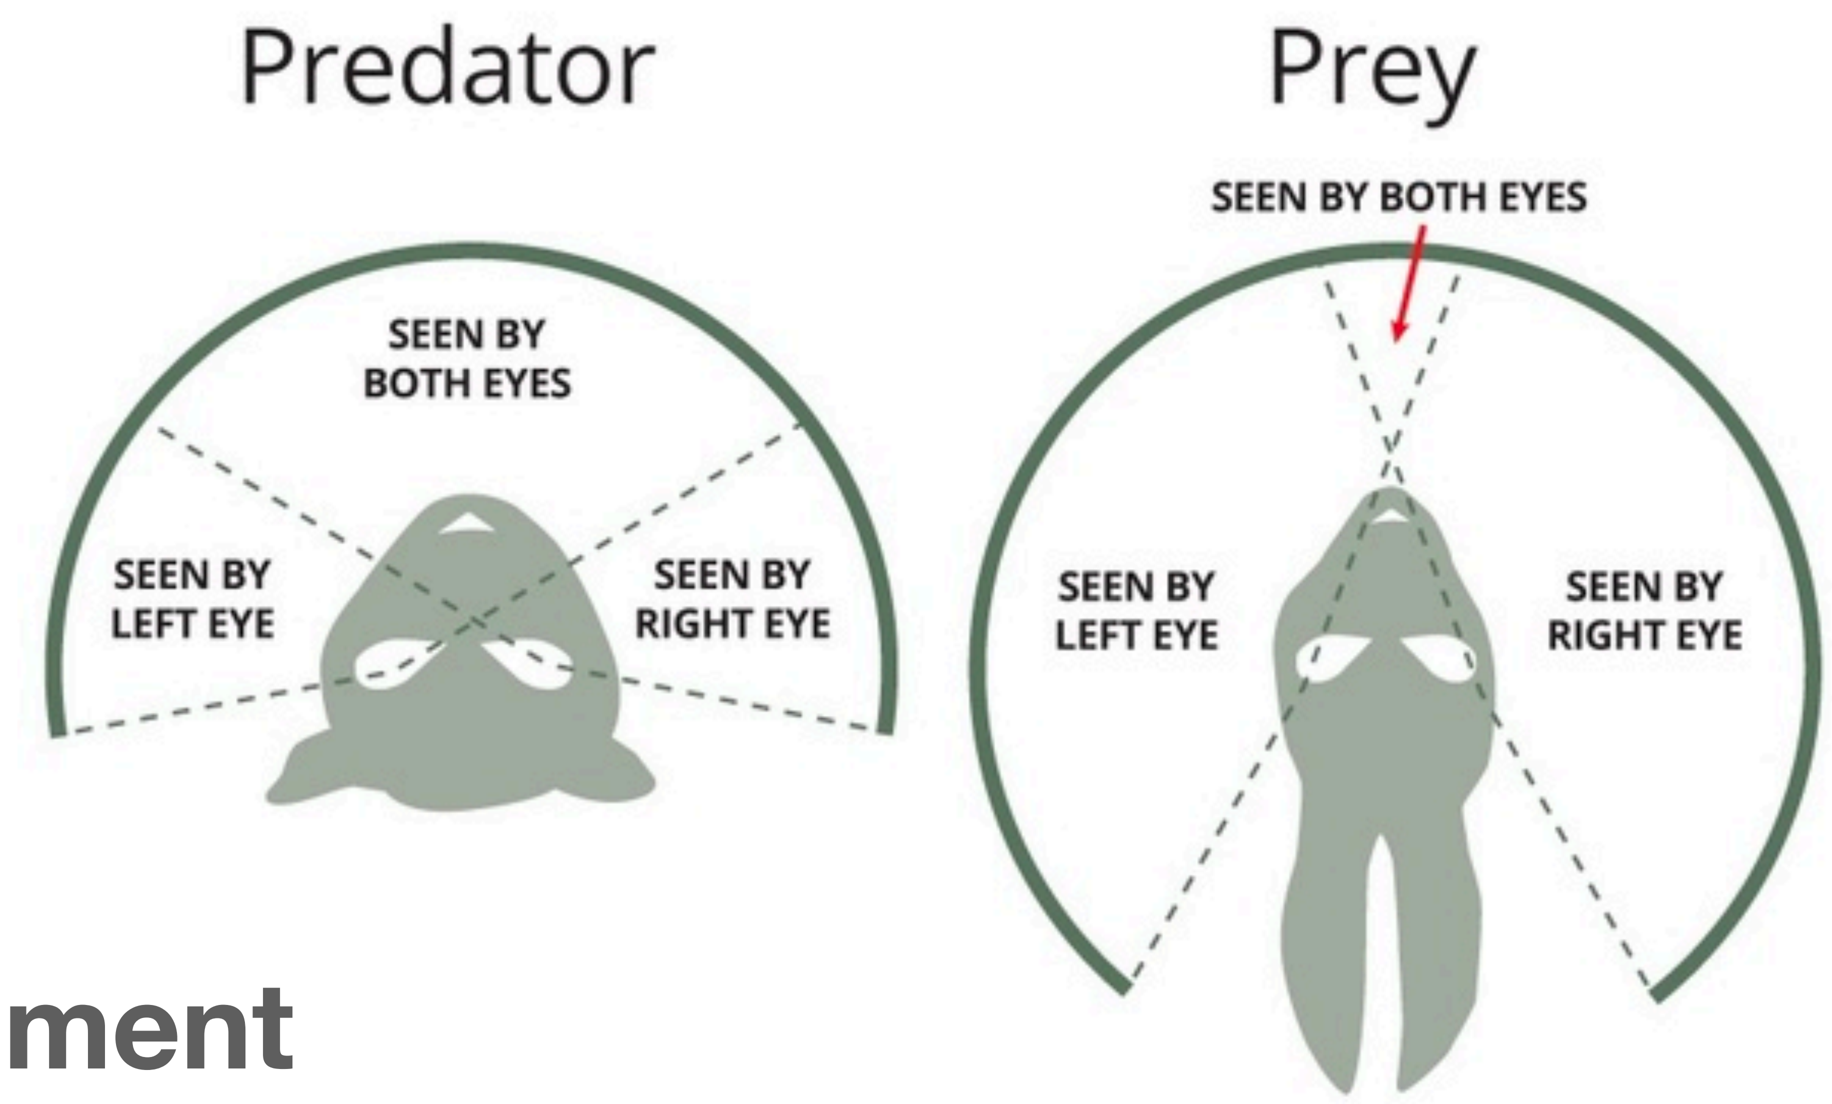

Placement

Self-Adaptive Manipulation

Geometry → **Other physical properties**
(deformation, contact mechanisms)

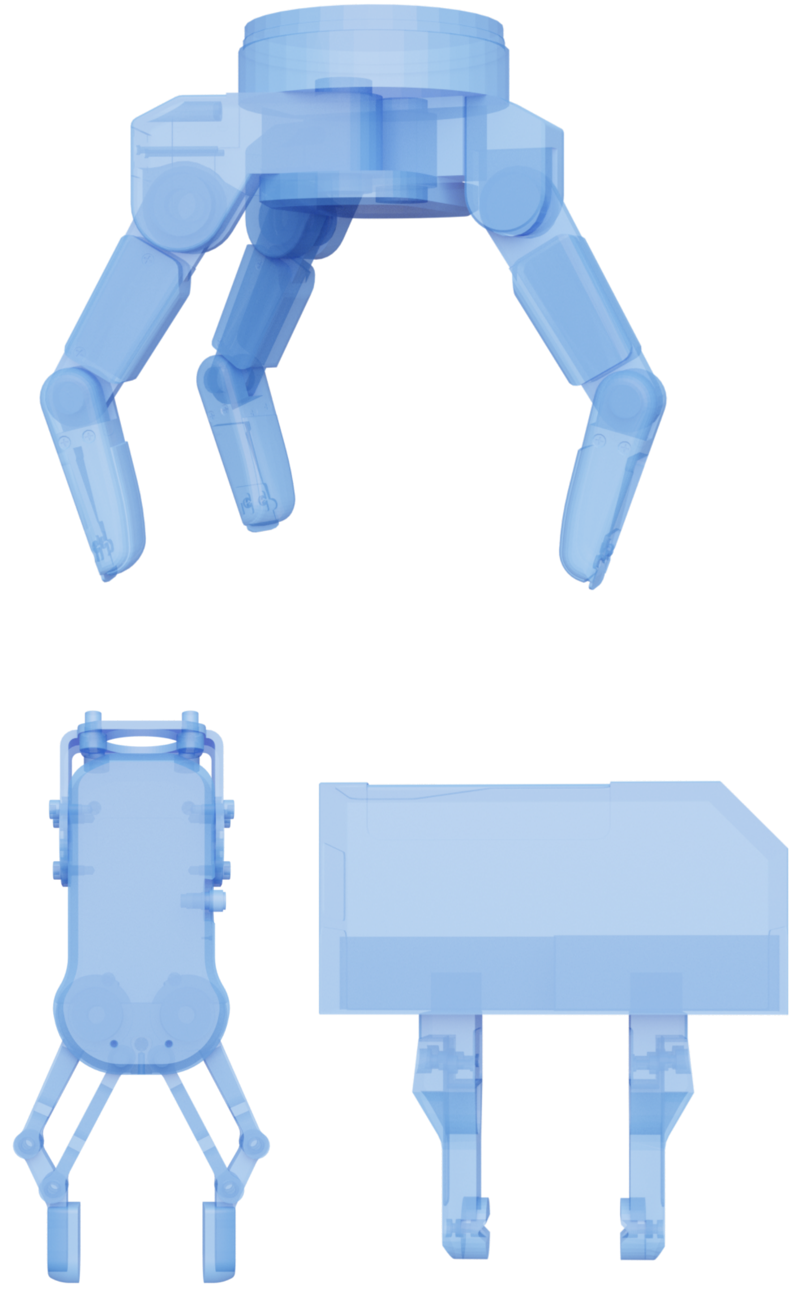

Embodied Intelligence

Fixed Hardware

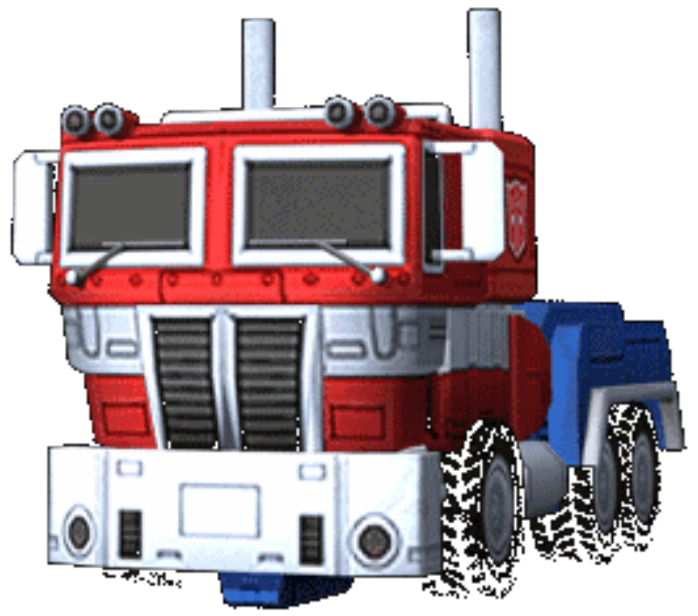

Hardware as an integral part of robot behavior

✓ **Adaptable**
Environments

✓ **Resilient**
Potential Damages

✓ **Versatile**
Tasks

Acknowledgment

[1] AdaGrasp: Learning an Adaptive Gripper-Aware Grasping Policy, ICRA 2021

[2] Fit2Form: 3D Generative Model for Robot Gripper Form Design. Conference on Robotic Learning (CoRL) 2020

Zhenjia Xu

Beichun Qi

Shubham Agrawal

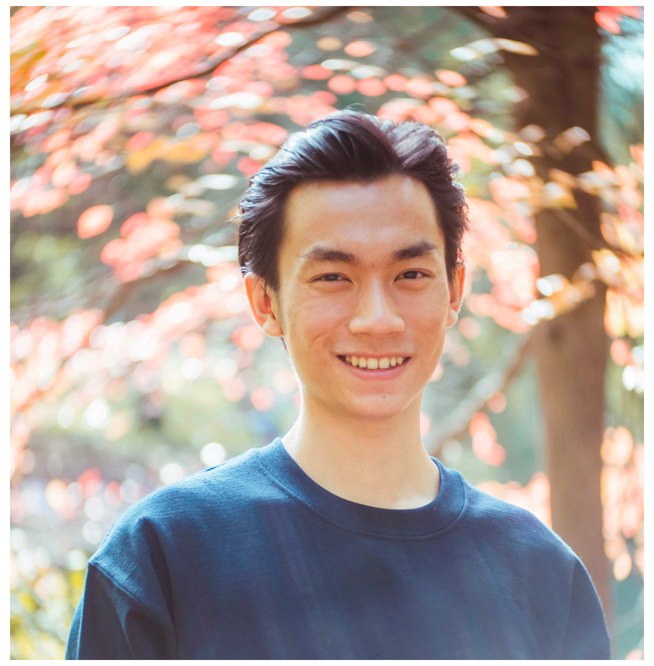

Huy Ha