UC San Diego

SAPIEN Manipulation Skill Challenge: a New Arena for Embodied AI

Hao Su (苏昊)

Embodied AI (具身人工智能)

Intangible AI:

- AI without a body
- Watches without acting

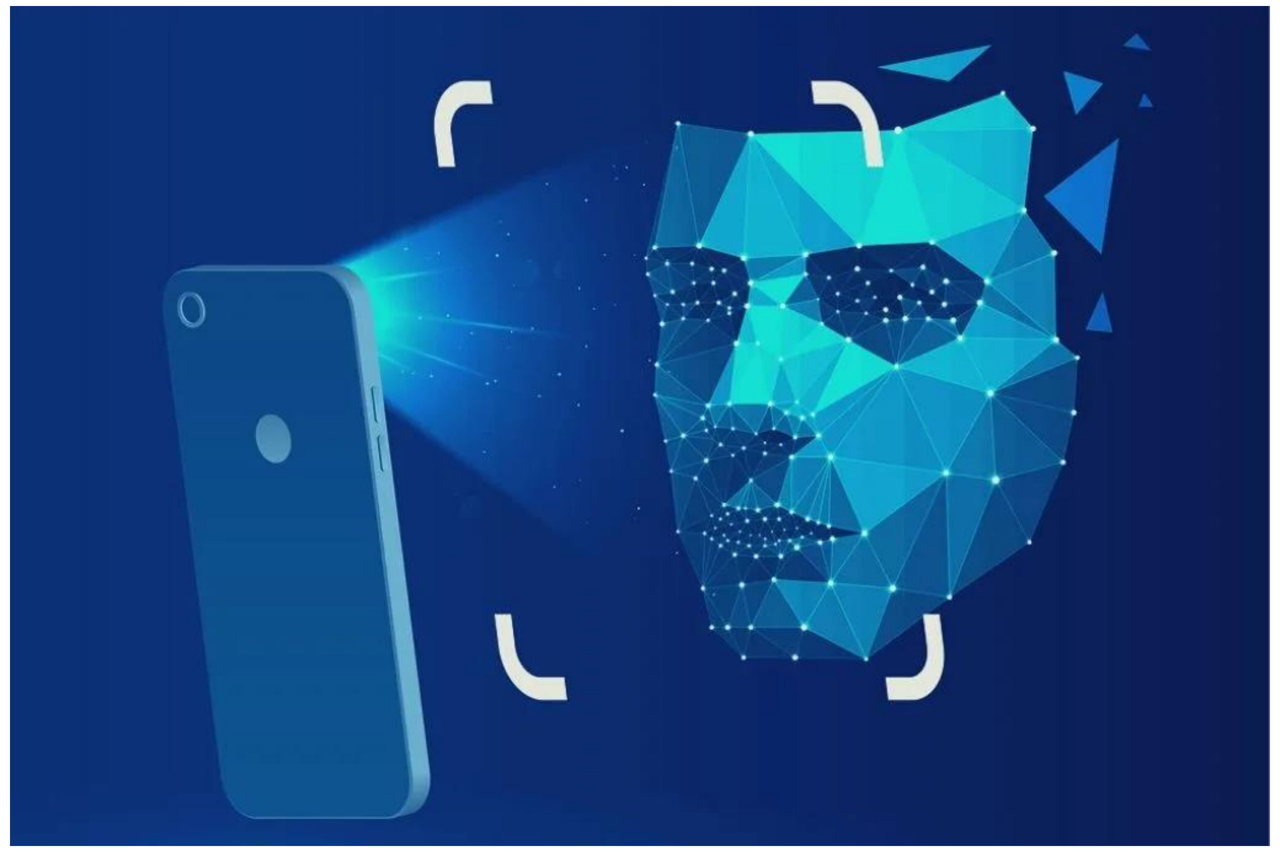

Embodied AI:

- AI with body
- **Interact** with the physical world

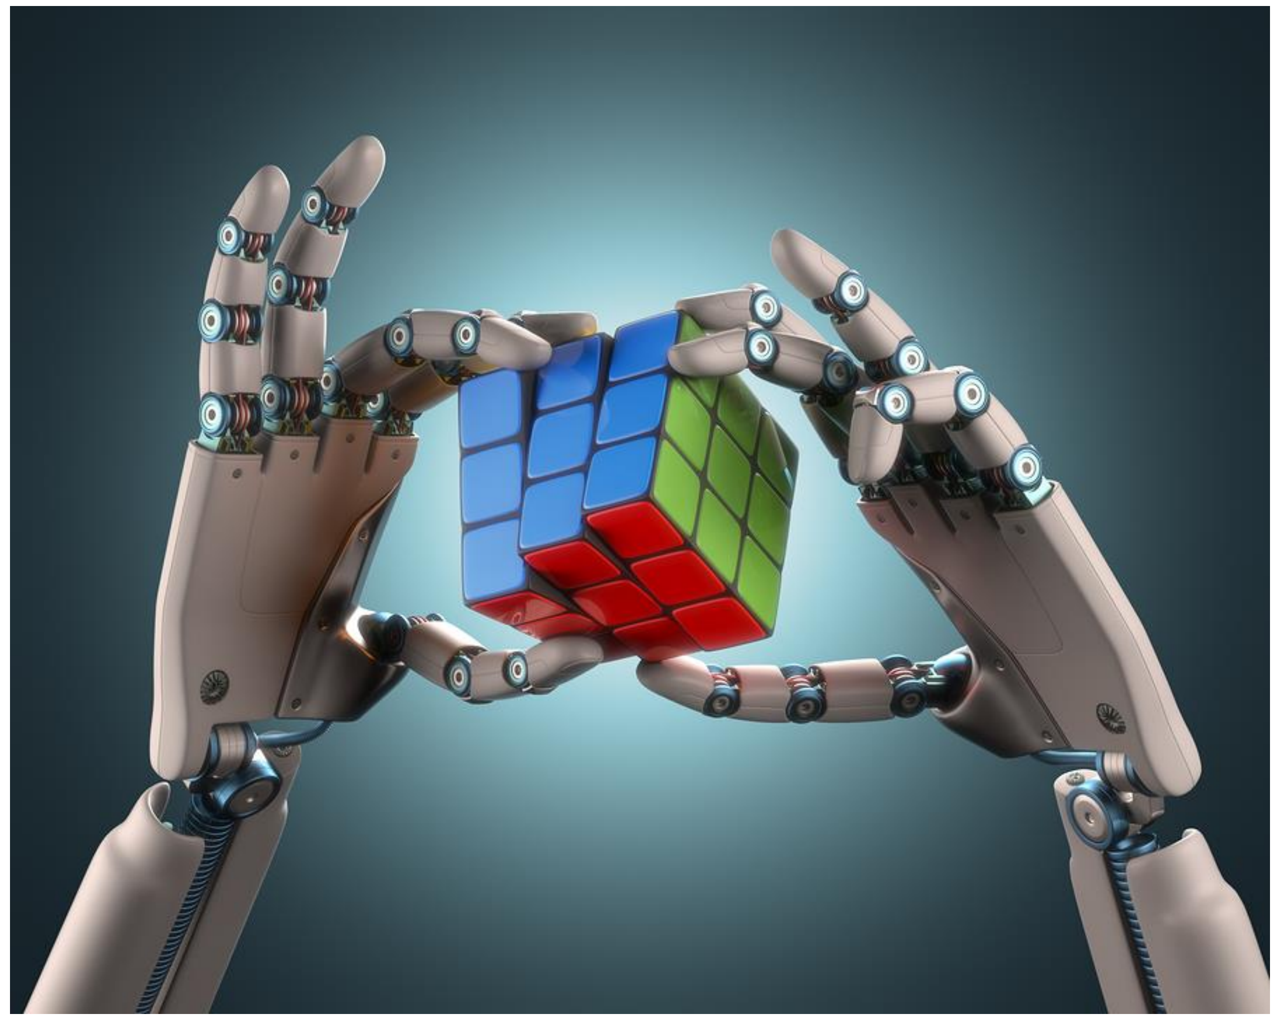

Exemple Tasks of Embodied AI

Autonomous driving[1]

Home service[2]

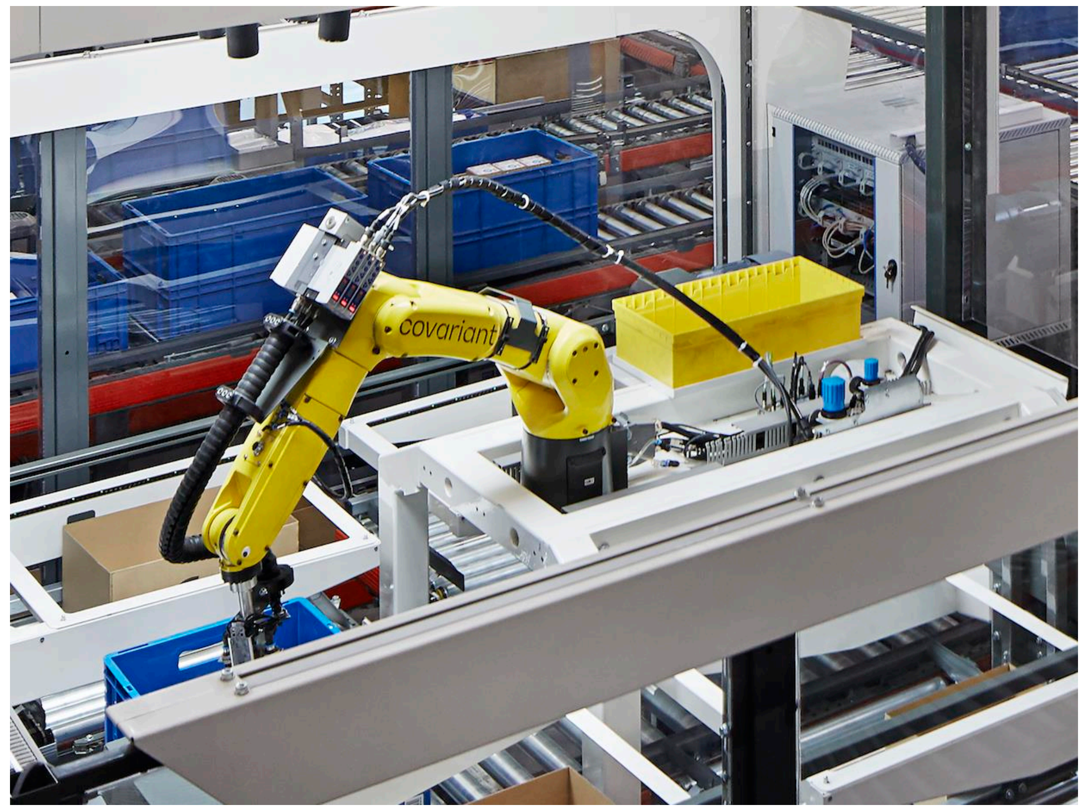

Smart manufacturing[3]

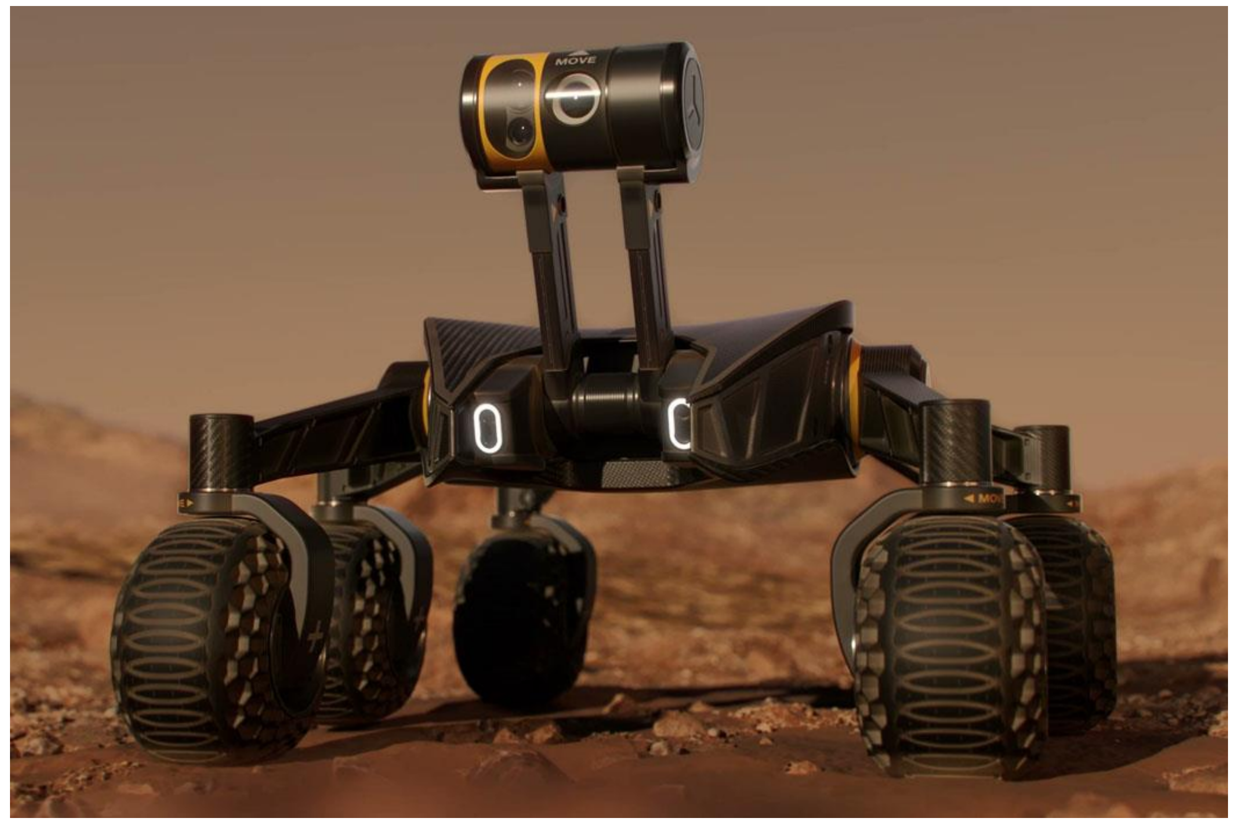

Explorer[4]

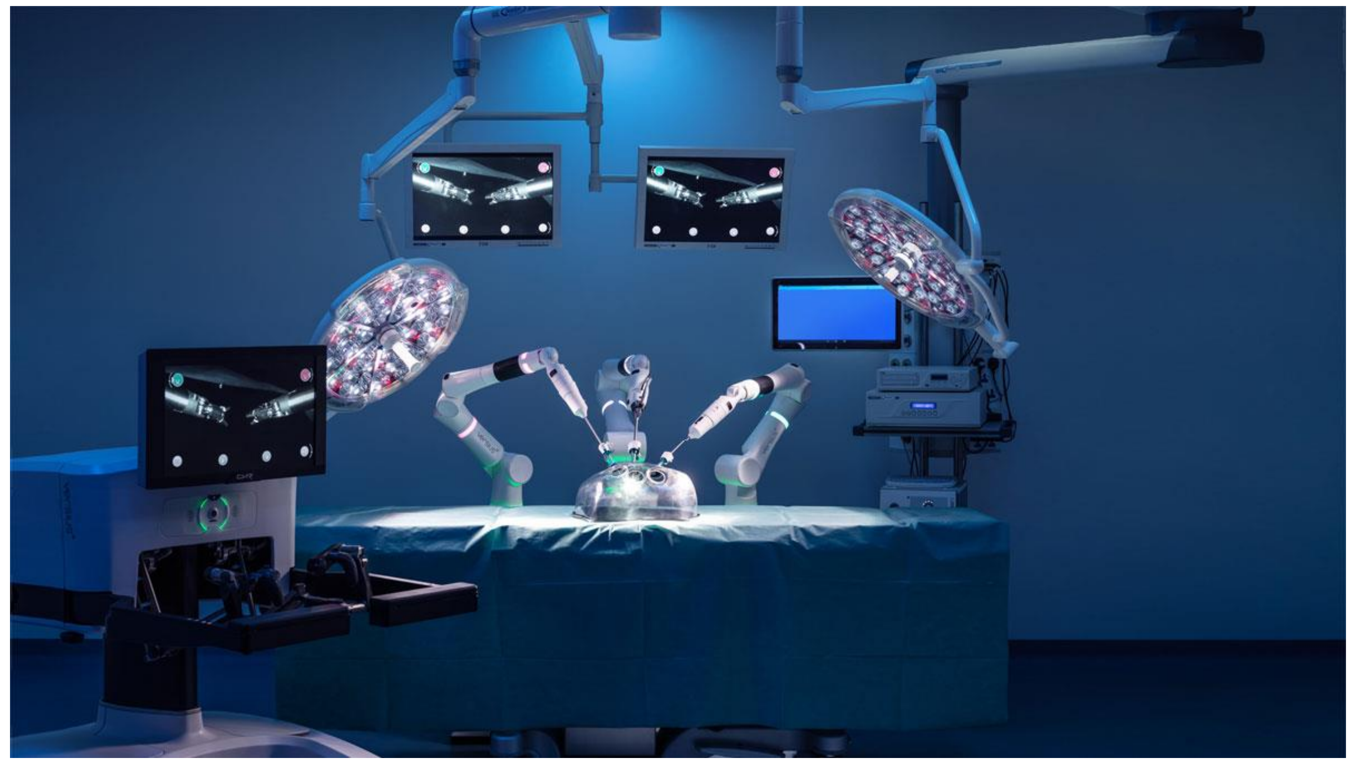

Healthcare[5]

1. Image from adobe stock
2. https://fuentitech.com/facebook-wants-to-help-train-robots-to-put-out-trash-and-unload-the-dishwasher/108043/
3. https://spectrum.ieee.org/automaton/robotics/industrial-robots/covariant-ai-gigantic-neural-network-to-automate-warehouse-picking
4. https://www.yankodesign.com/2020/11/21/this-ai-enabled-mars-exploration-rover-is-as-adorable-as-wall-e/
5. https://medtech.pharmaintelligence.informa.com/MT123233/CMR-Surgical-Launches-Robotic-System-To-Rival-Intuitives-Da-Vinci

Why Embodied AI?

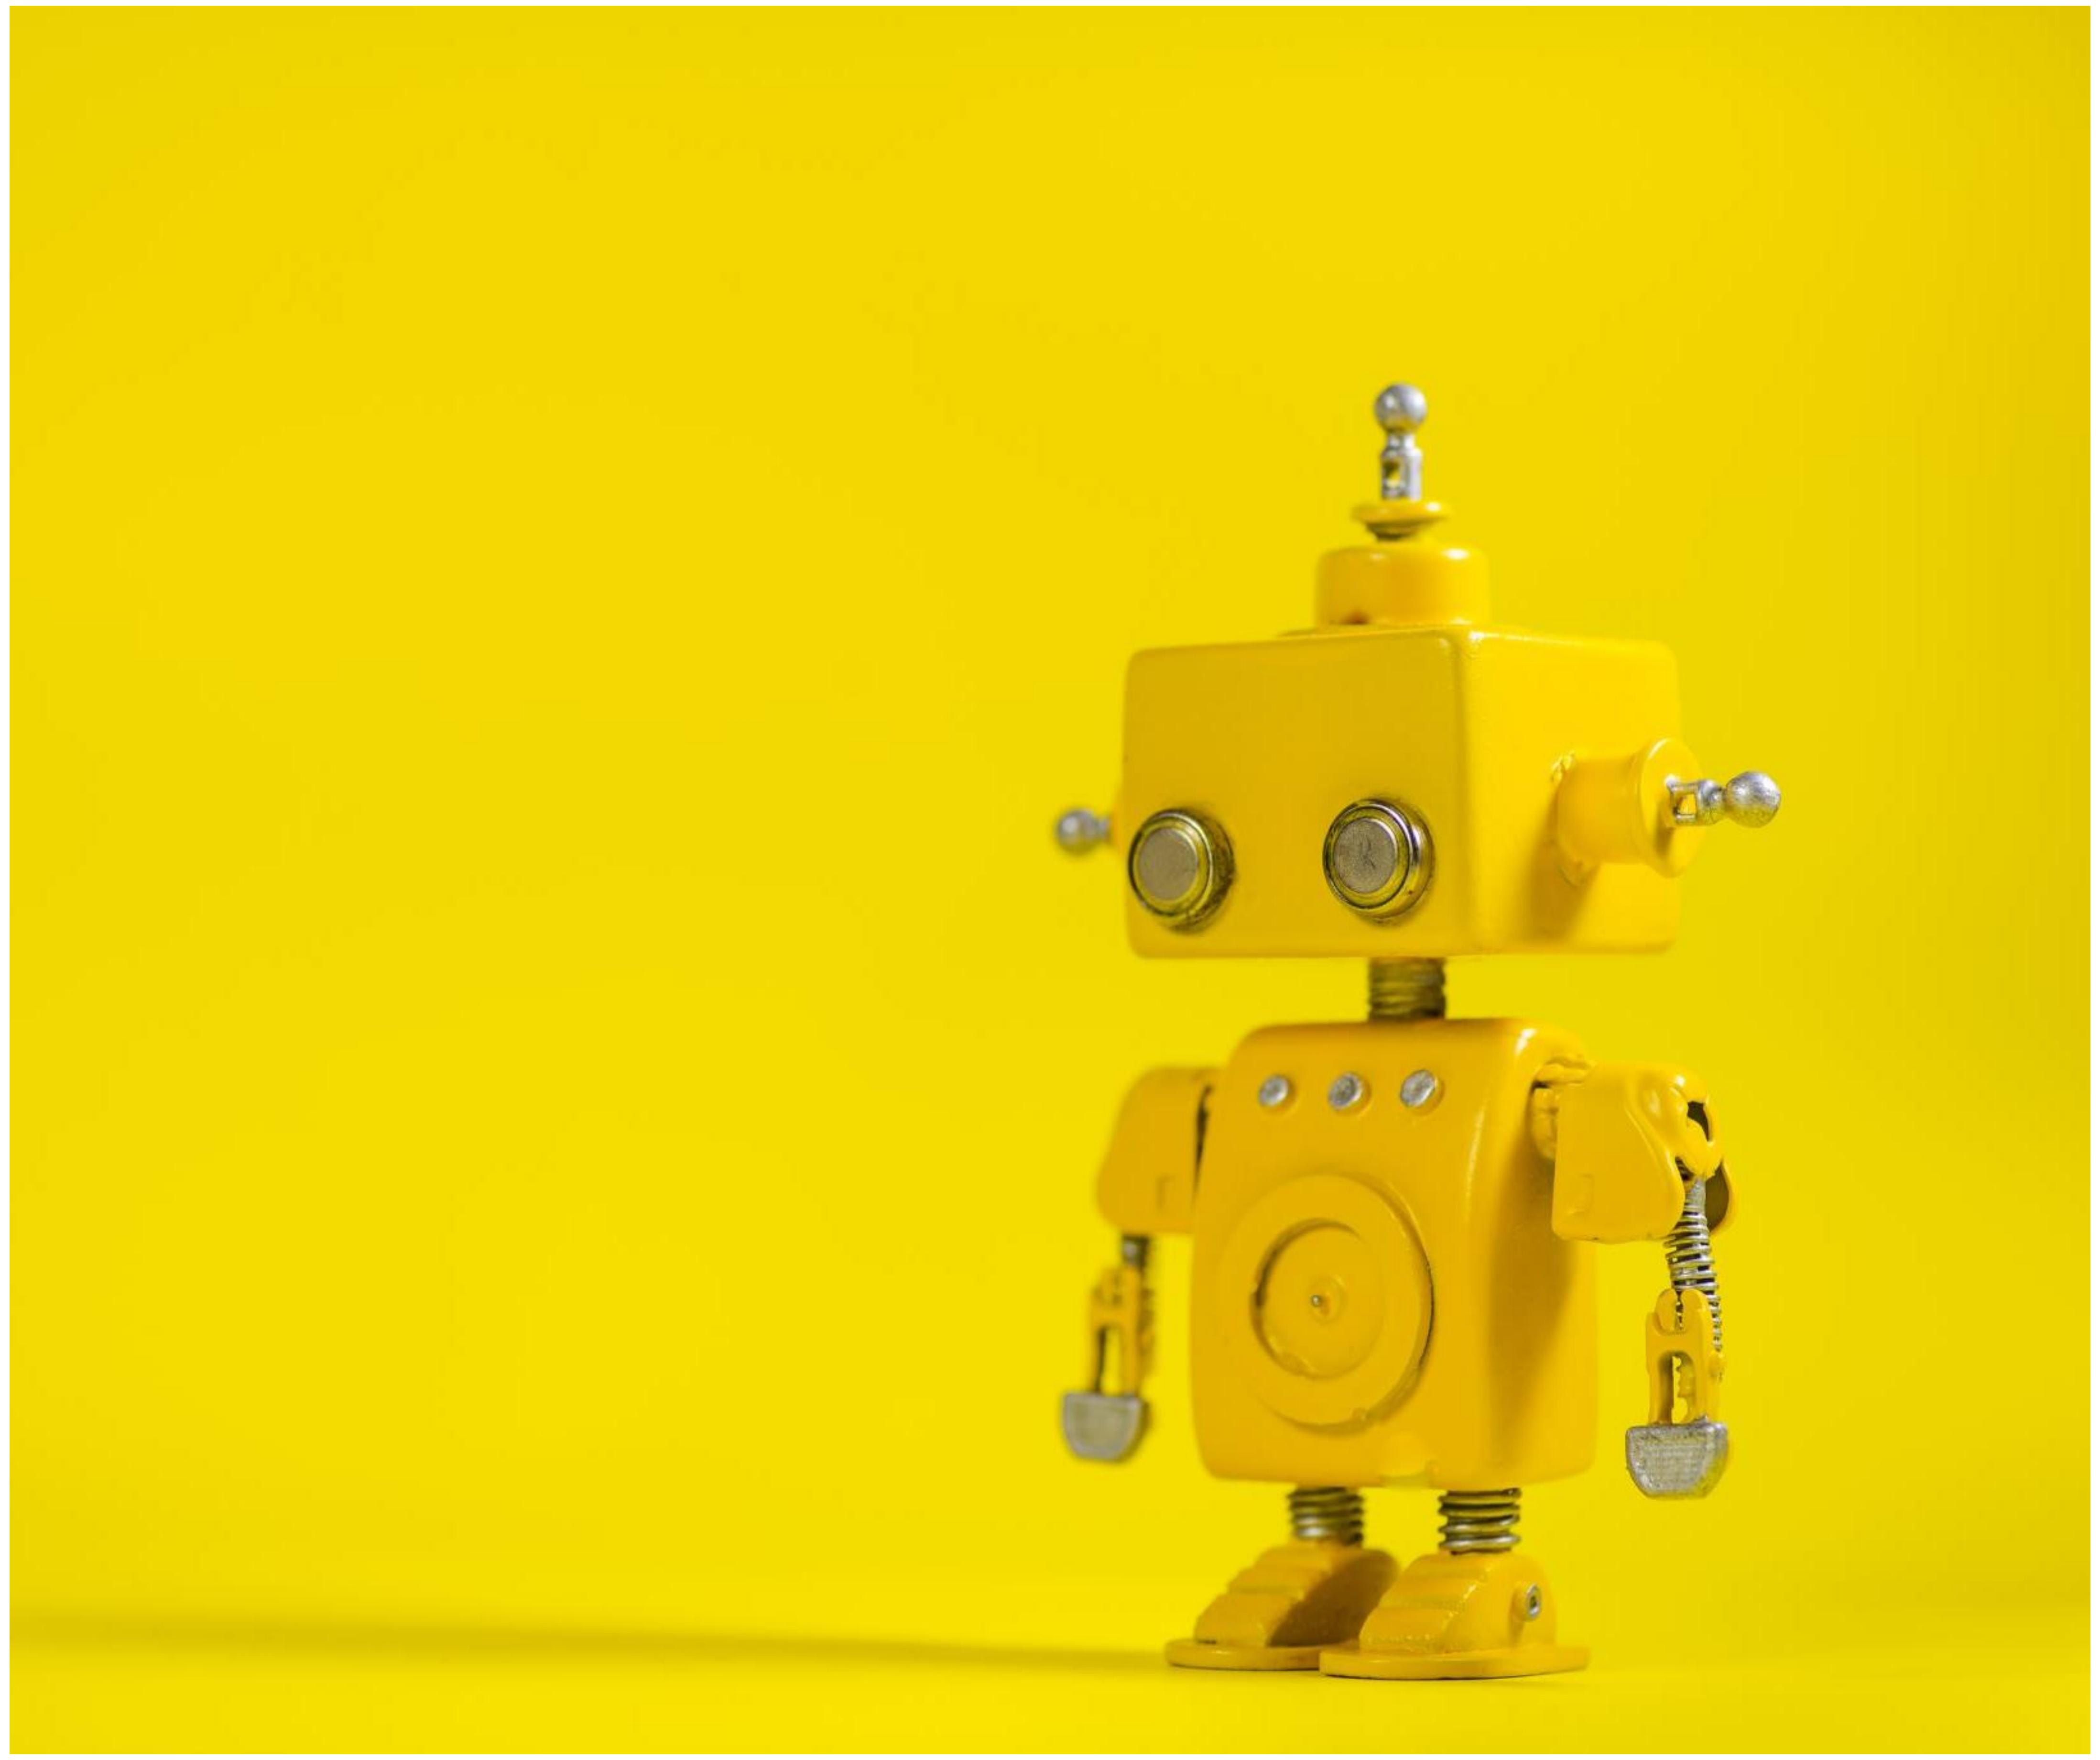

Case Study: Learning to play football

Image source: https://www.thesoccerstore.co.uk/blog/football-goals/best-kids-garden-football-goals/

Case Study: Learning to play football

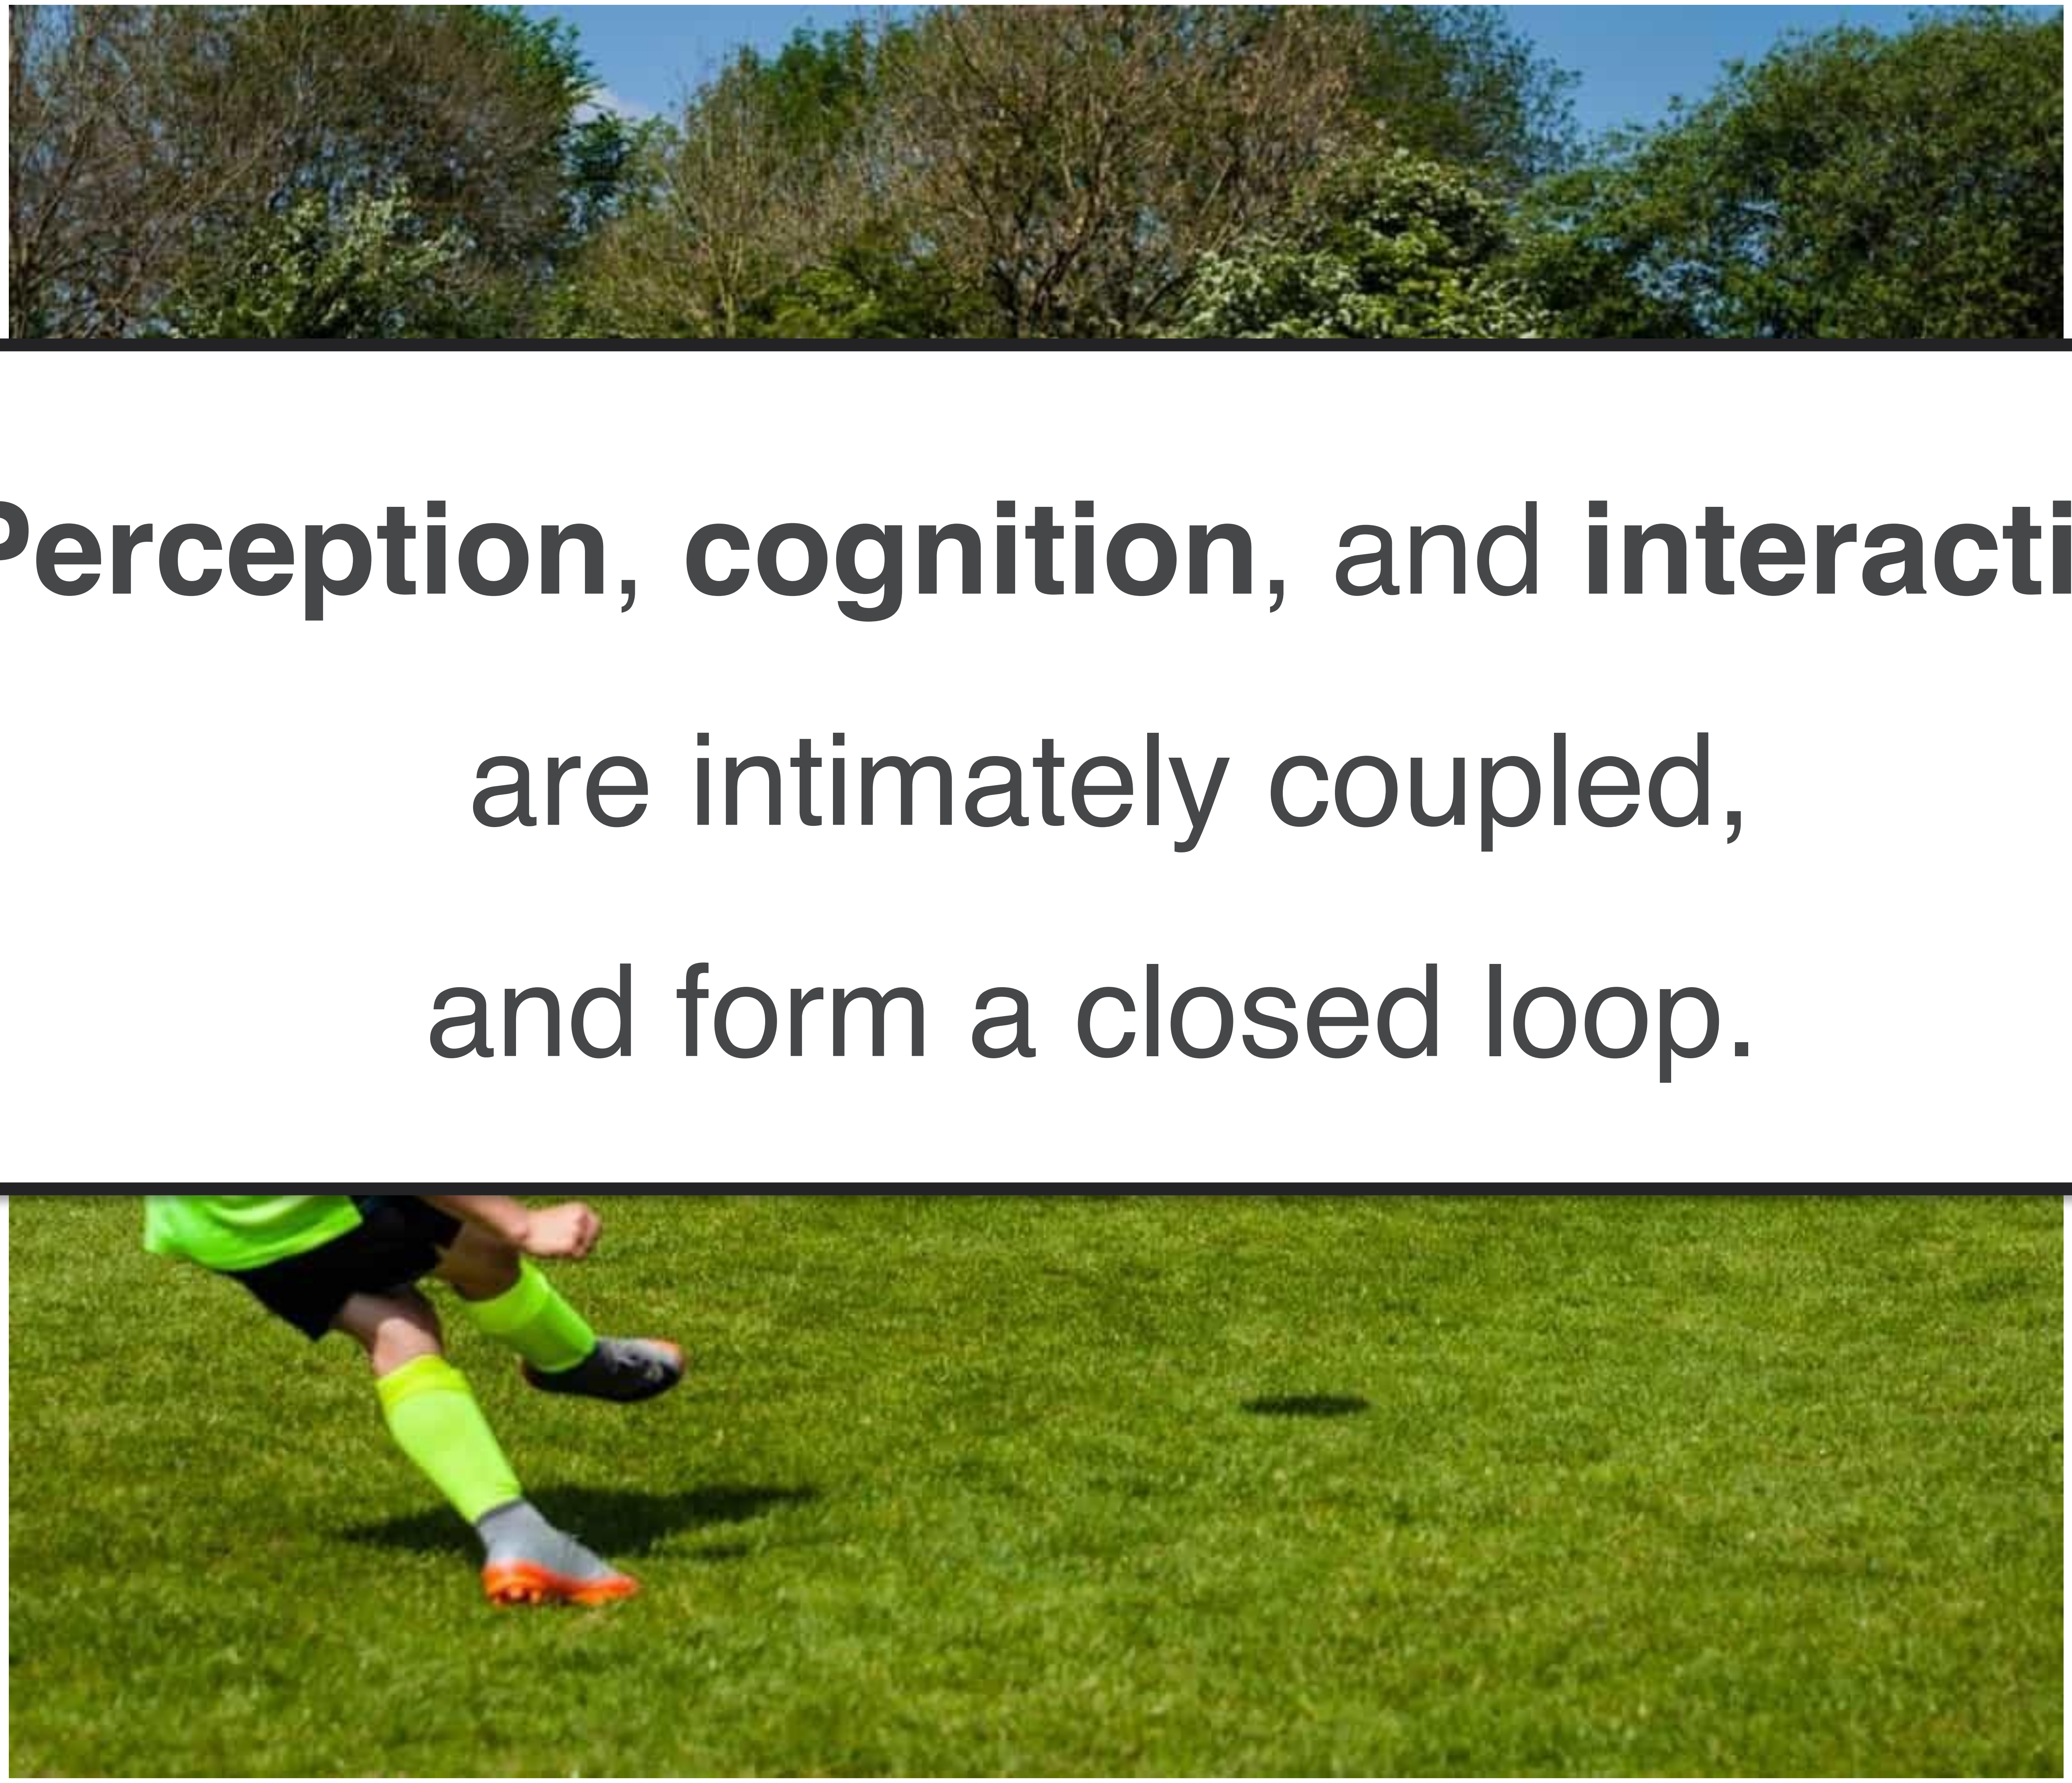

Perception, **cognition**, and **interaction**

are intimately coupled,

and form a closed loop.

Image source: https://www.thesoccerstore.co.uk/blog/football-goals/best-kids-garden-football-goals/

Embodiment Hypothesis:

"intelligence emerges in the interaction of an agent with an environment and as a result of sensorimotor activity"

"智能在智能体与环境的交互中涌现，是感觉运动行为的结果"

The Development of Embodied Cognition: Six Lessons from Babies

Linda Smith
Psychology Department
Indiana University
Bloomington, IN 47405
smith4@Indiana.edu

Michael Gasser
Computer Science Department
Indiana University
Bloomington, IN 47405
gasser@Indiana.edu

Abstract The embodiment hypothesis is the idea that intelligence emerges in the interaction of an agent with an environment and as a result of sensorimotor activity. We offer six lessons for *developing* embodied intelligent agents suggested by research in developmental psychology. We argue that starting as a baby grounded in a physical, social, and linguistic world is crucial to the development of the flexible and inventive intelligence that characterizes humankind.

Keywords
Development, cognition, language, embodiment, motor control

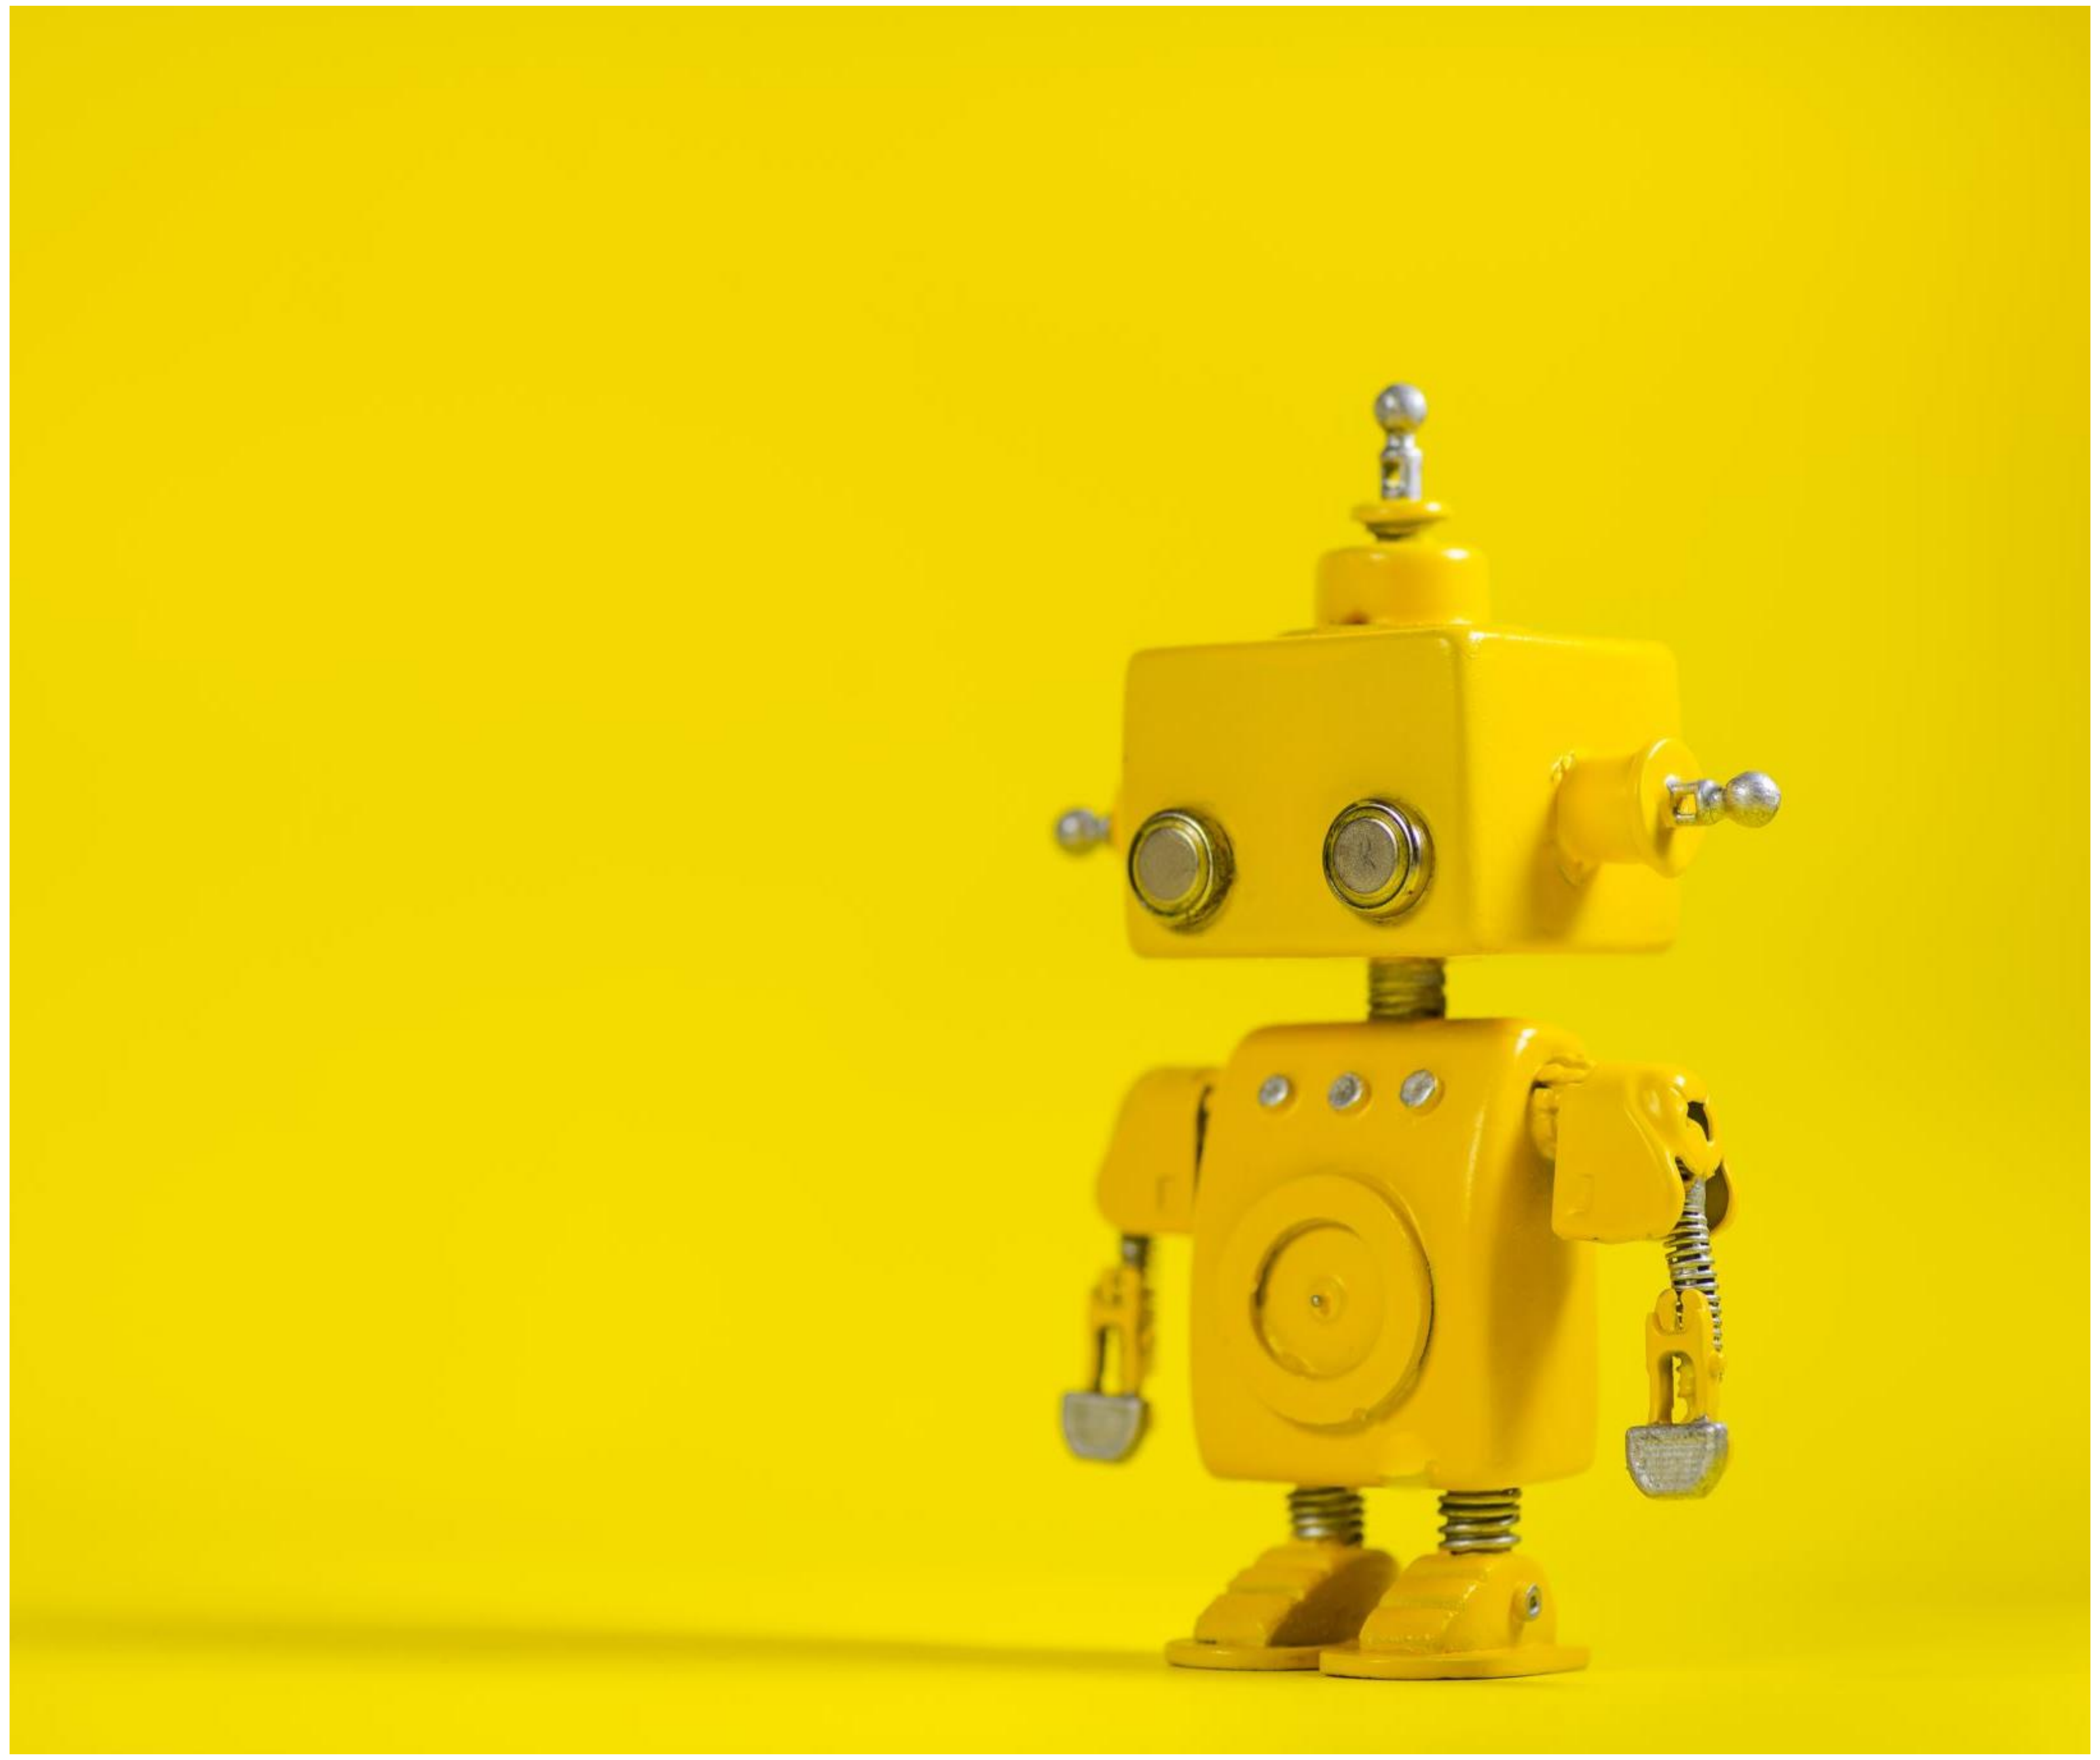

Why Embodied AI Now?

Dependencies of Embodied AI Research

Neural network: 3D deep learning architecture

Data source: Low-cost interactive environments

Framework: Closed-loop learning framework

Dependencies of Embodied AI Research

Neural network: 3D deep learning architecture

–Recommend Prof. Huang's course and my course

Data source: Low-cost interactive environments

Framework: Closed-loop learning framework

Dependencies of Embodied AI Research

Neural network: 3D deep learning architecture

– *Recommend Prof. Huang's course and my course*

Data source: Low-cost interactive environments

– *My main topic today*

Framework: Closed-loop learning framework

– *Will touch upon this topic in the end*

Outline

- **Low-cost interactive environments**
  - SAPIEN++: Realistic Simulator for Manipulation Research
  - ManiSkill: SAPIEN Manipulation Challenge
- Frameworks of Learning-from-Demonstrations

SAPIEN++

Realistic Simulator for Manipulation Research

Options for Collecting Interaction Data

- Collect data in real world
  - Costly (hardware, human labor)
  - Dangerous

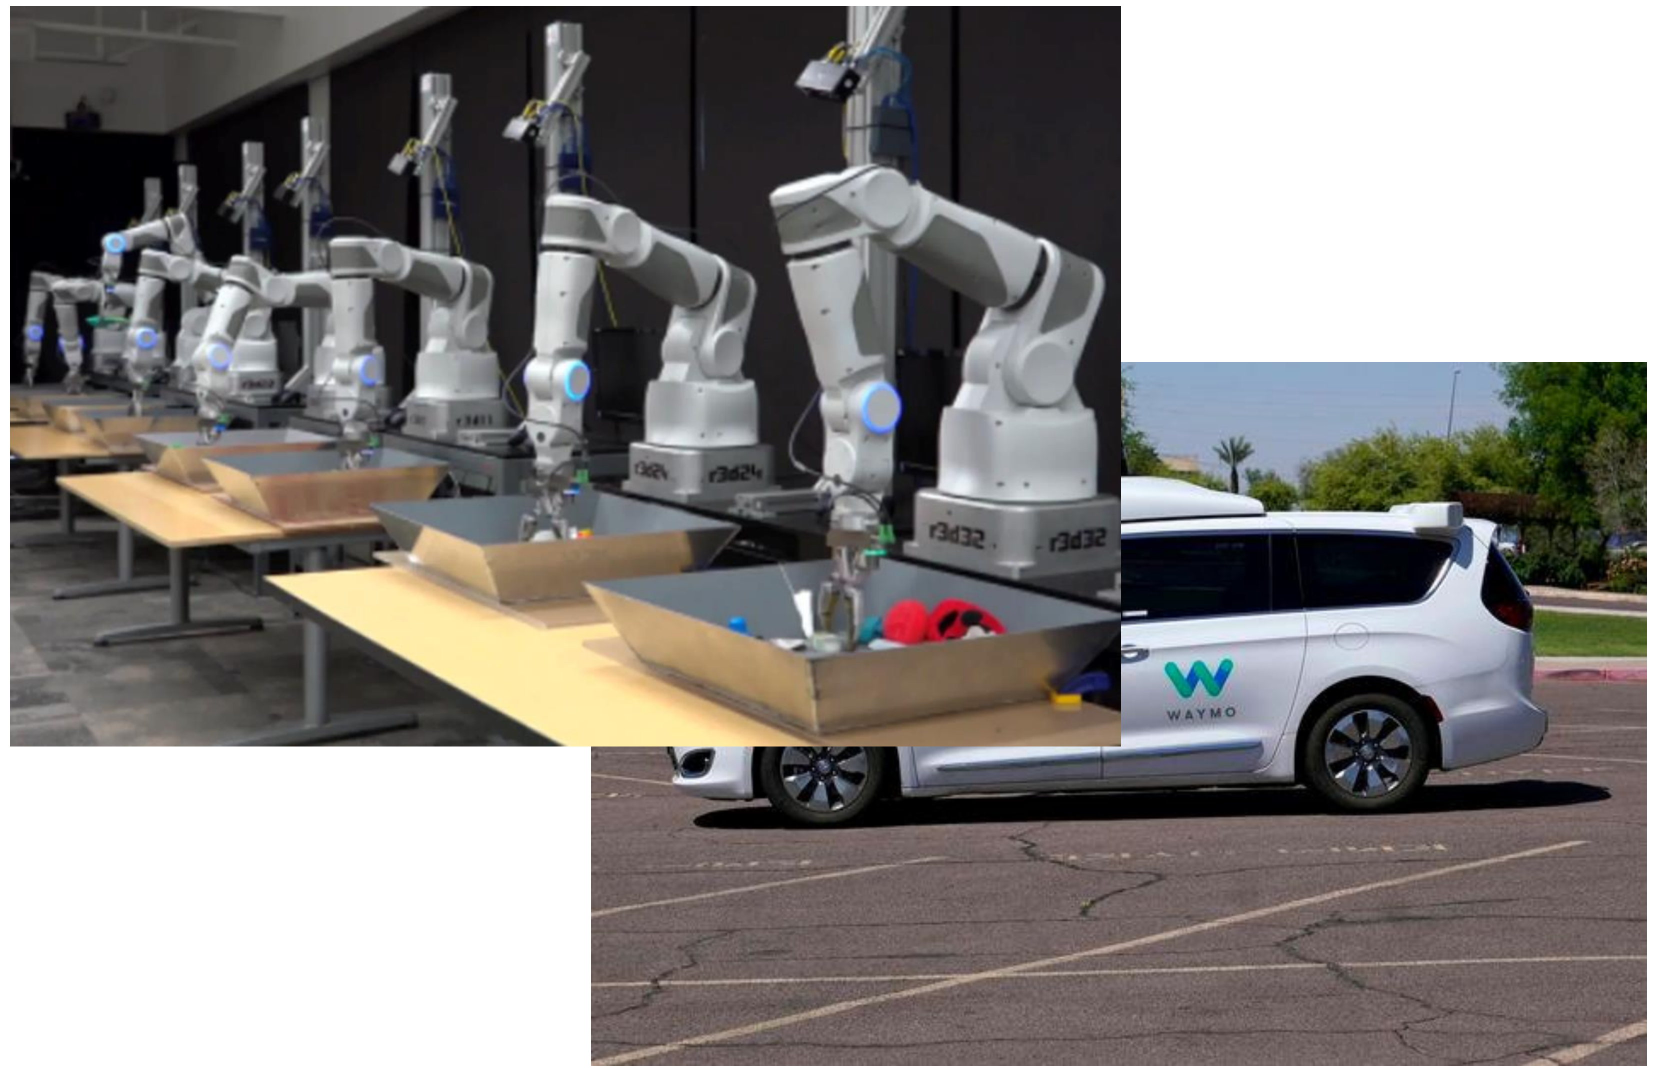

Want to study embodied AI but cannot afford the device and time?

Worry about research reproducibility?

Options for Collecting Interaction Data

- Collect data in real world
  - Costly (hardware, human labor)
  - Dangerous

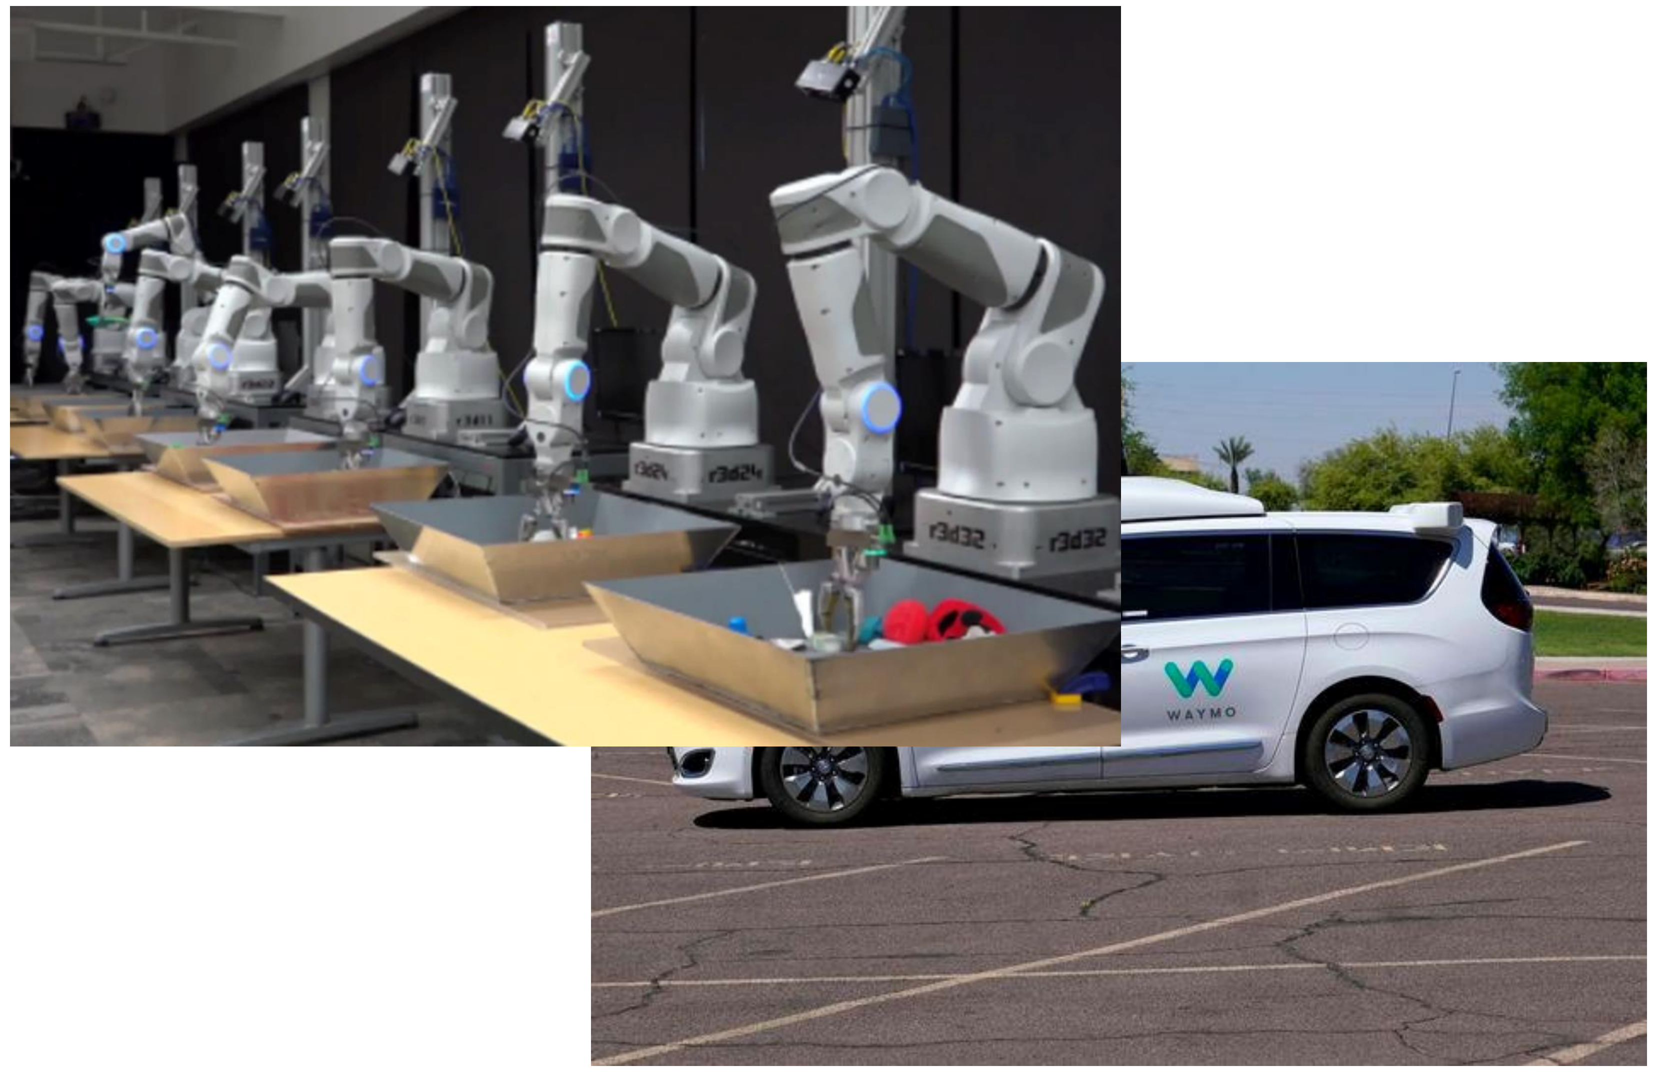

- Collect data in a simulator
  - Scalable
  - Safe

Previous Version: SAPIEN (CVPR 2020 Oral)

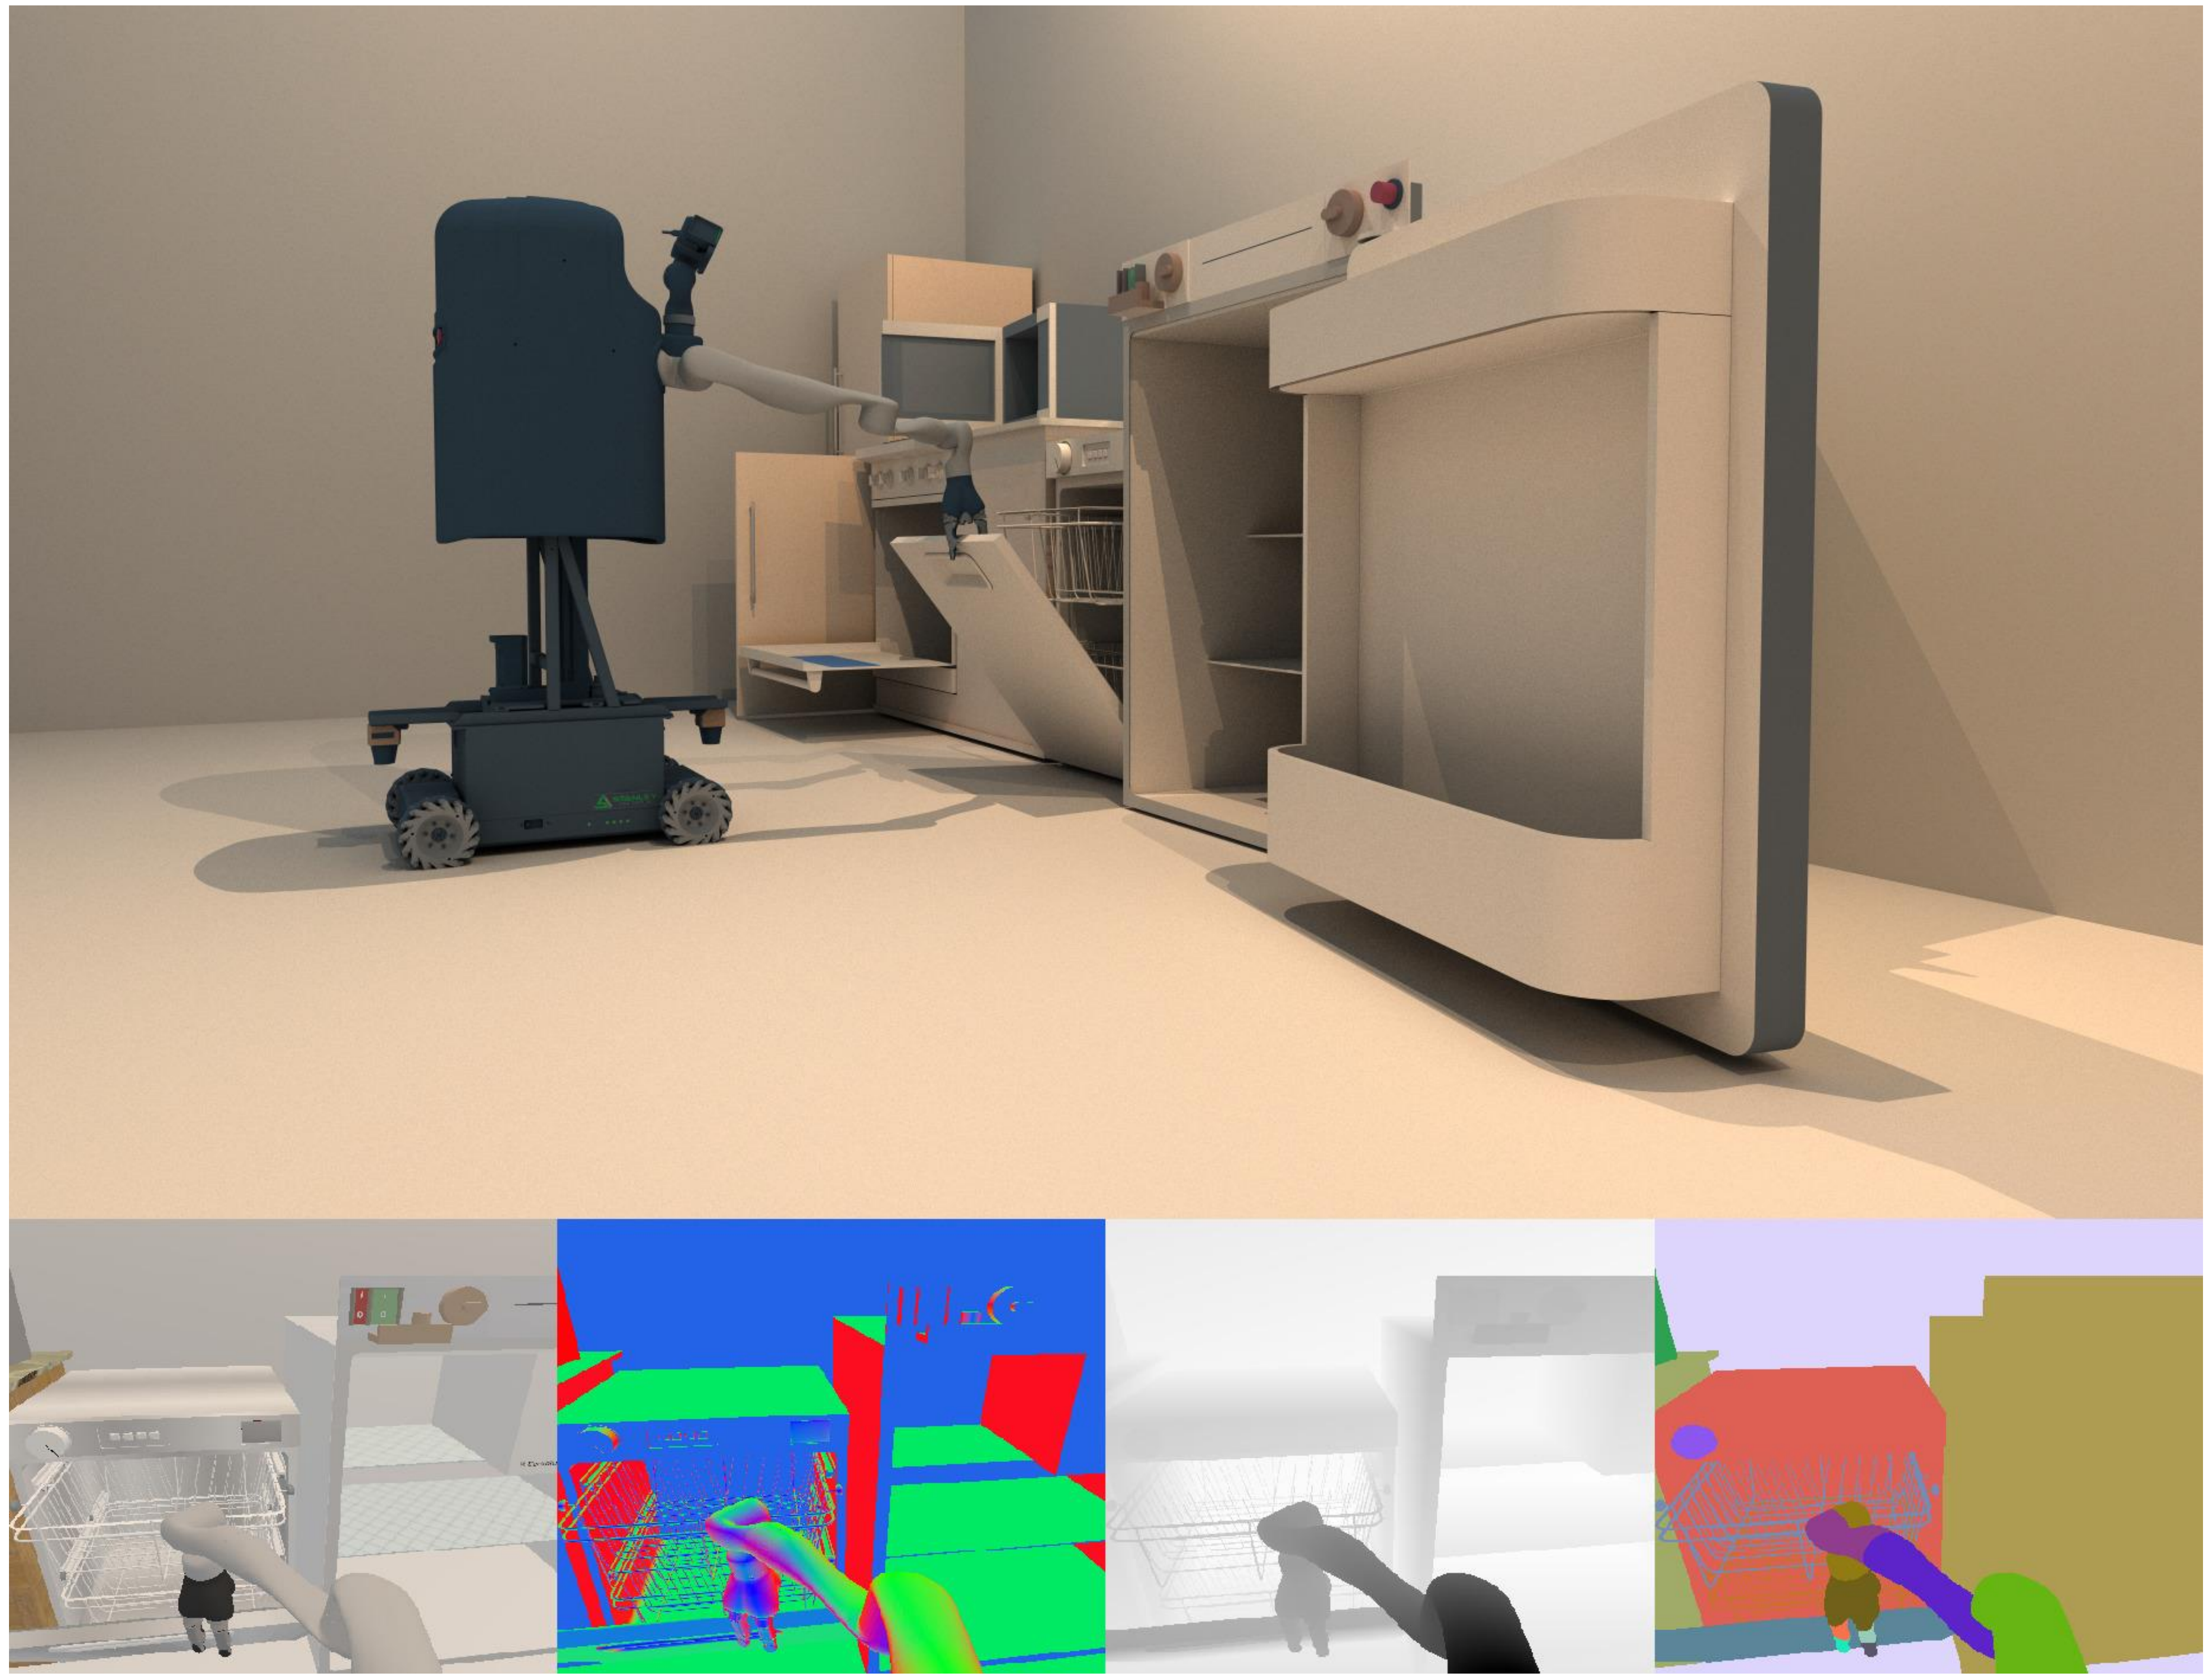

A **S**imul**A**ted **P**art-based **I**nteractive **EN**vironment

Sapien

[seɪ-piən]:

Derives from an old Latin word meaning **‘Wise’** (智慧).

Used together and written as **‘Homo sapien’** (智人), it describes a species of man or human being.

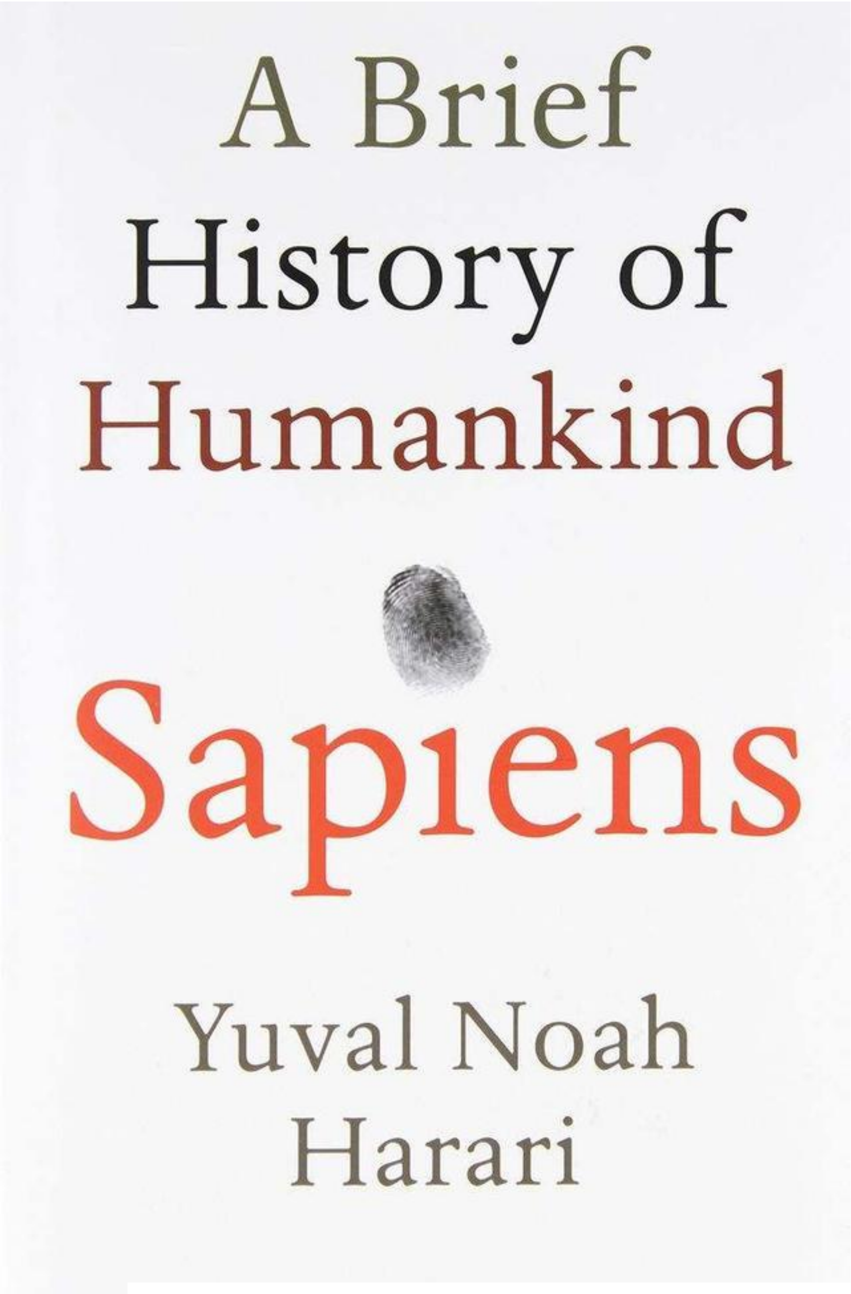

(人类简史)

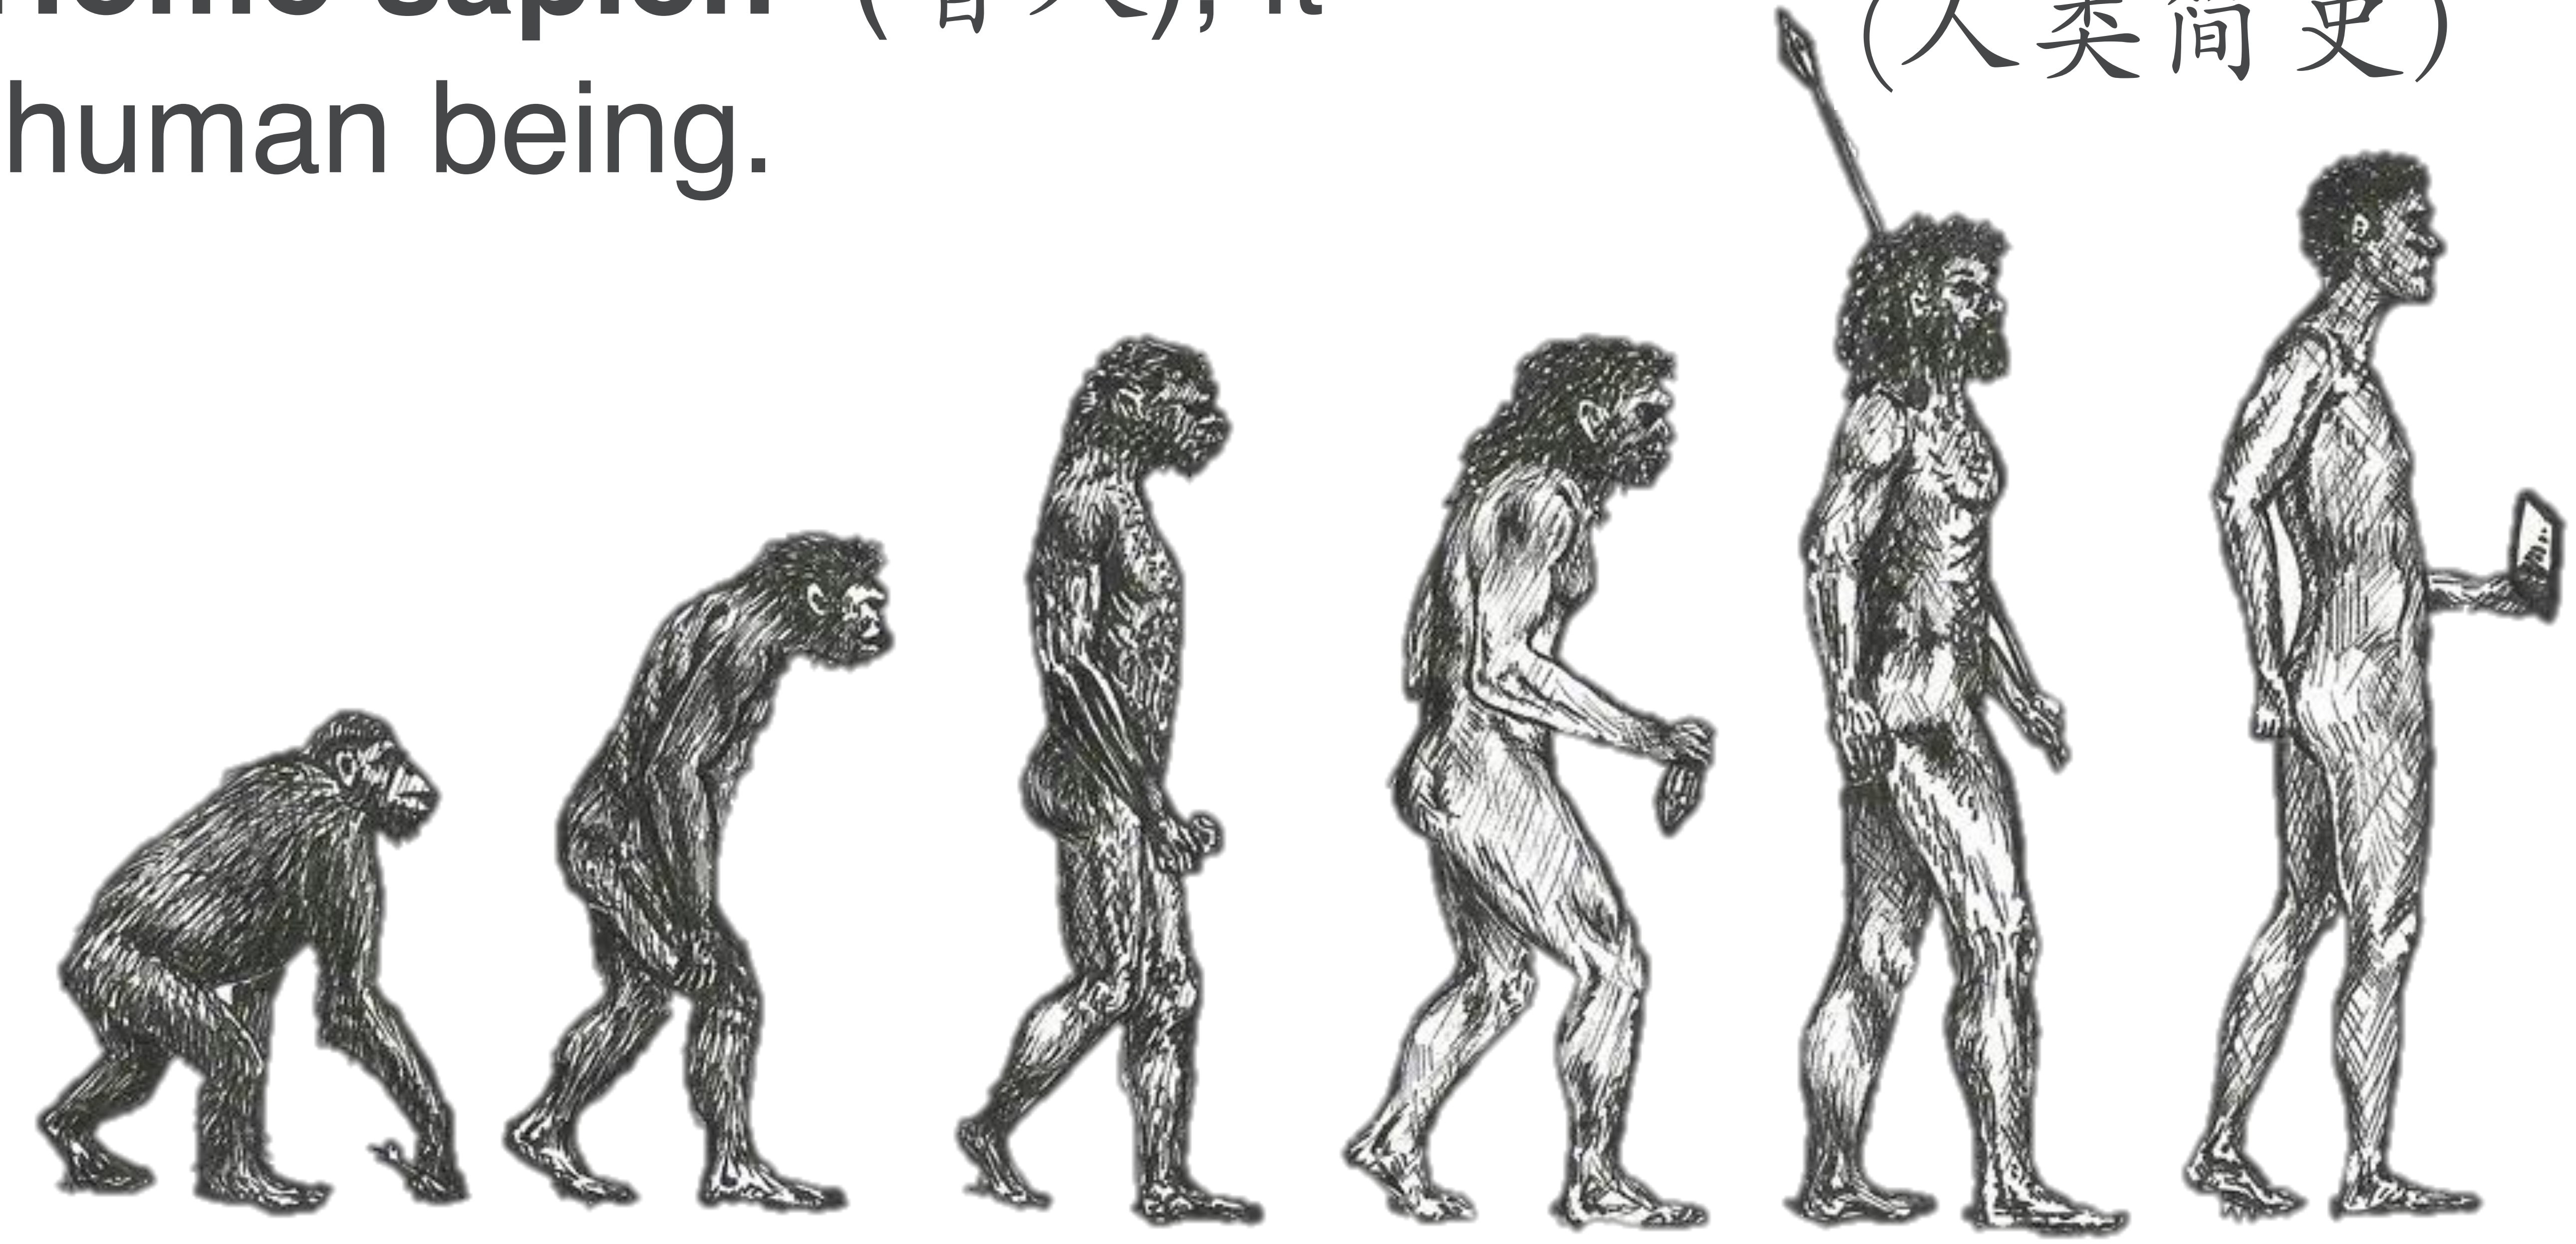

Figure from http://isciencemag.co.uk/features/the-rise-of-homo-sapiens

SAPIEN++

Realistic simulation

Manipulation-oriented

User-friendly

Realistic simulation

Mind the Sim2Real Domain Gap

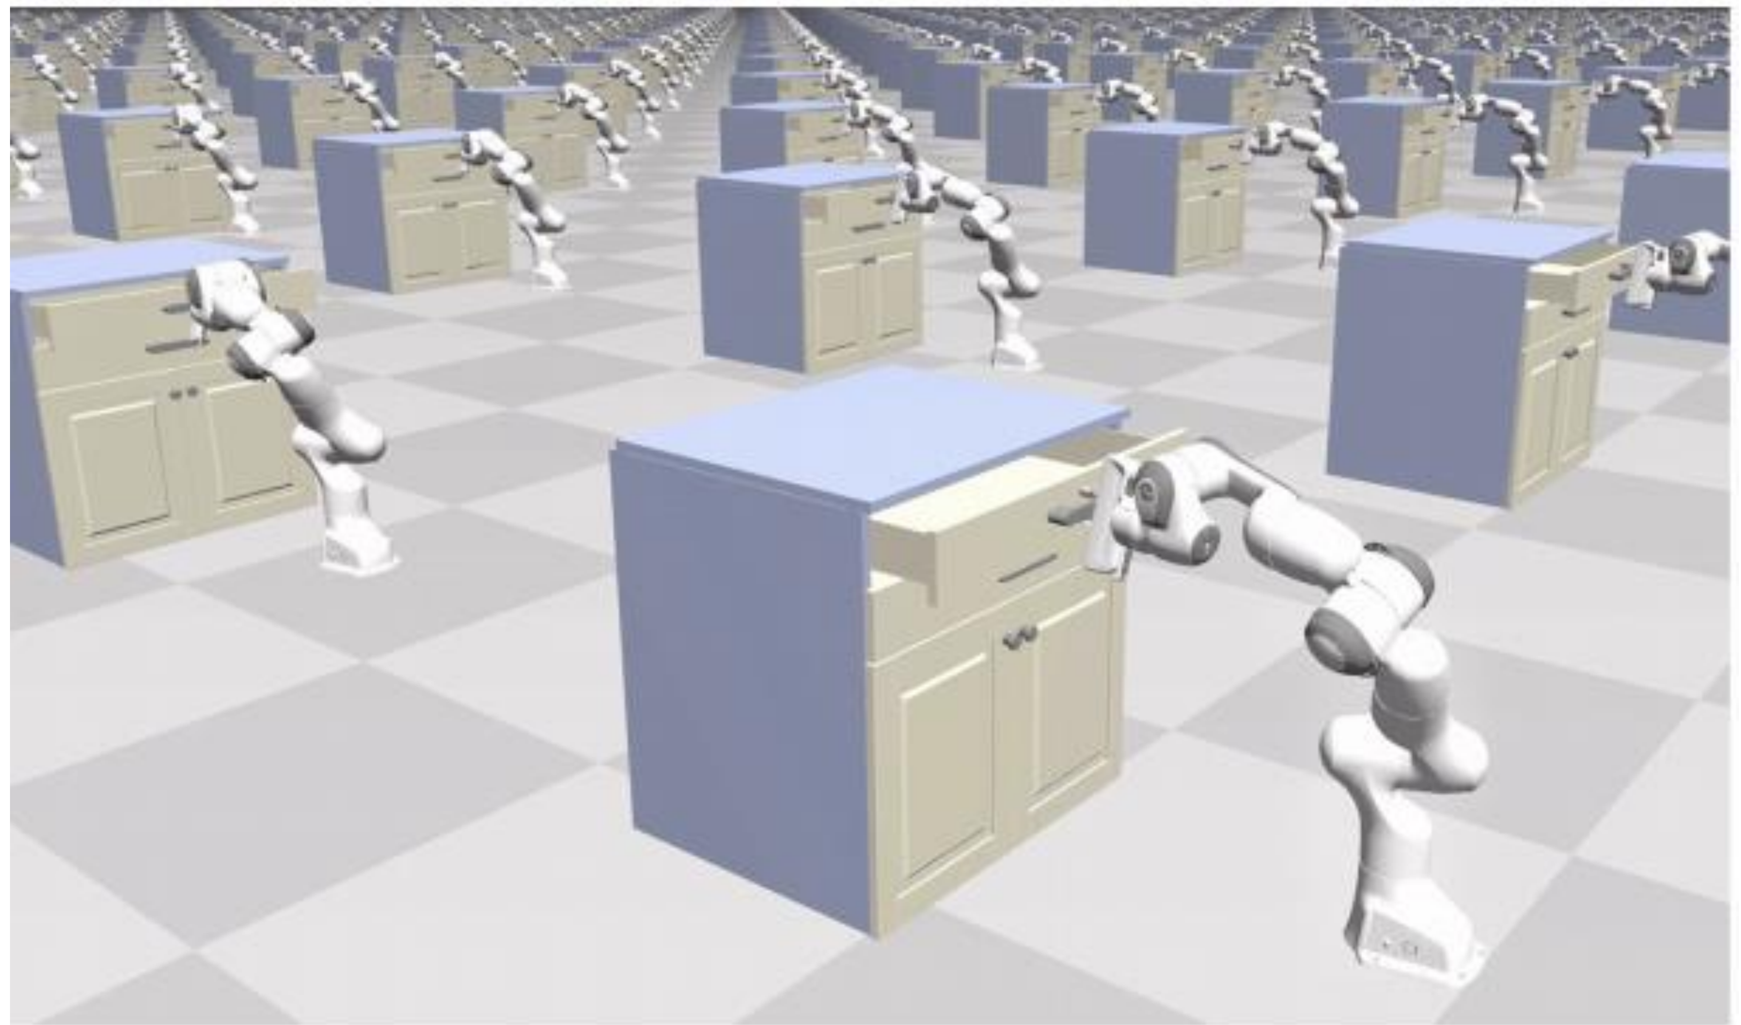

Where you afford to learn:
"simulator"

Where your robot works:
"real world"

Image credit: phdcomics.com, Adapting Simulation Randomization with Real World Experience

SAPIEN++ provide you with

more realistic visual input

Ray-tracing-based rendering

Empowered by NVIDIA OptiX DL-based Denoiser
(16 SPP, 3 bounce, 60 FPS for static scenes)

Transferability of SAPIEN

and more realistic dynamics

- Low-level physical simulation
- Physics-based contact model and robot control

Replaying the behavior of real robot in SAPIEN++ (velocity controller, xArm 7)

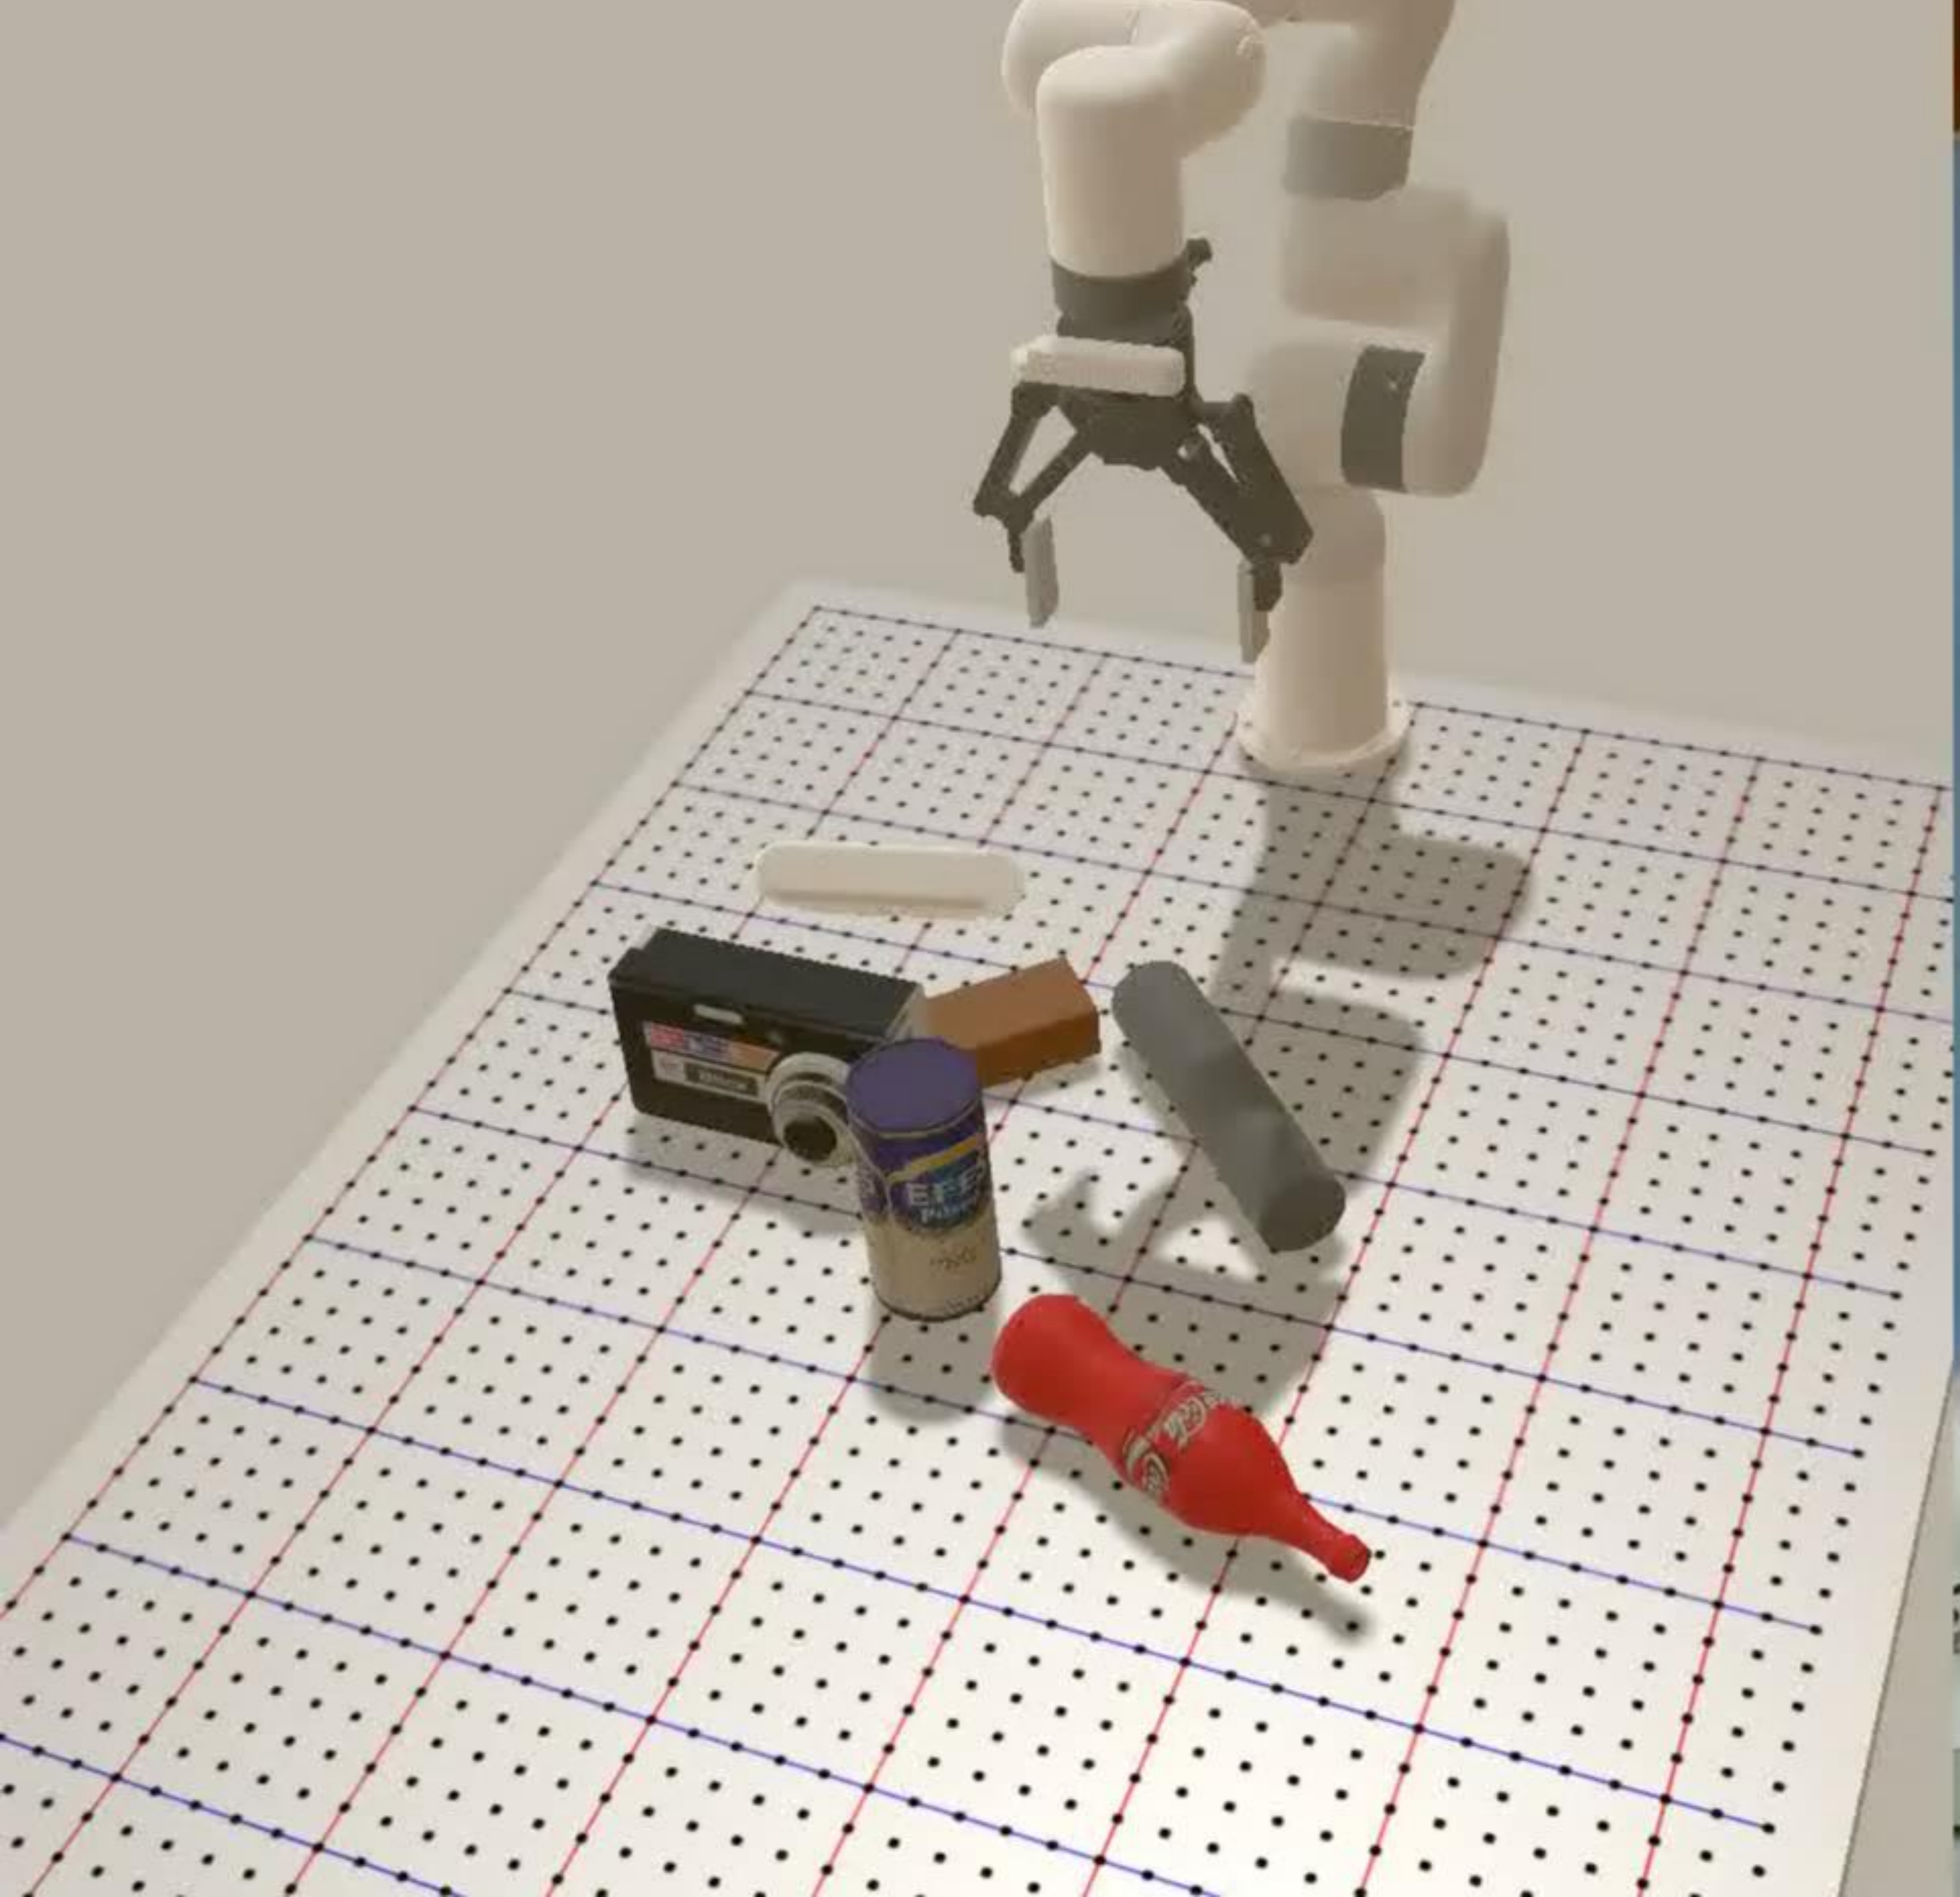

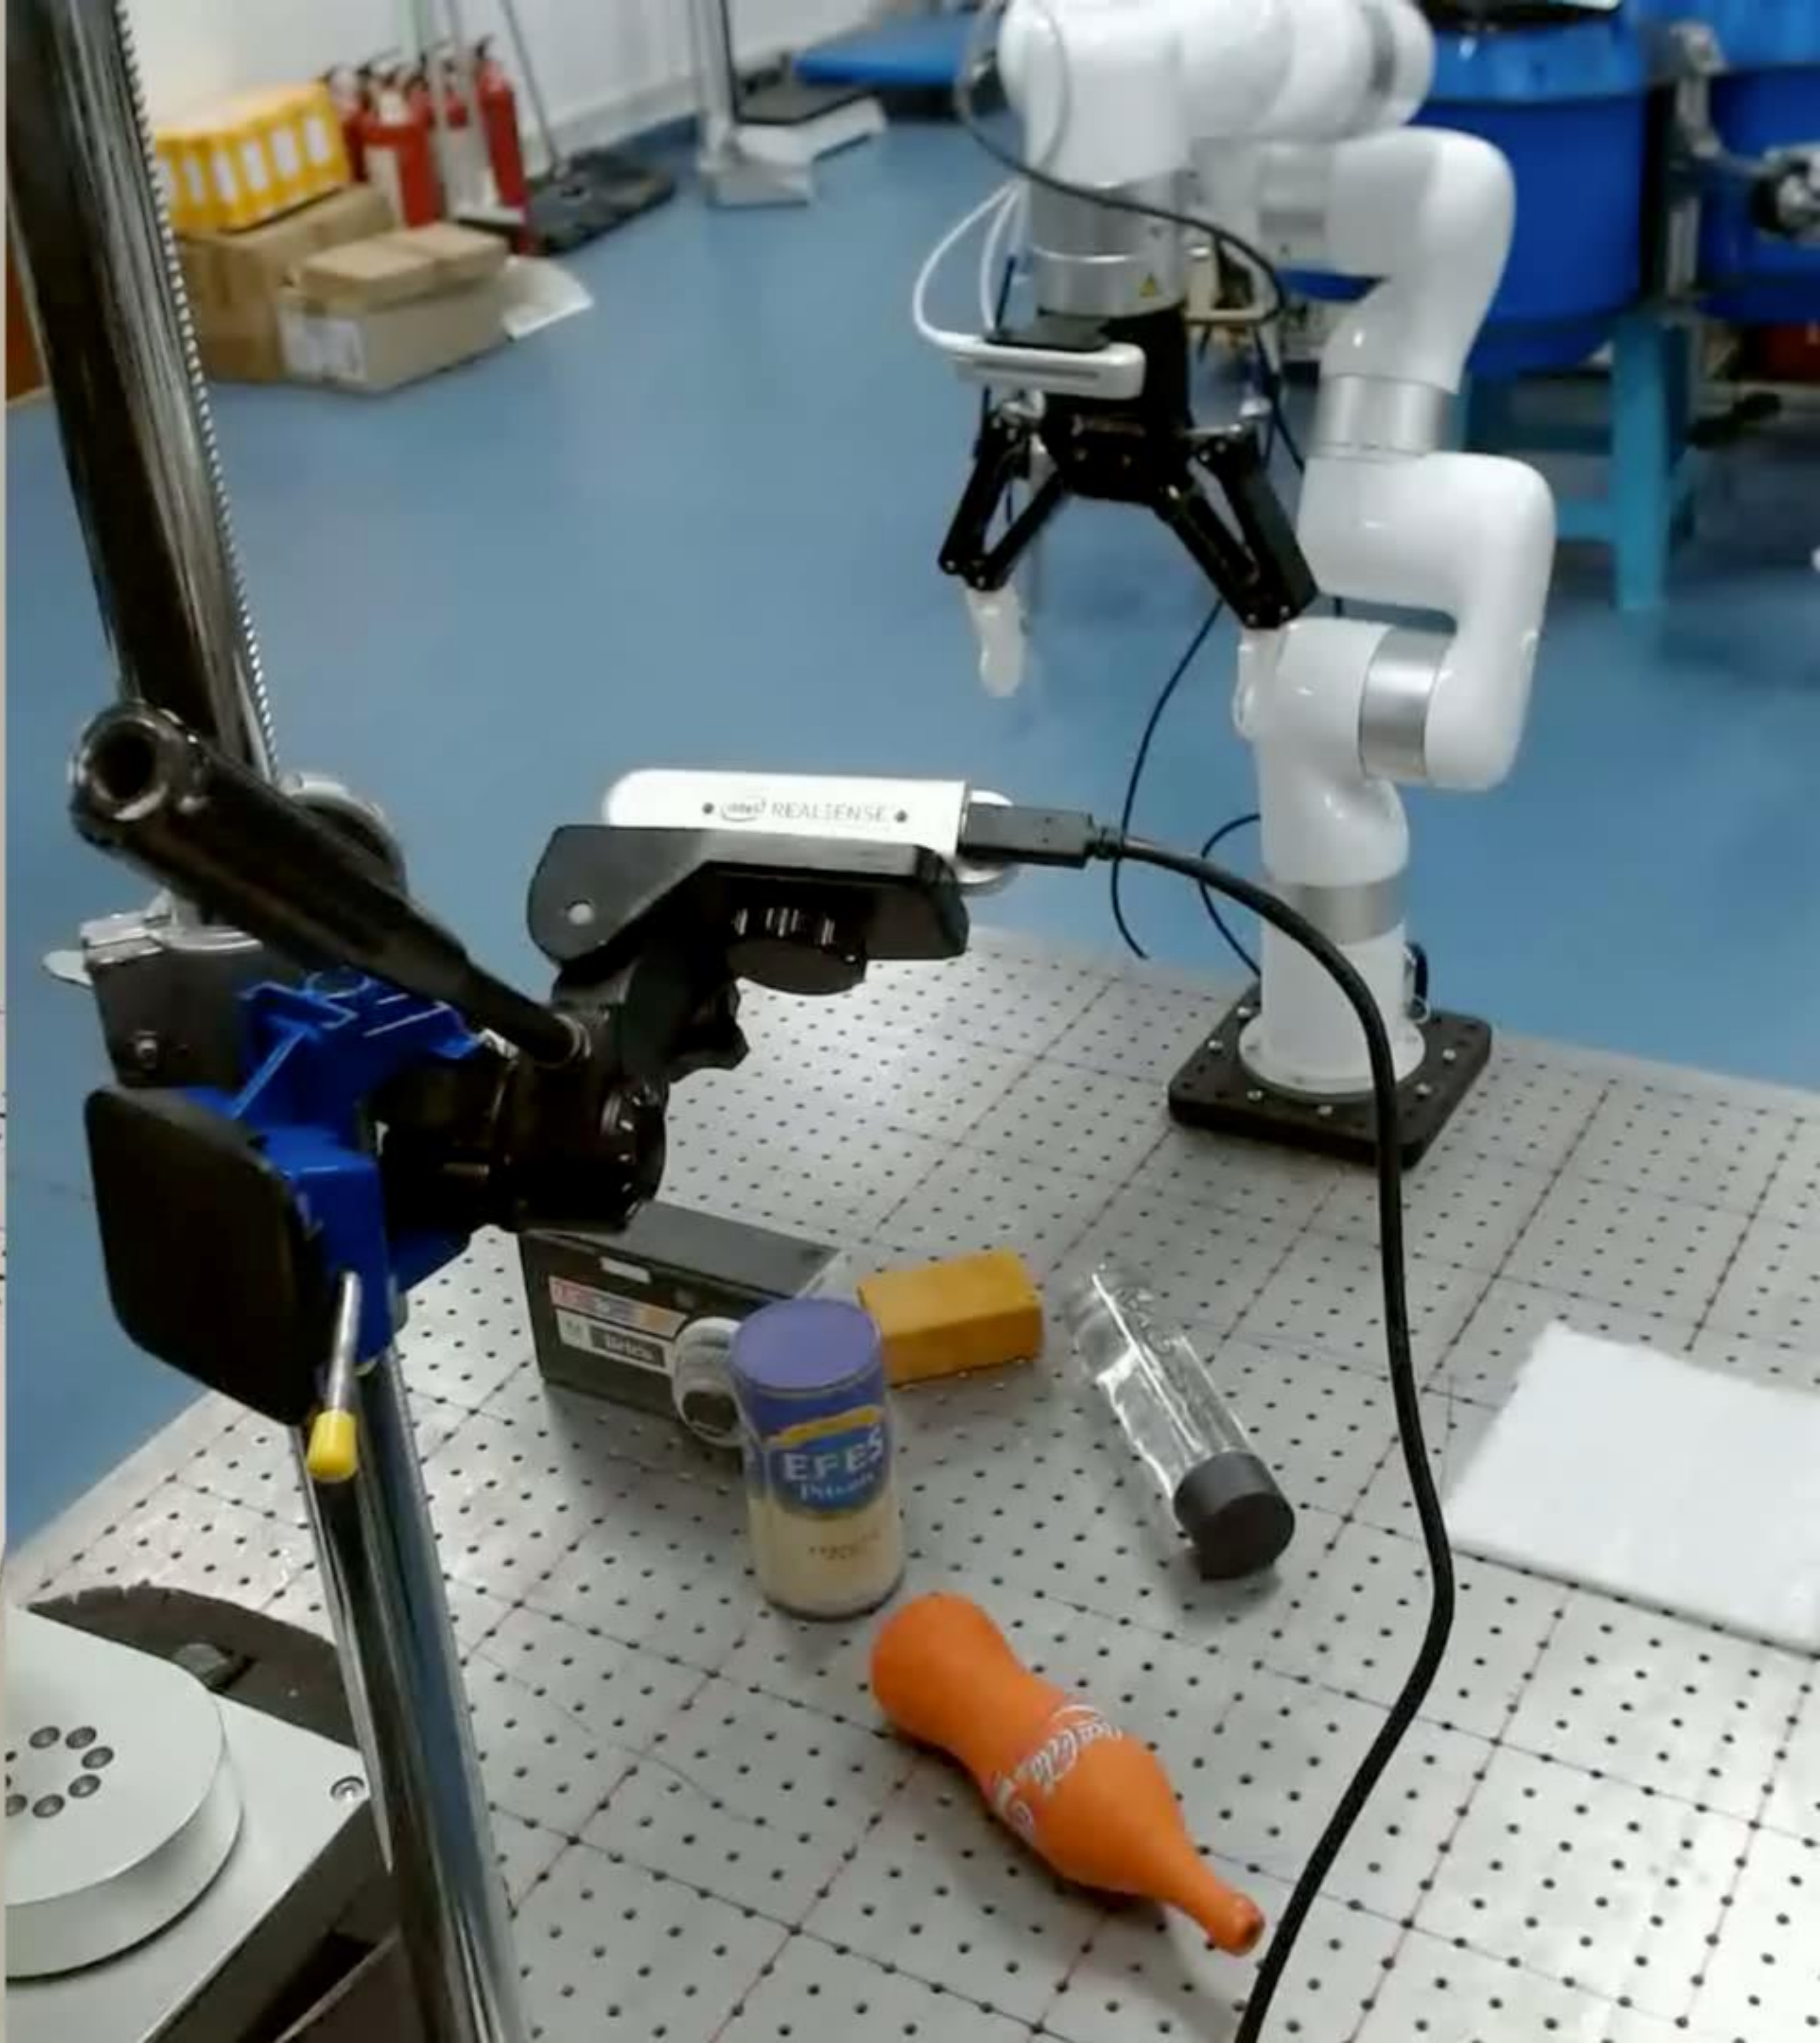

Manipulation-oriented

Off-the-shelf

manipulation environments

manipulation motion planner

Off-the-shelf Environments

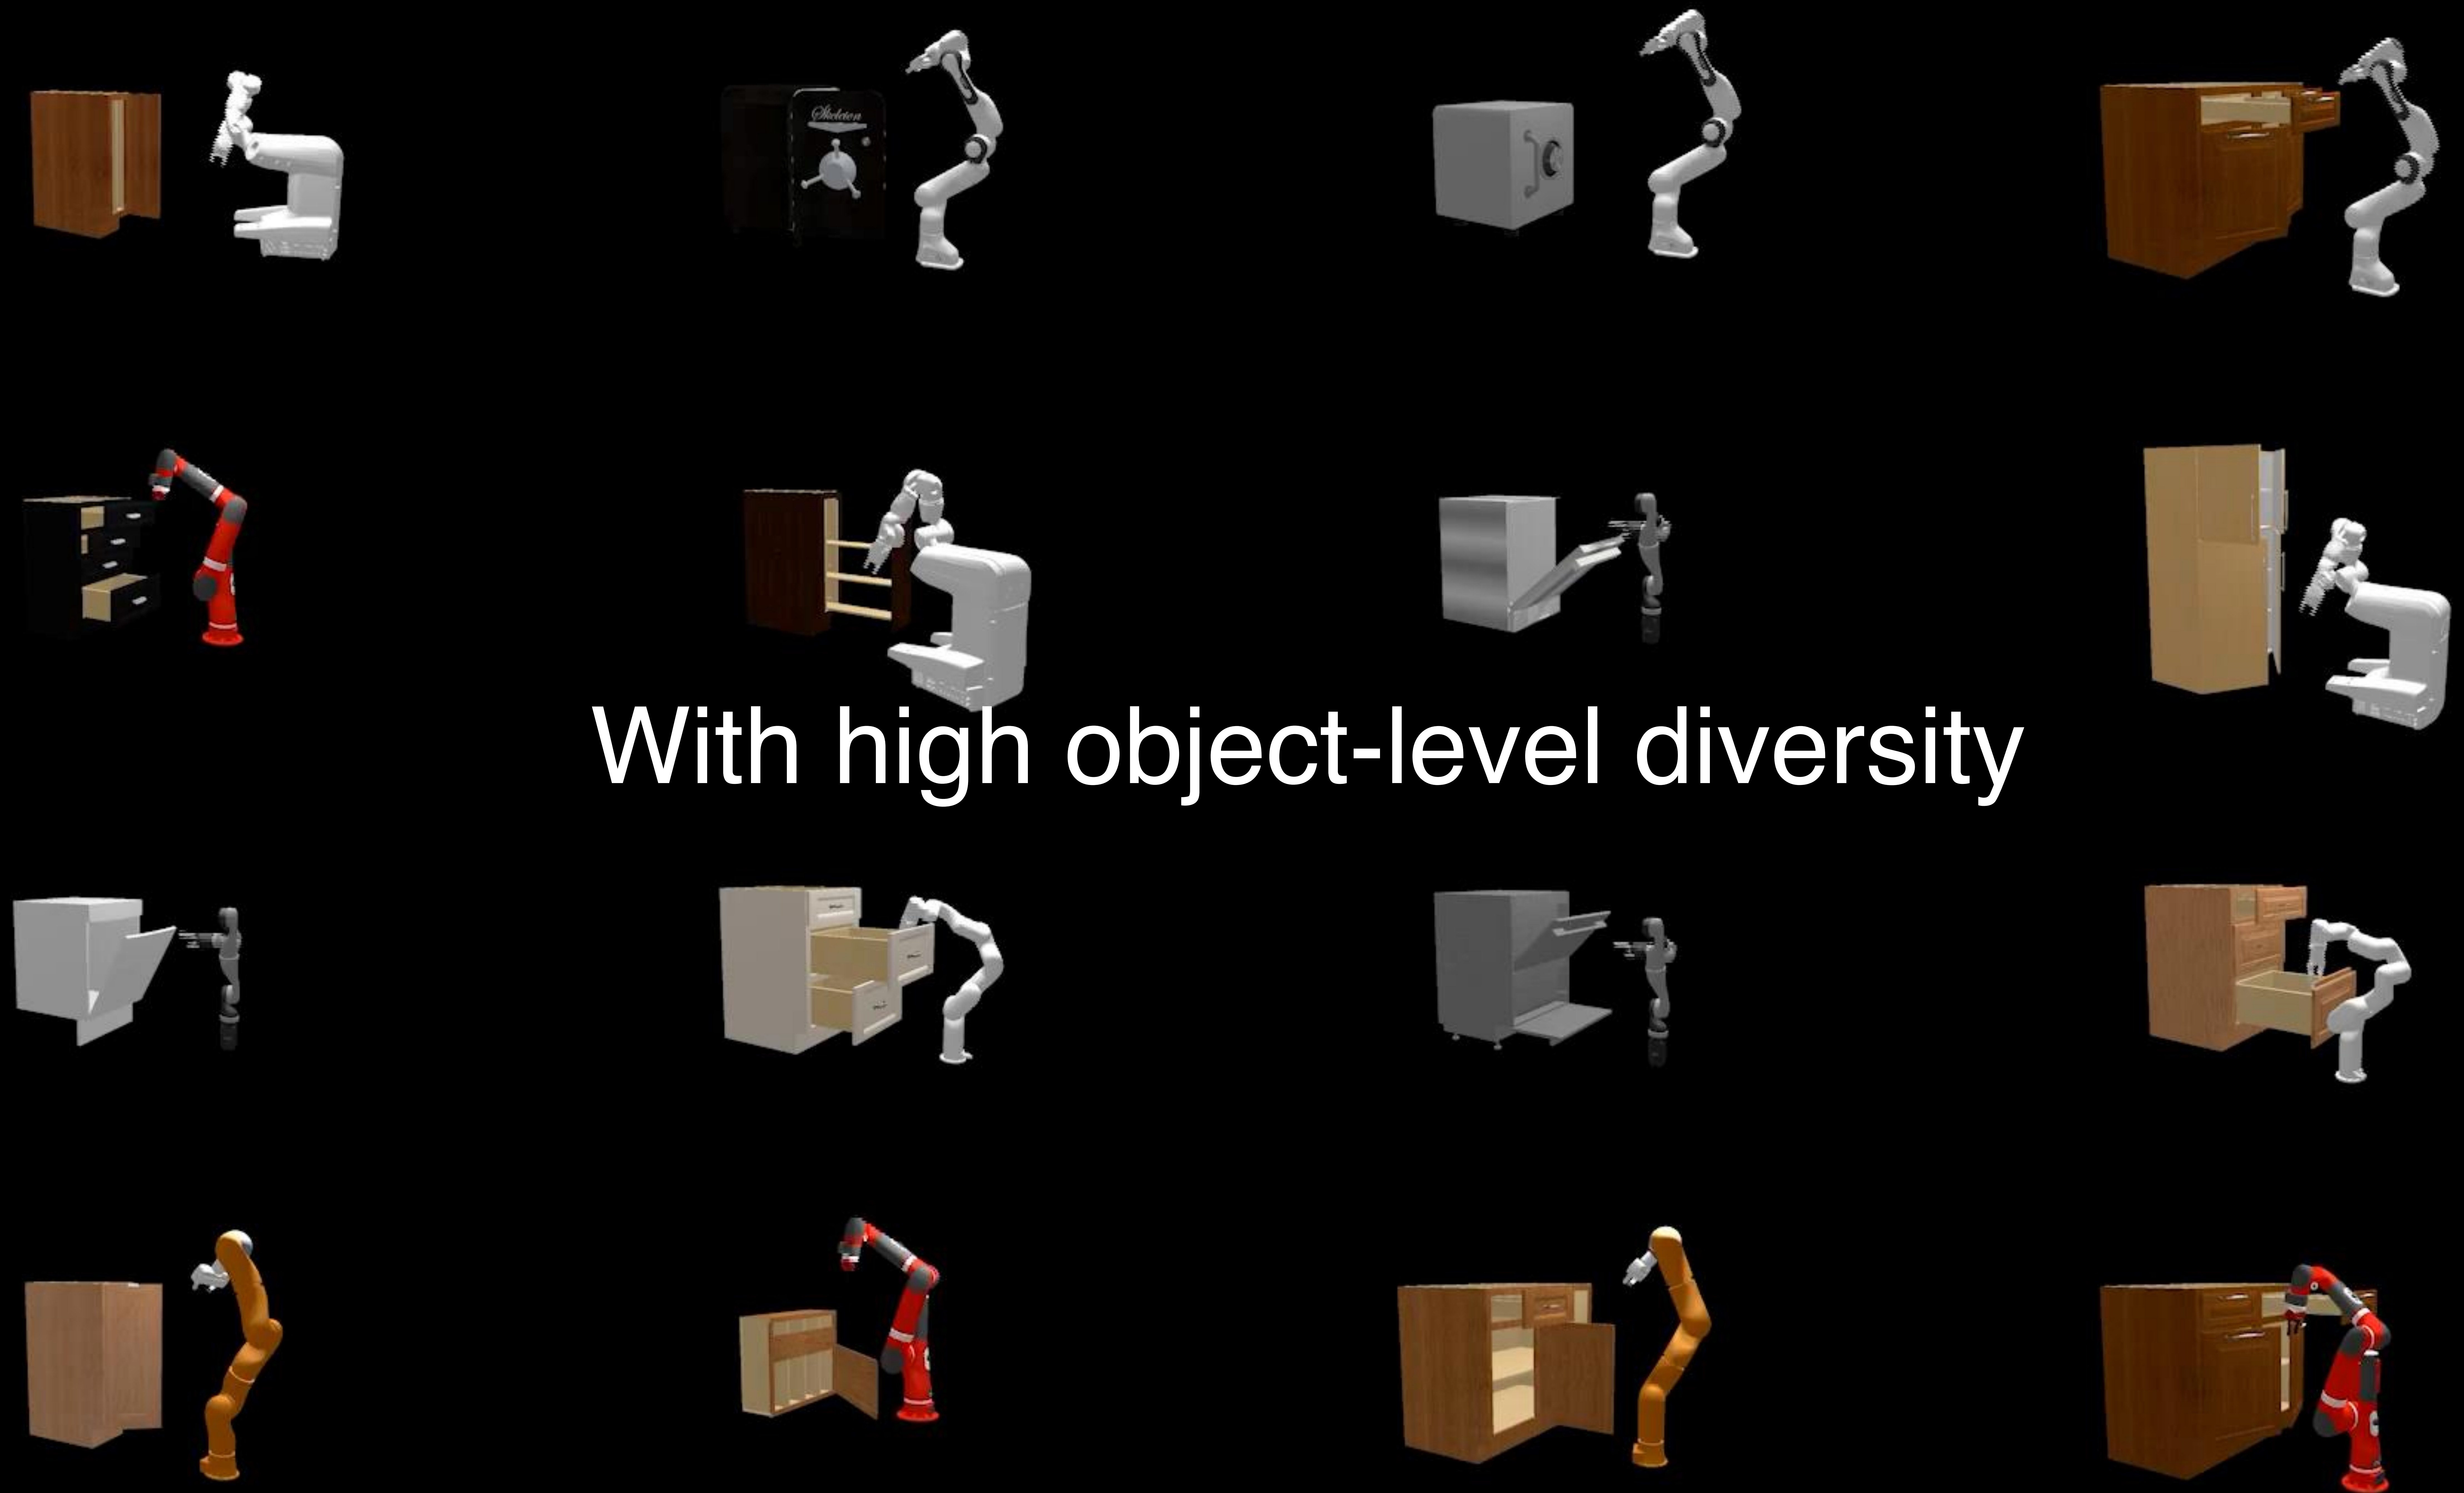
With high object-level diversity

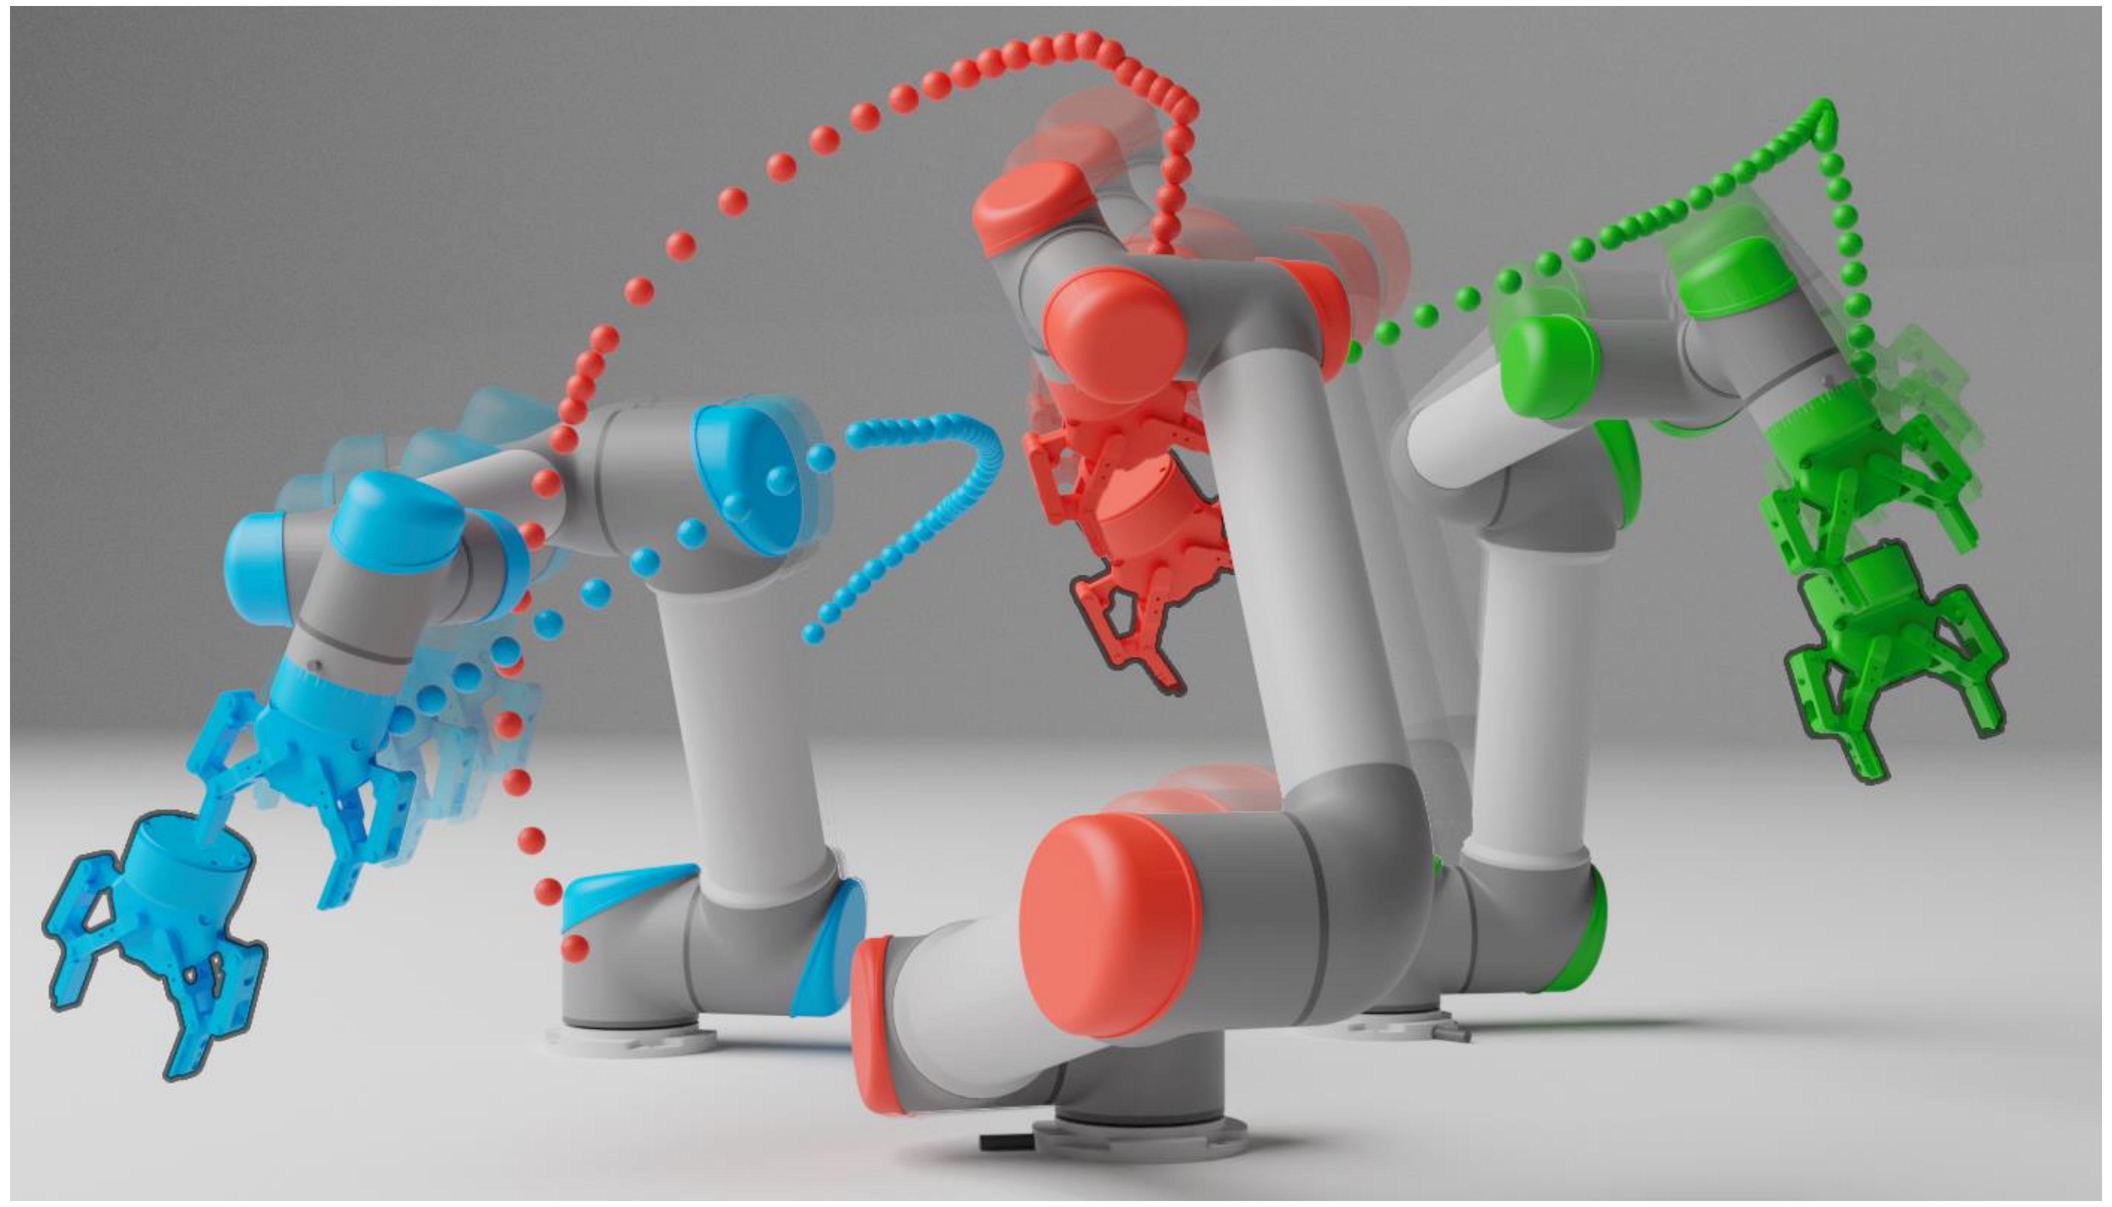

SOTA included

(based on OMPL + Pinocchio + FCL)

Compilation free

No ROS required

Image credit: ROS

```python
import sapien.core as sapien
from sapien_mp import Planner

planner = Planner(
  urdf='path/to/urdf',
  srdf='path/to/srdf',
  move_group='ee',
)

goal_pose = ...
current_qpos = ...

trajectory = planner.plan(
  goal_pose,
  current_qpos,
)
```

Build a planner with a few lines of **Python** codes

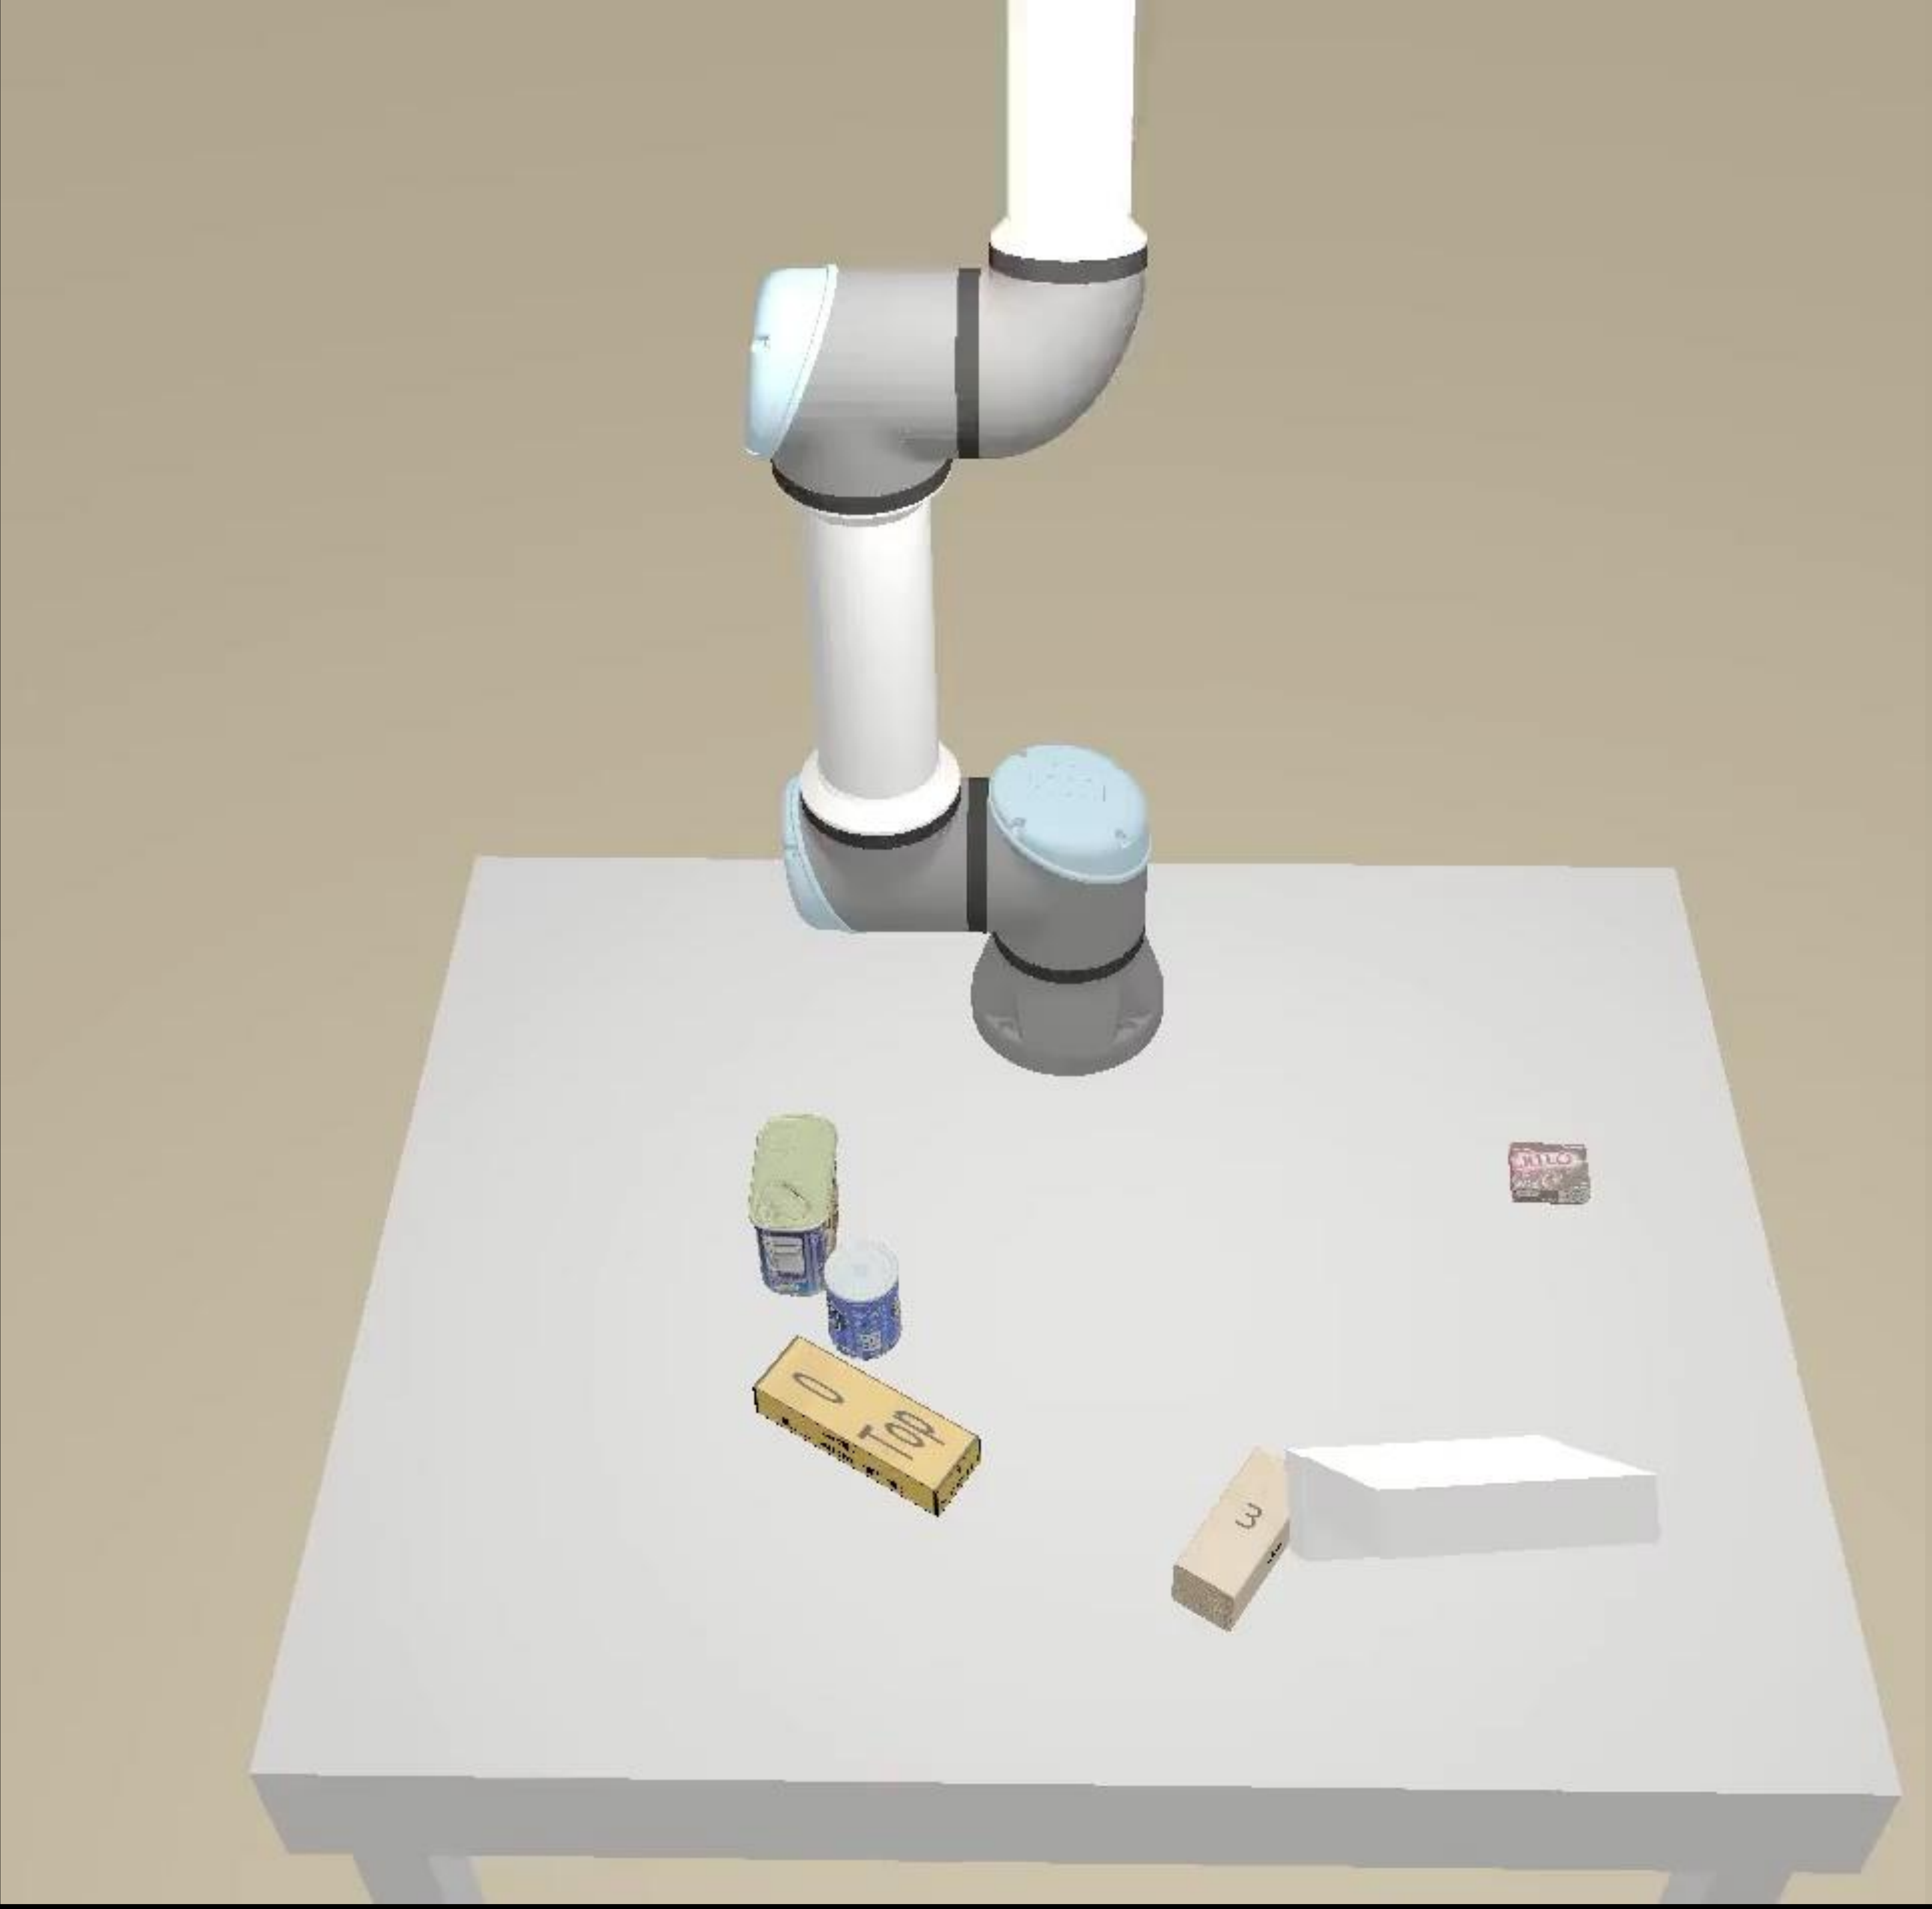

SAPIEN-mp

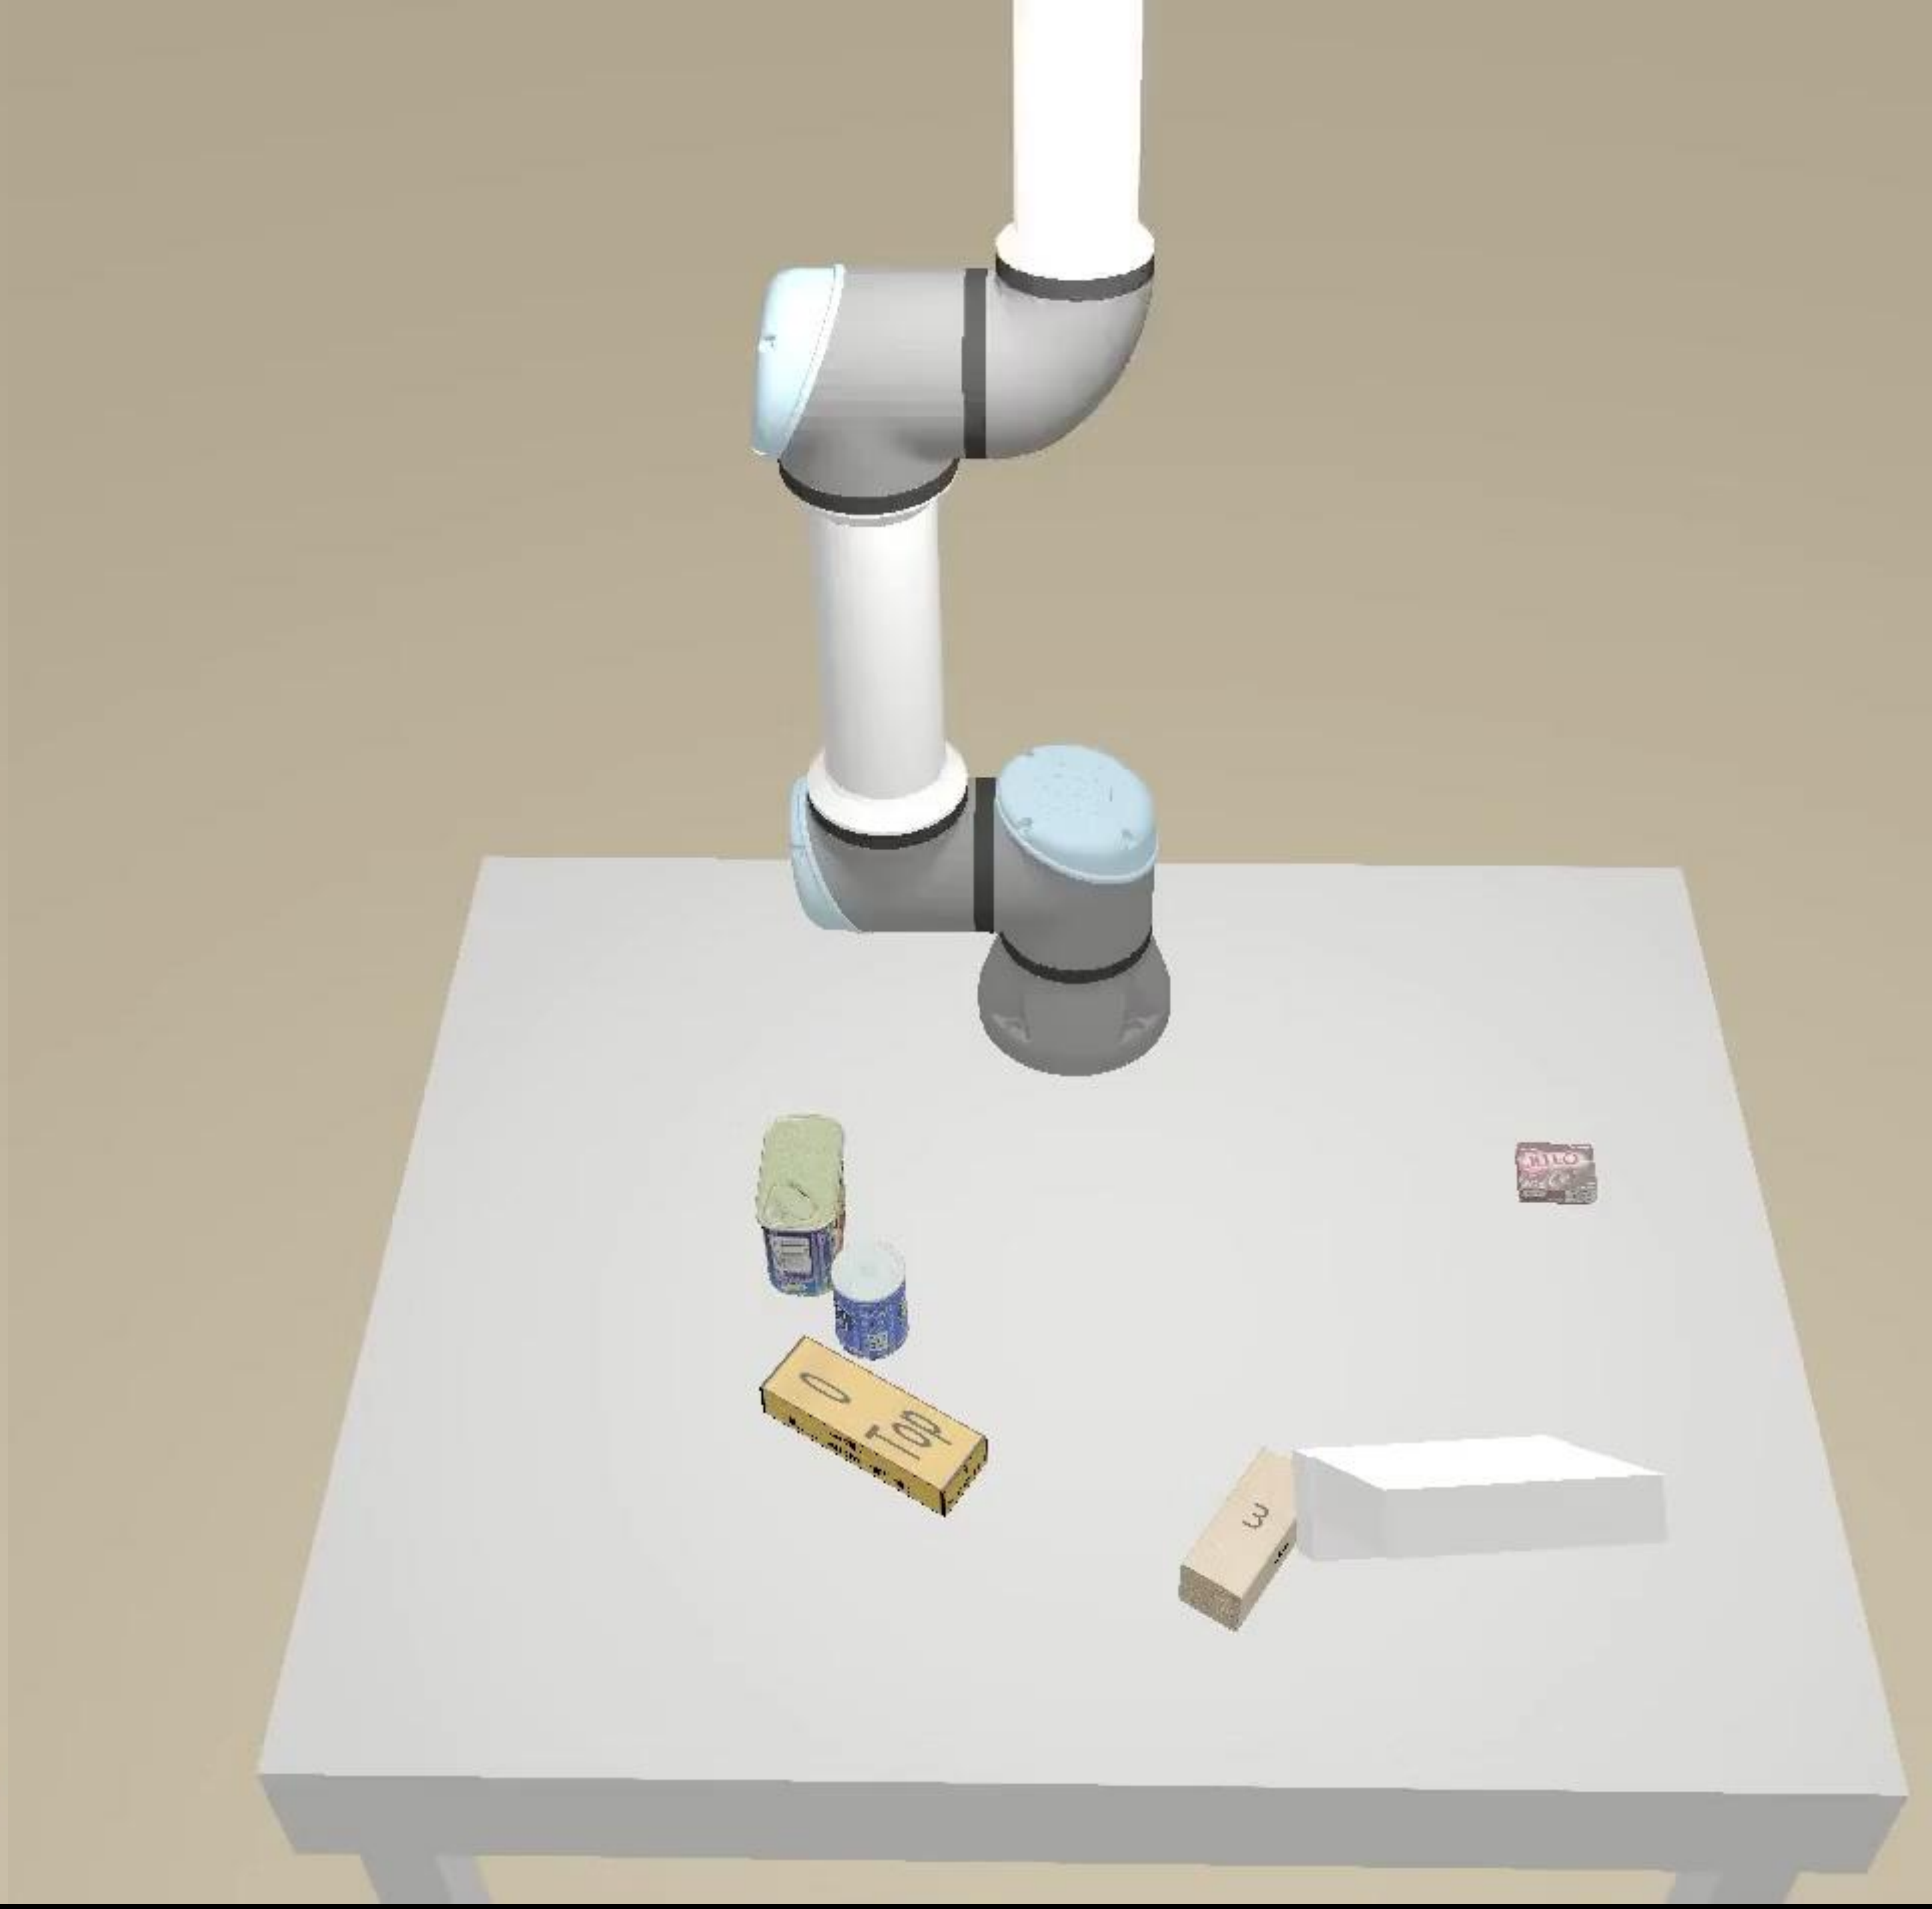

MoveIt (ROS required)

User-friendly

It is easy to set up SAPIEN:

```
pip install sapien
```

Many Examples Provided

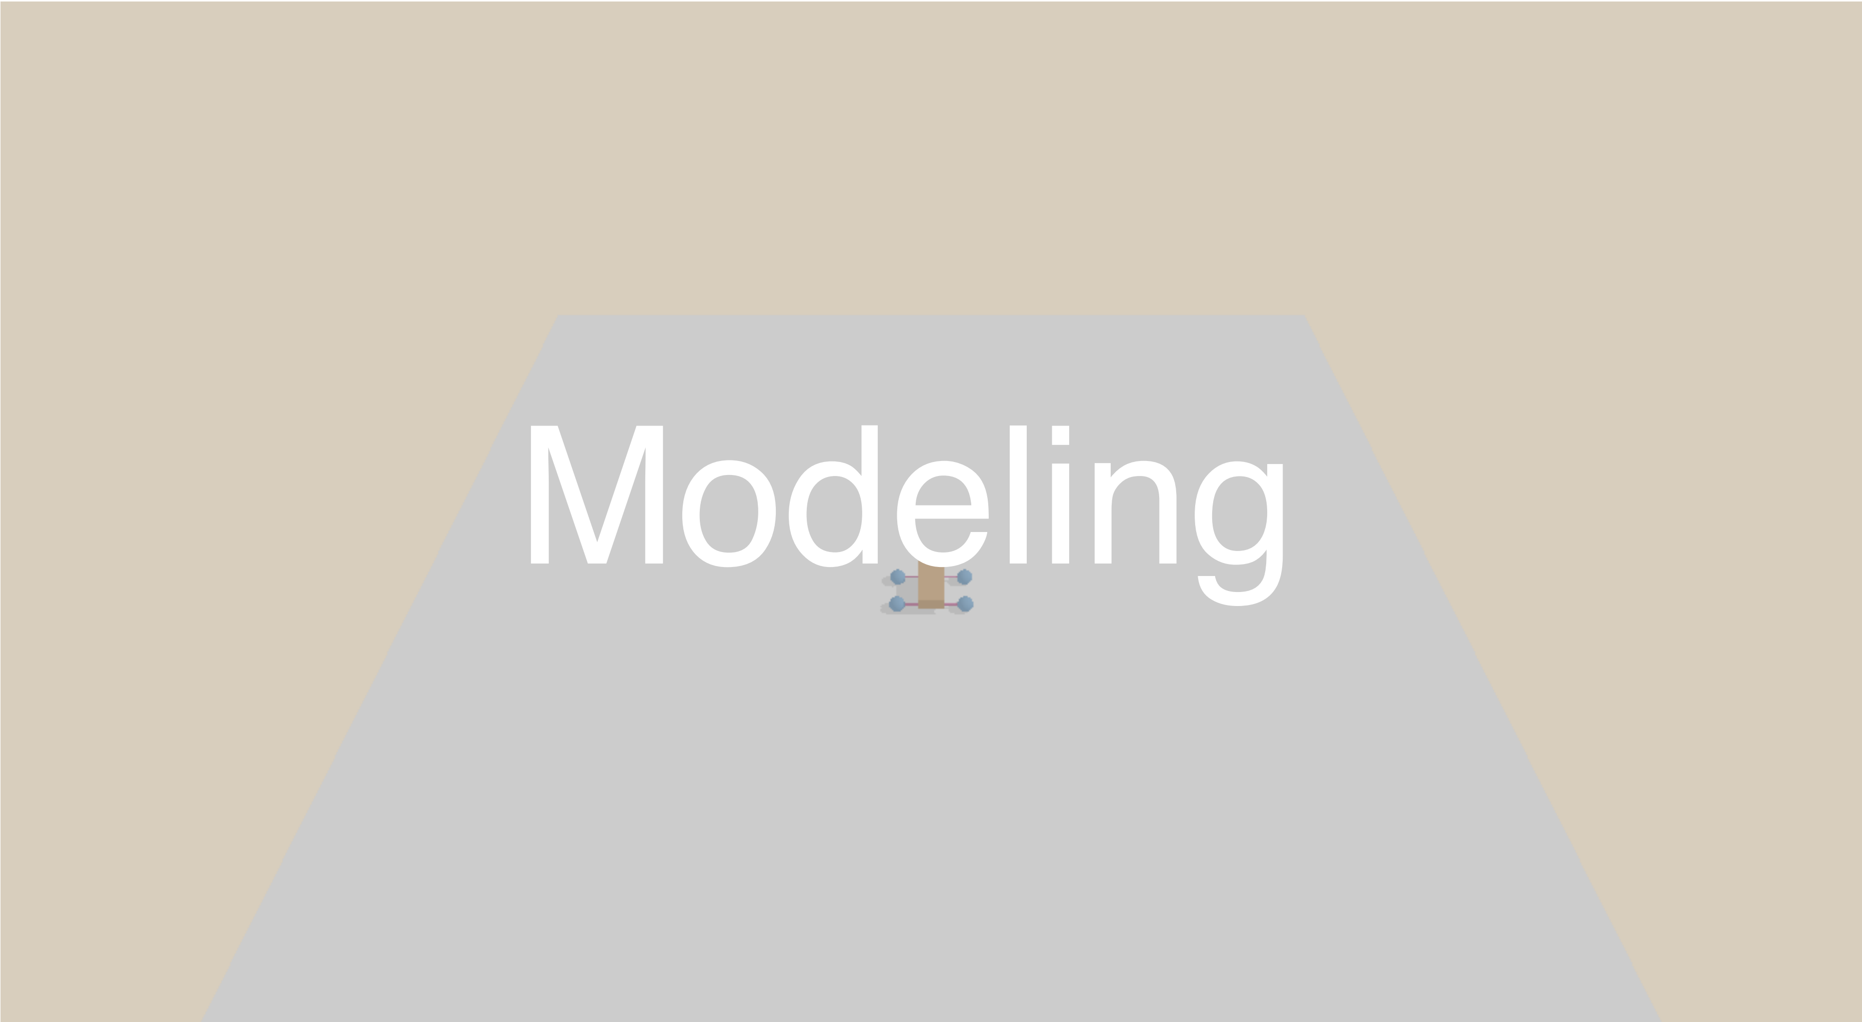

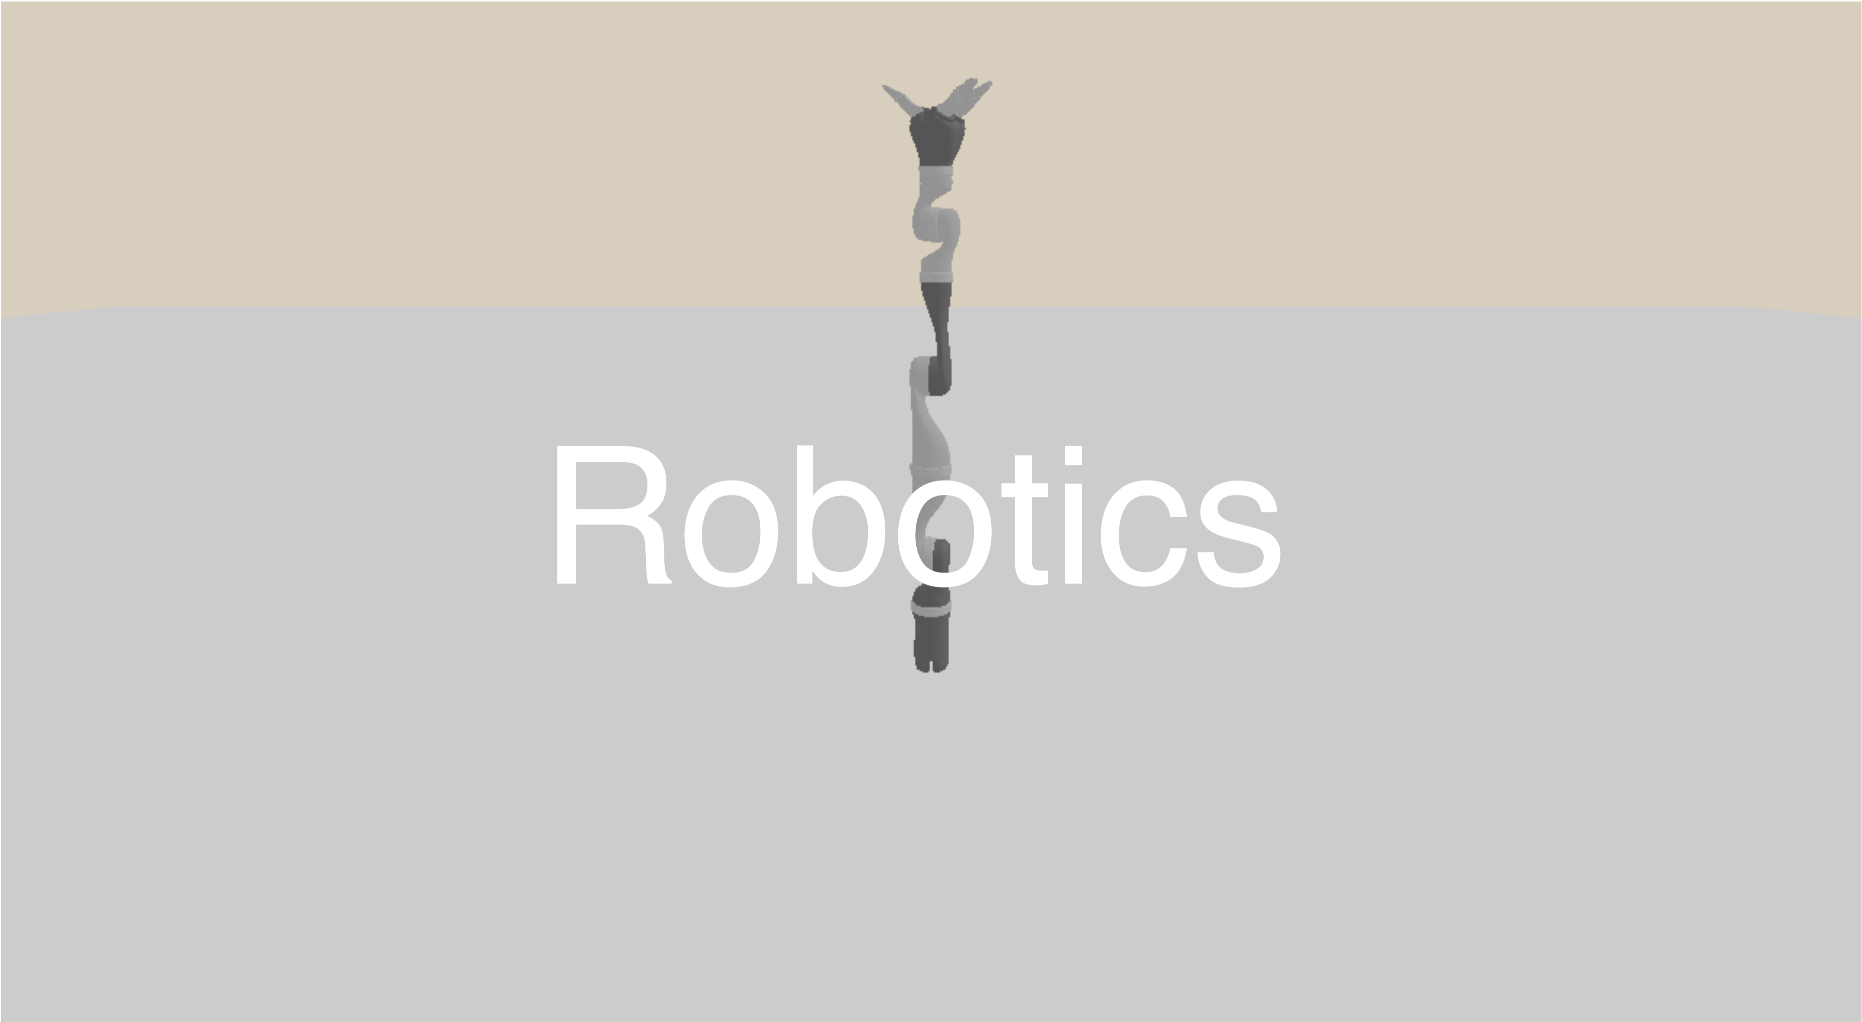

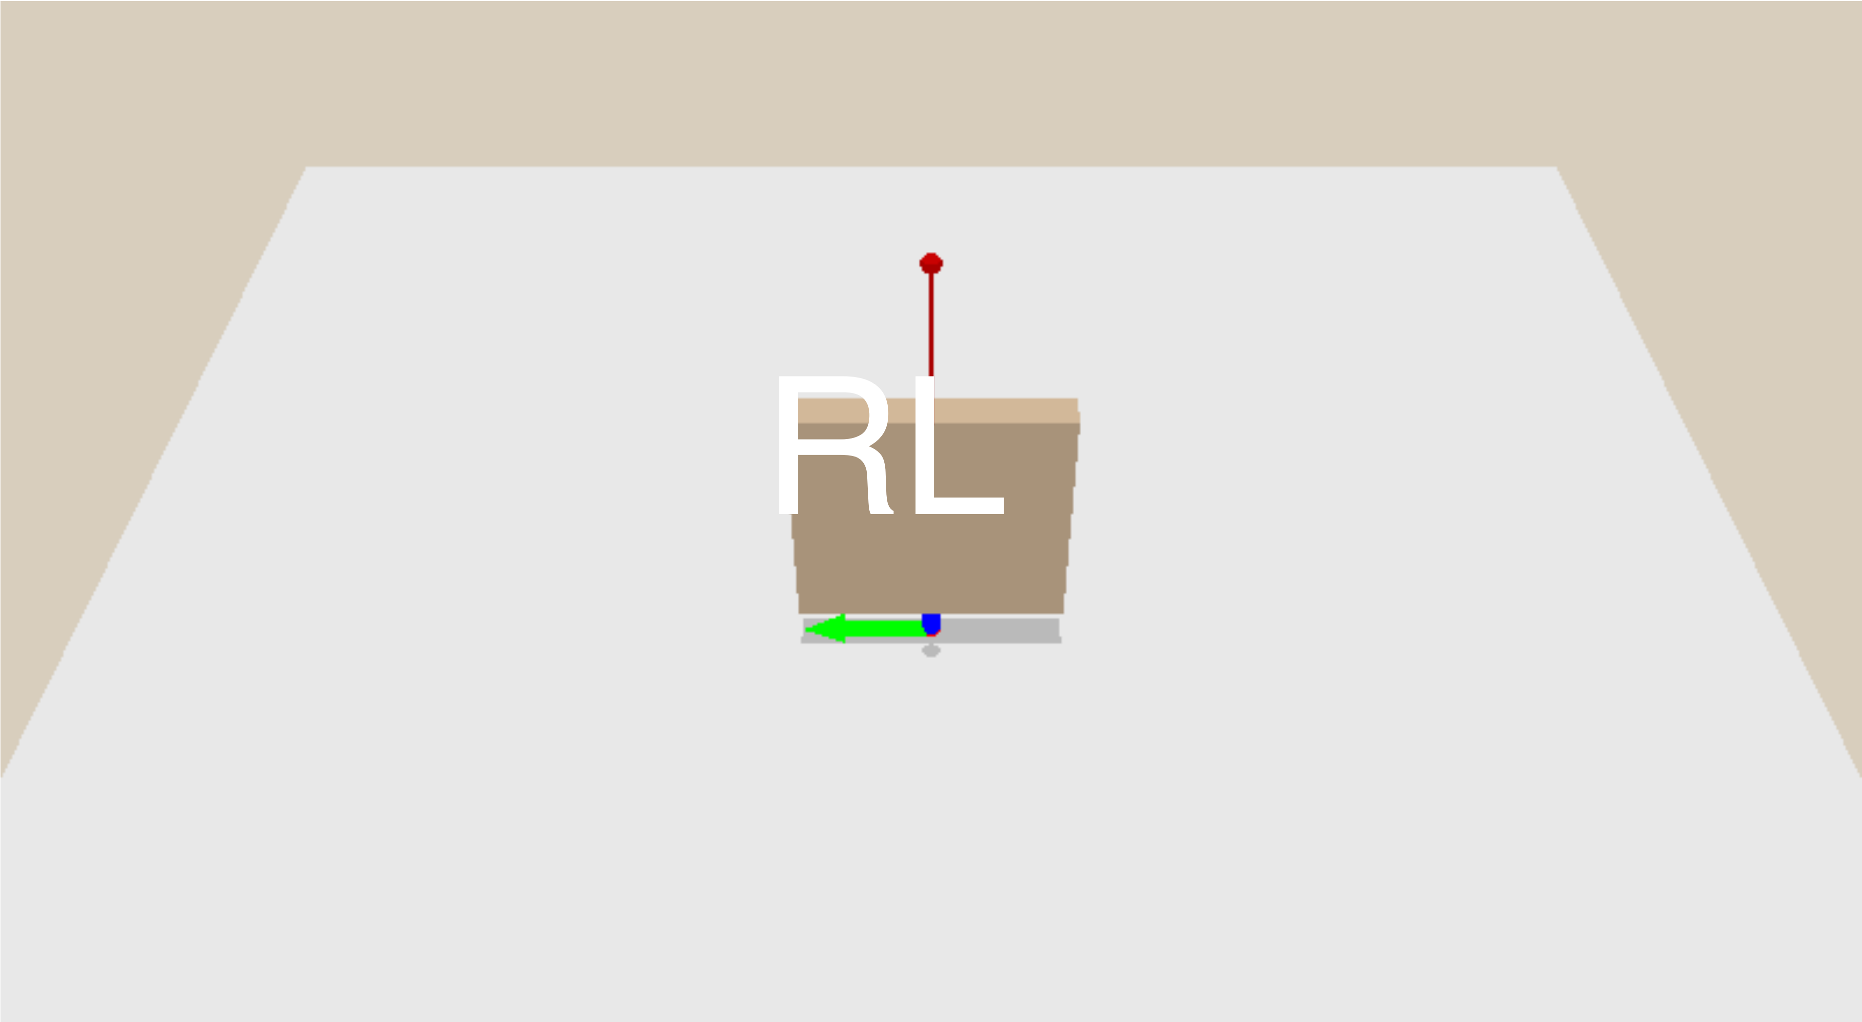

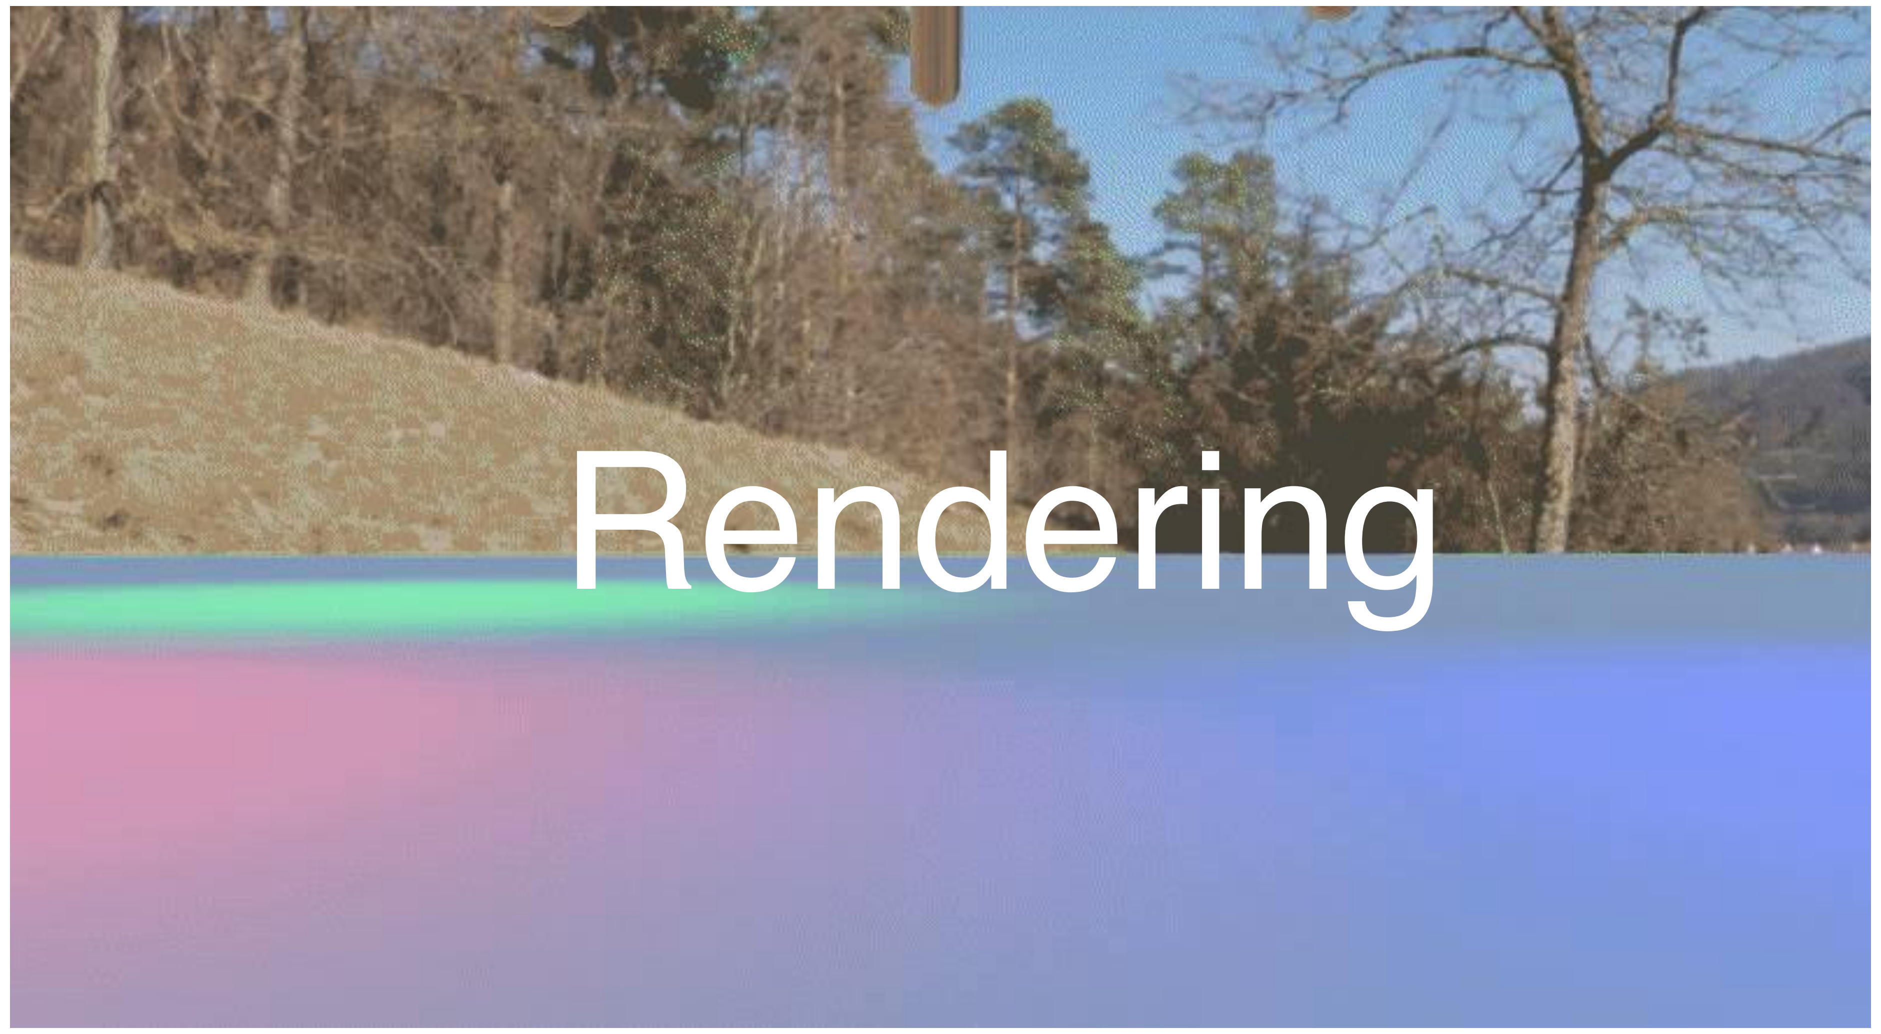

https://sapien.ucsd.edu/docs/latest/index.html

Powered by SAPIEN

ManiSkill Challenge

SAPIEN Manipulation Skill Challenge

UC San Diego
SU Lab

Berkeley
UNIVERSITY OF CALIFORNIA
RAIL

Manipulation Skills

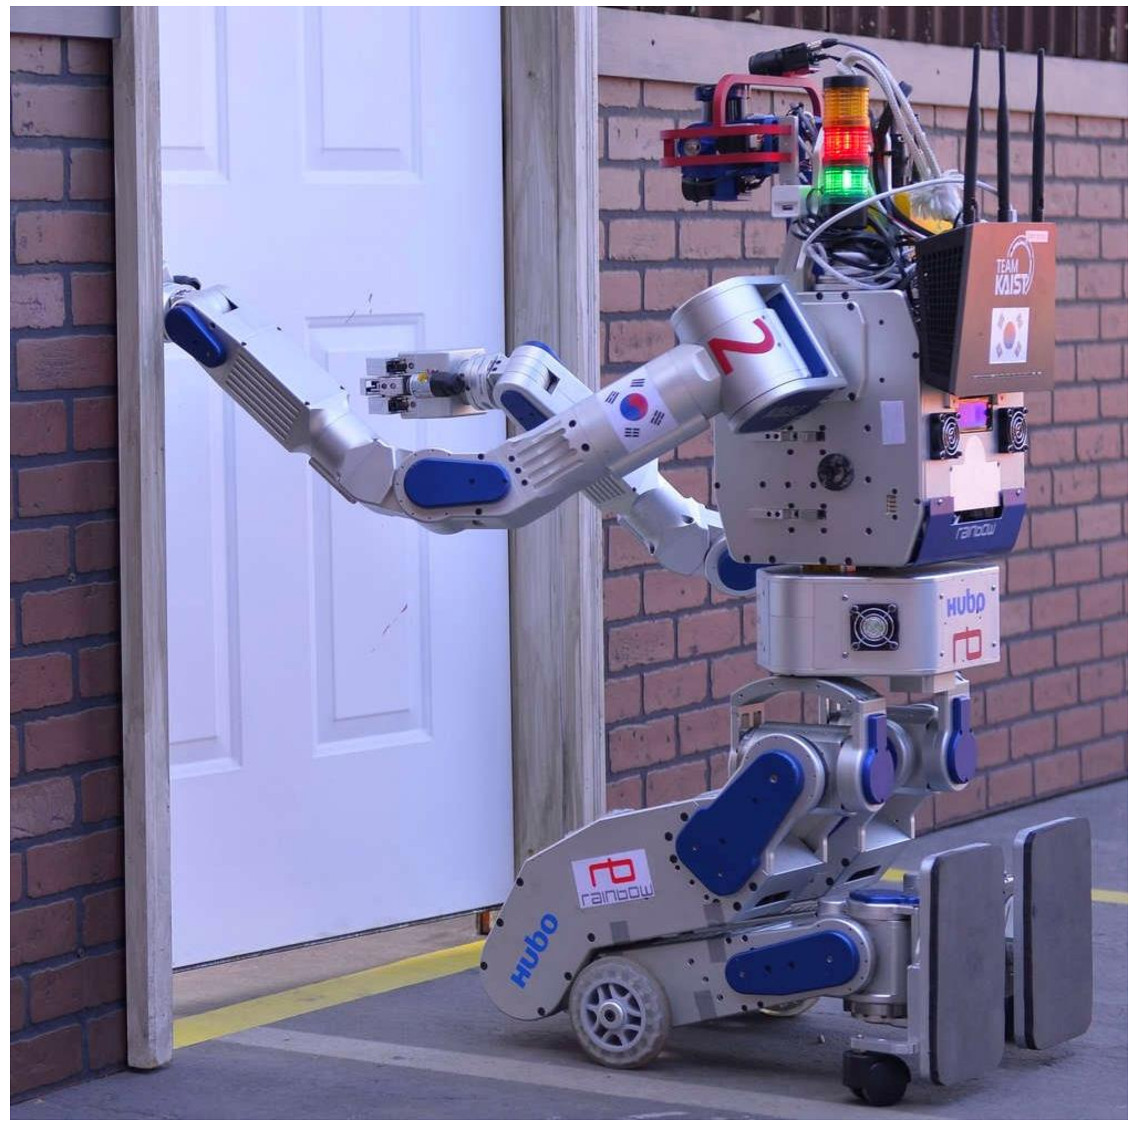

Opening

Pushing

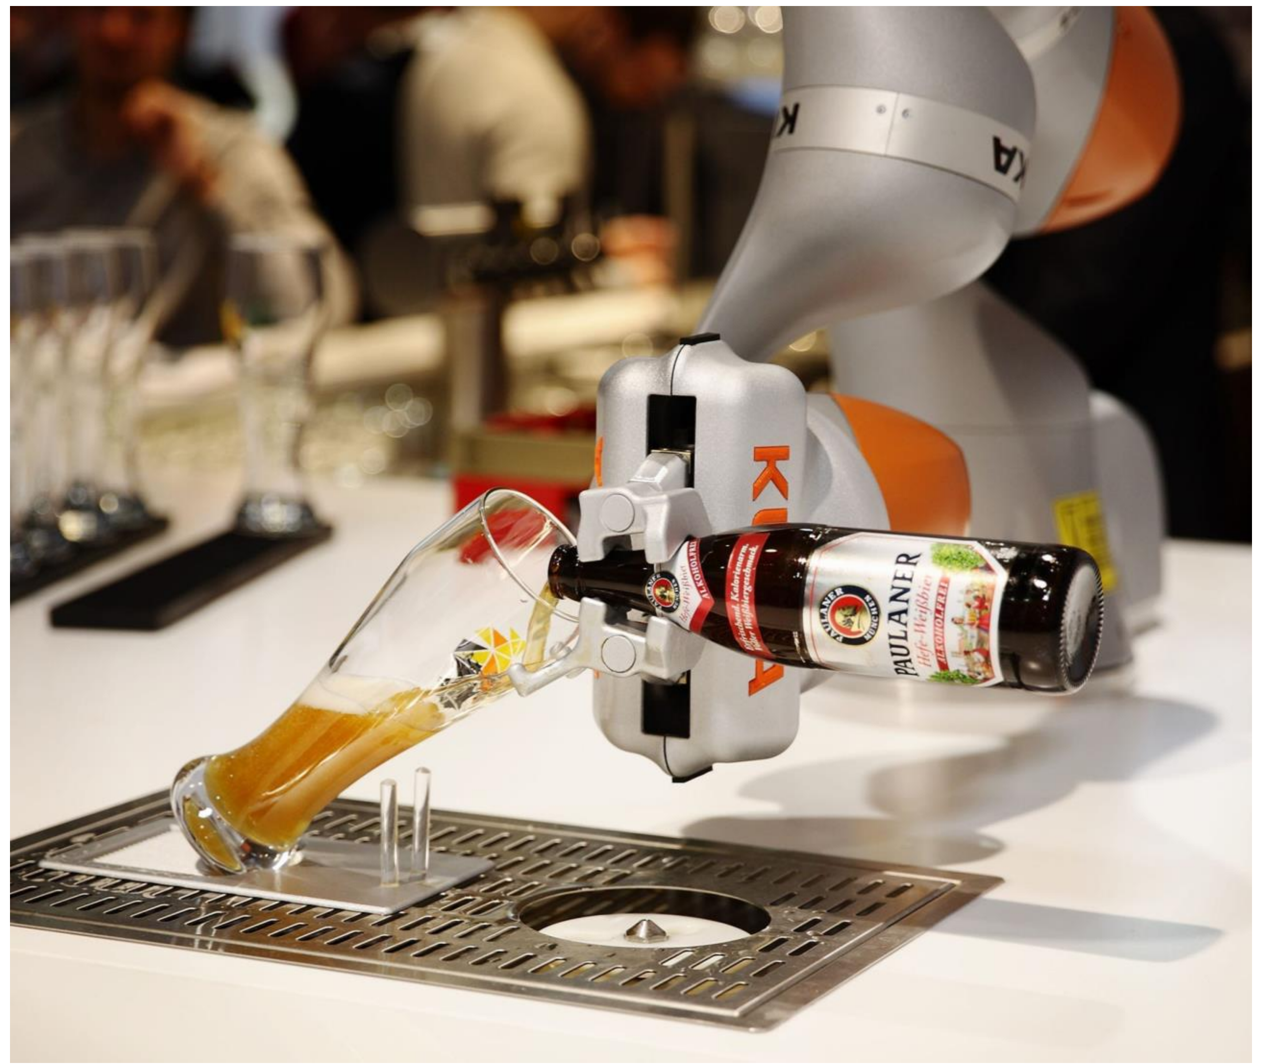

Pouring

Solve *Short-horizon* Tasks

- Goal: Move a chair out of a room

Skill 1: Open the door

Skill 2: Push the chair through the door

Key Features

- Object-centric skill organization
- Object-level generalizability metric
- Point cloud or RGBD input
- Easy to start with

Feature I: Object-Centric Skill Organization

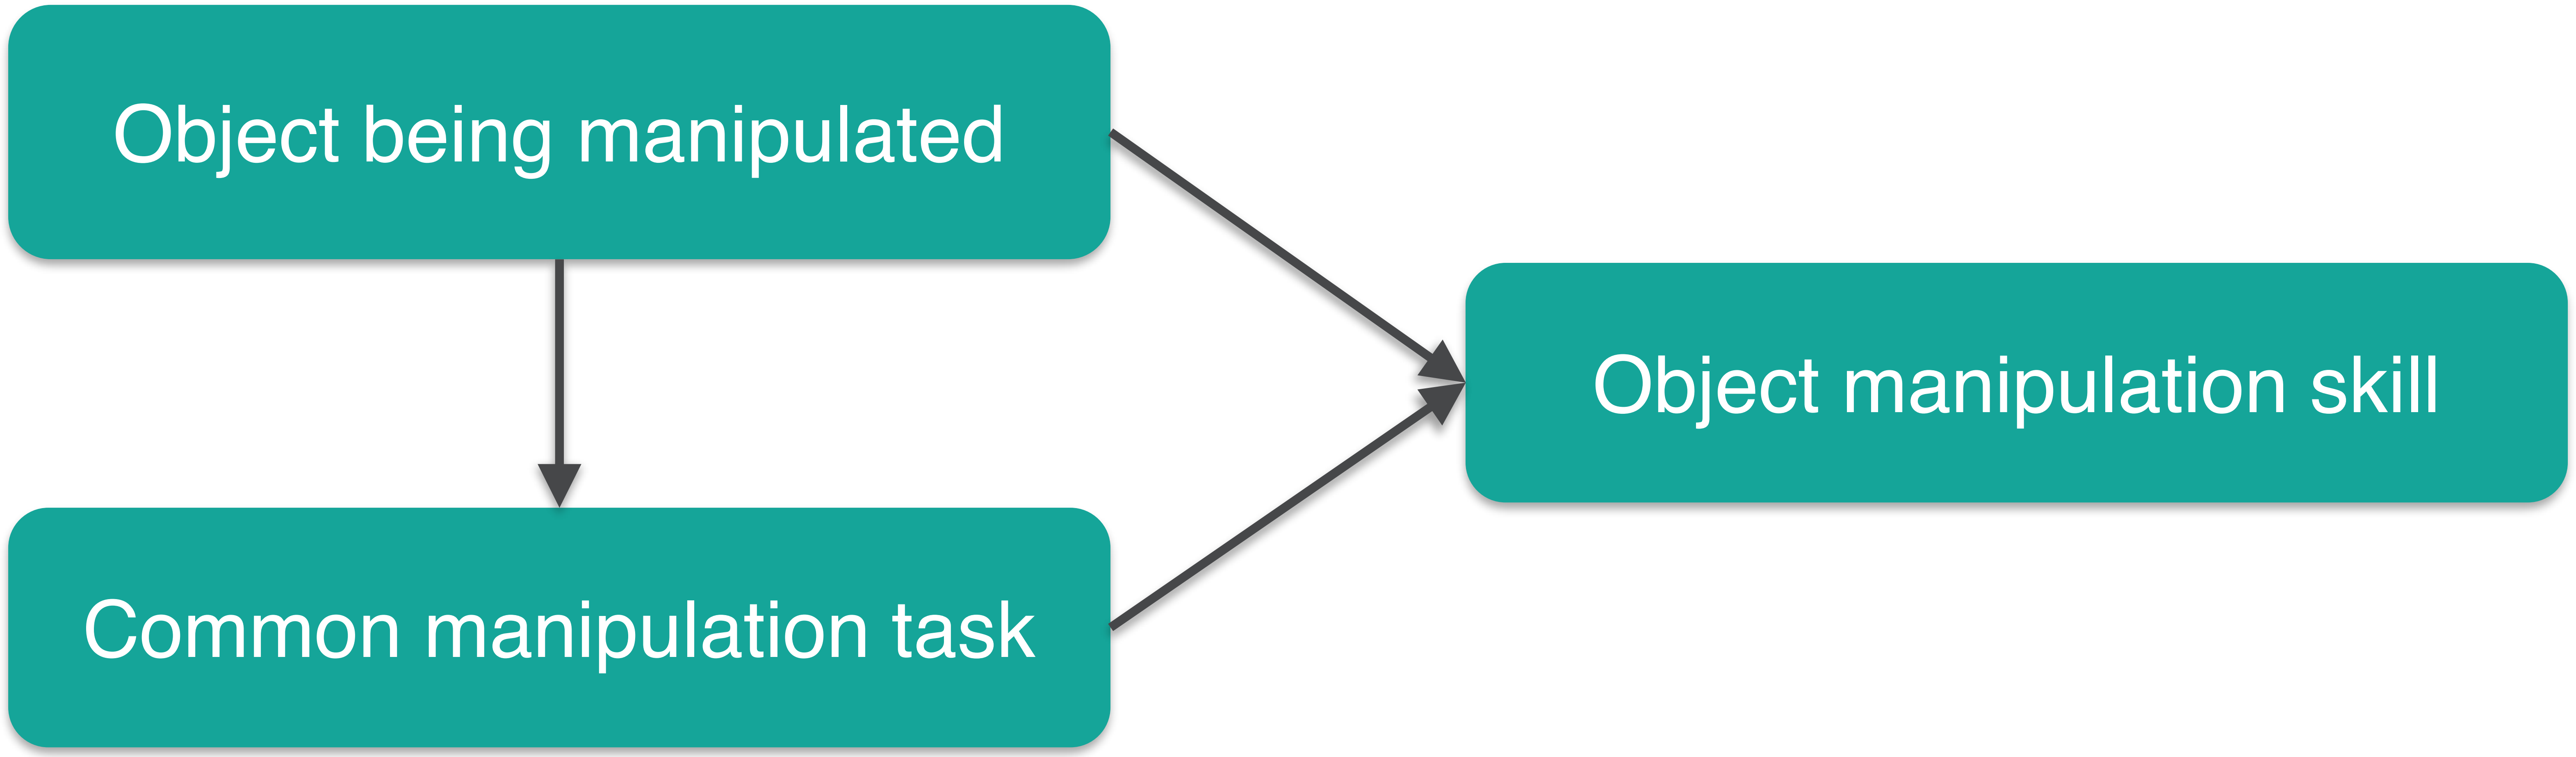

- Easy to leverage existing 3D datasets organized by semantic taxonomy
- Different from how robotics literature organizes skills
  - e.g., grasping, pushing, lifting
  - Our formulation restricts generalization at object-level

Feature II: Object-level Generalizability

- The agent needs to handle variations within an object category
- We train and test on different 3D objects.

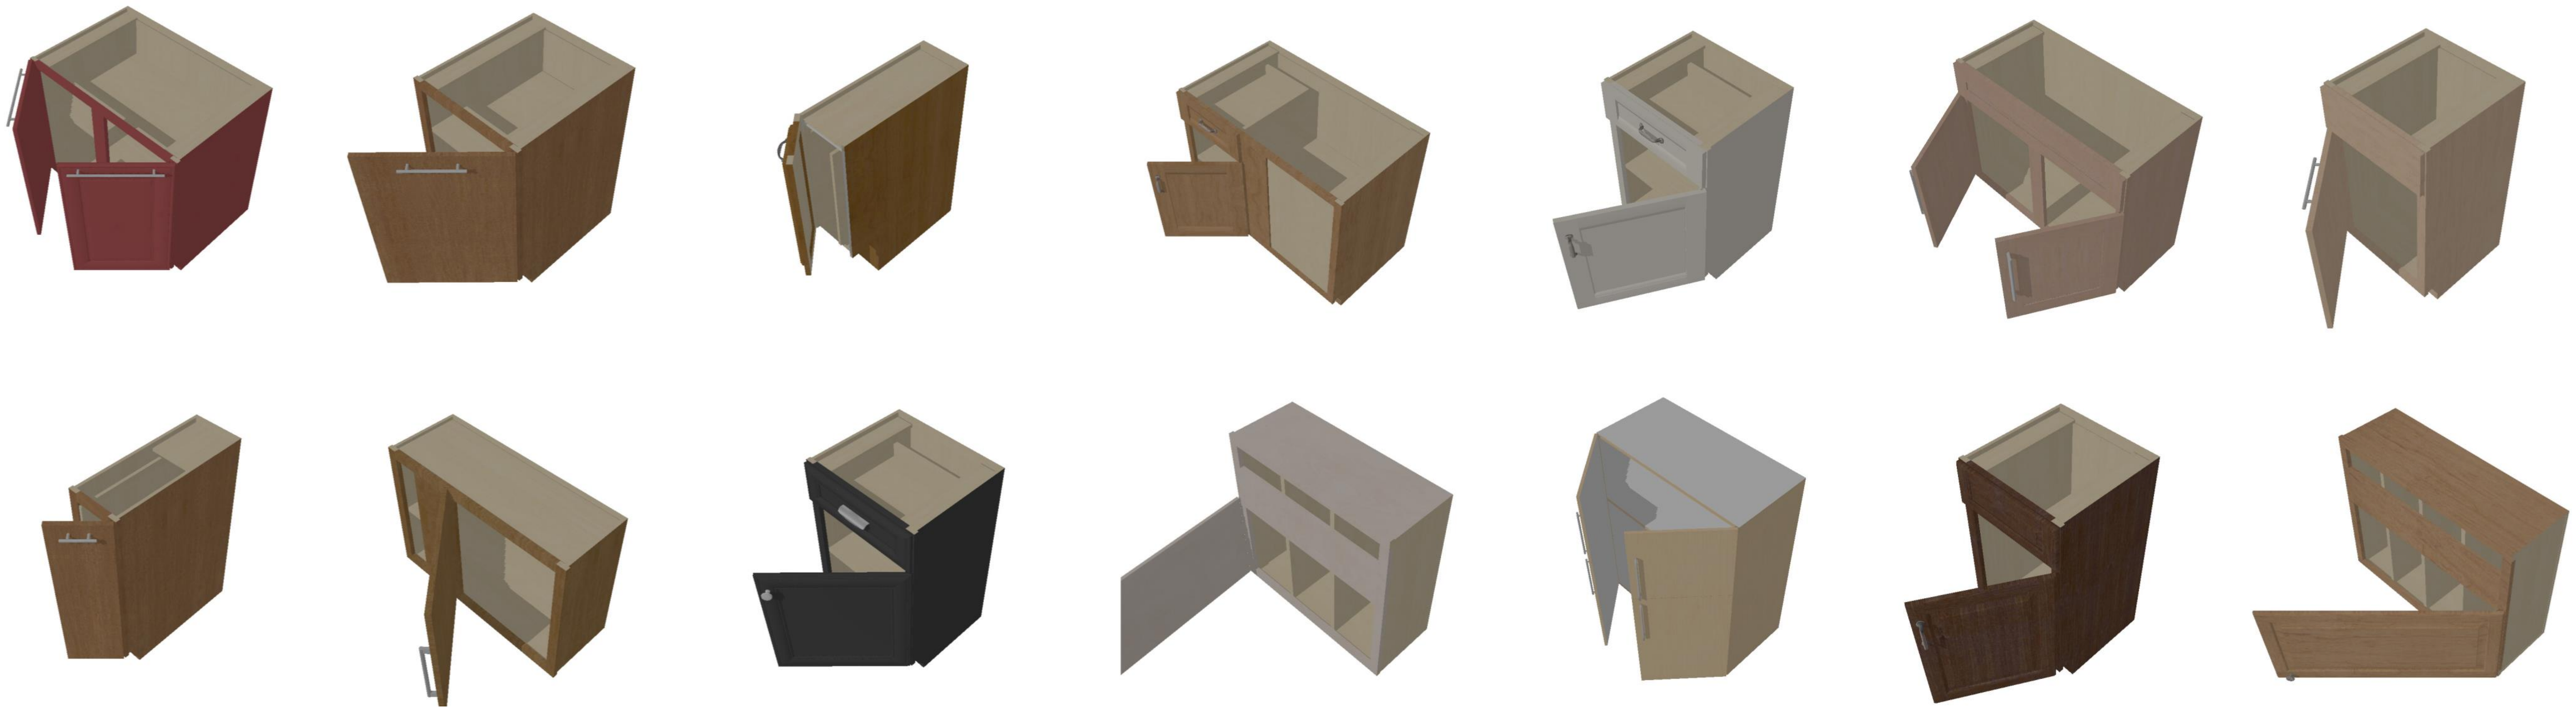

Benchmark for Generalizable Manipulation Skills

Our challenge focuses on 4 tasks

Each task targets a specific manipulation skill


Open cabinet

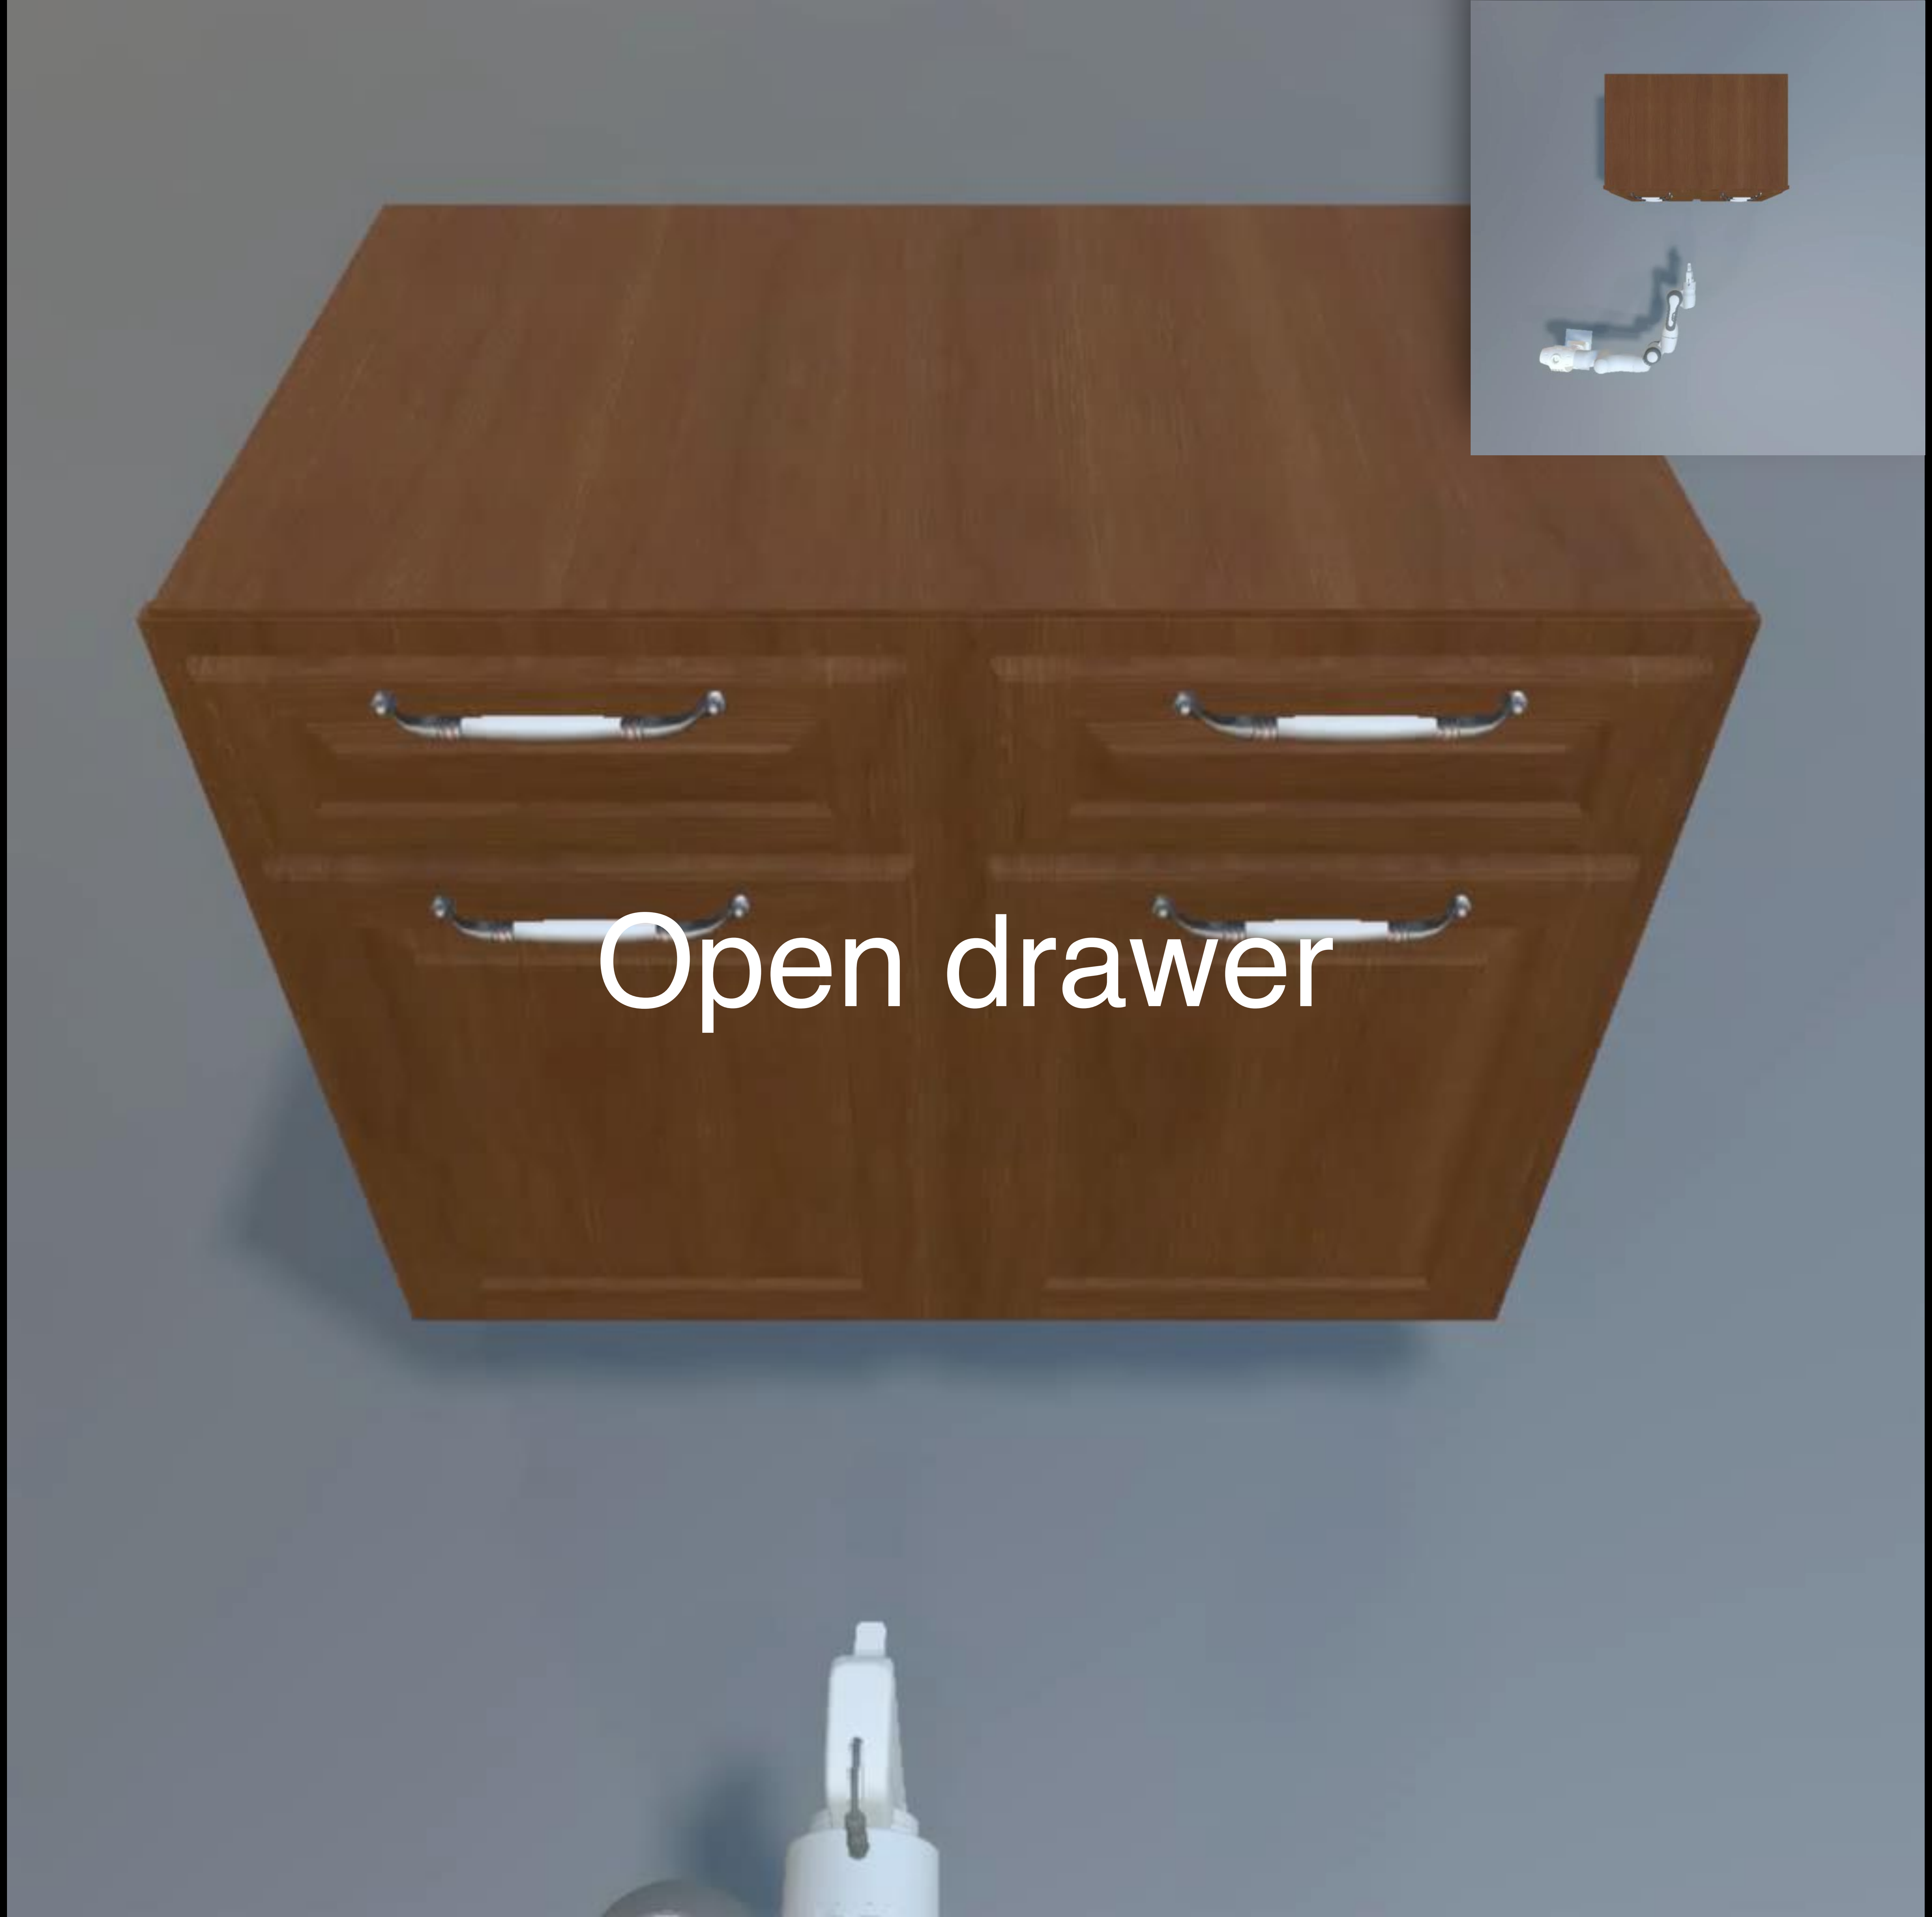
Open drawer


Push chair

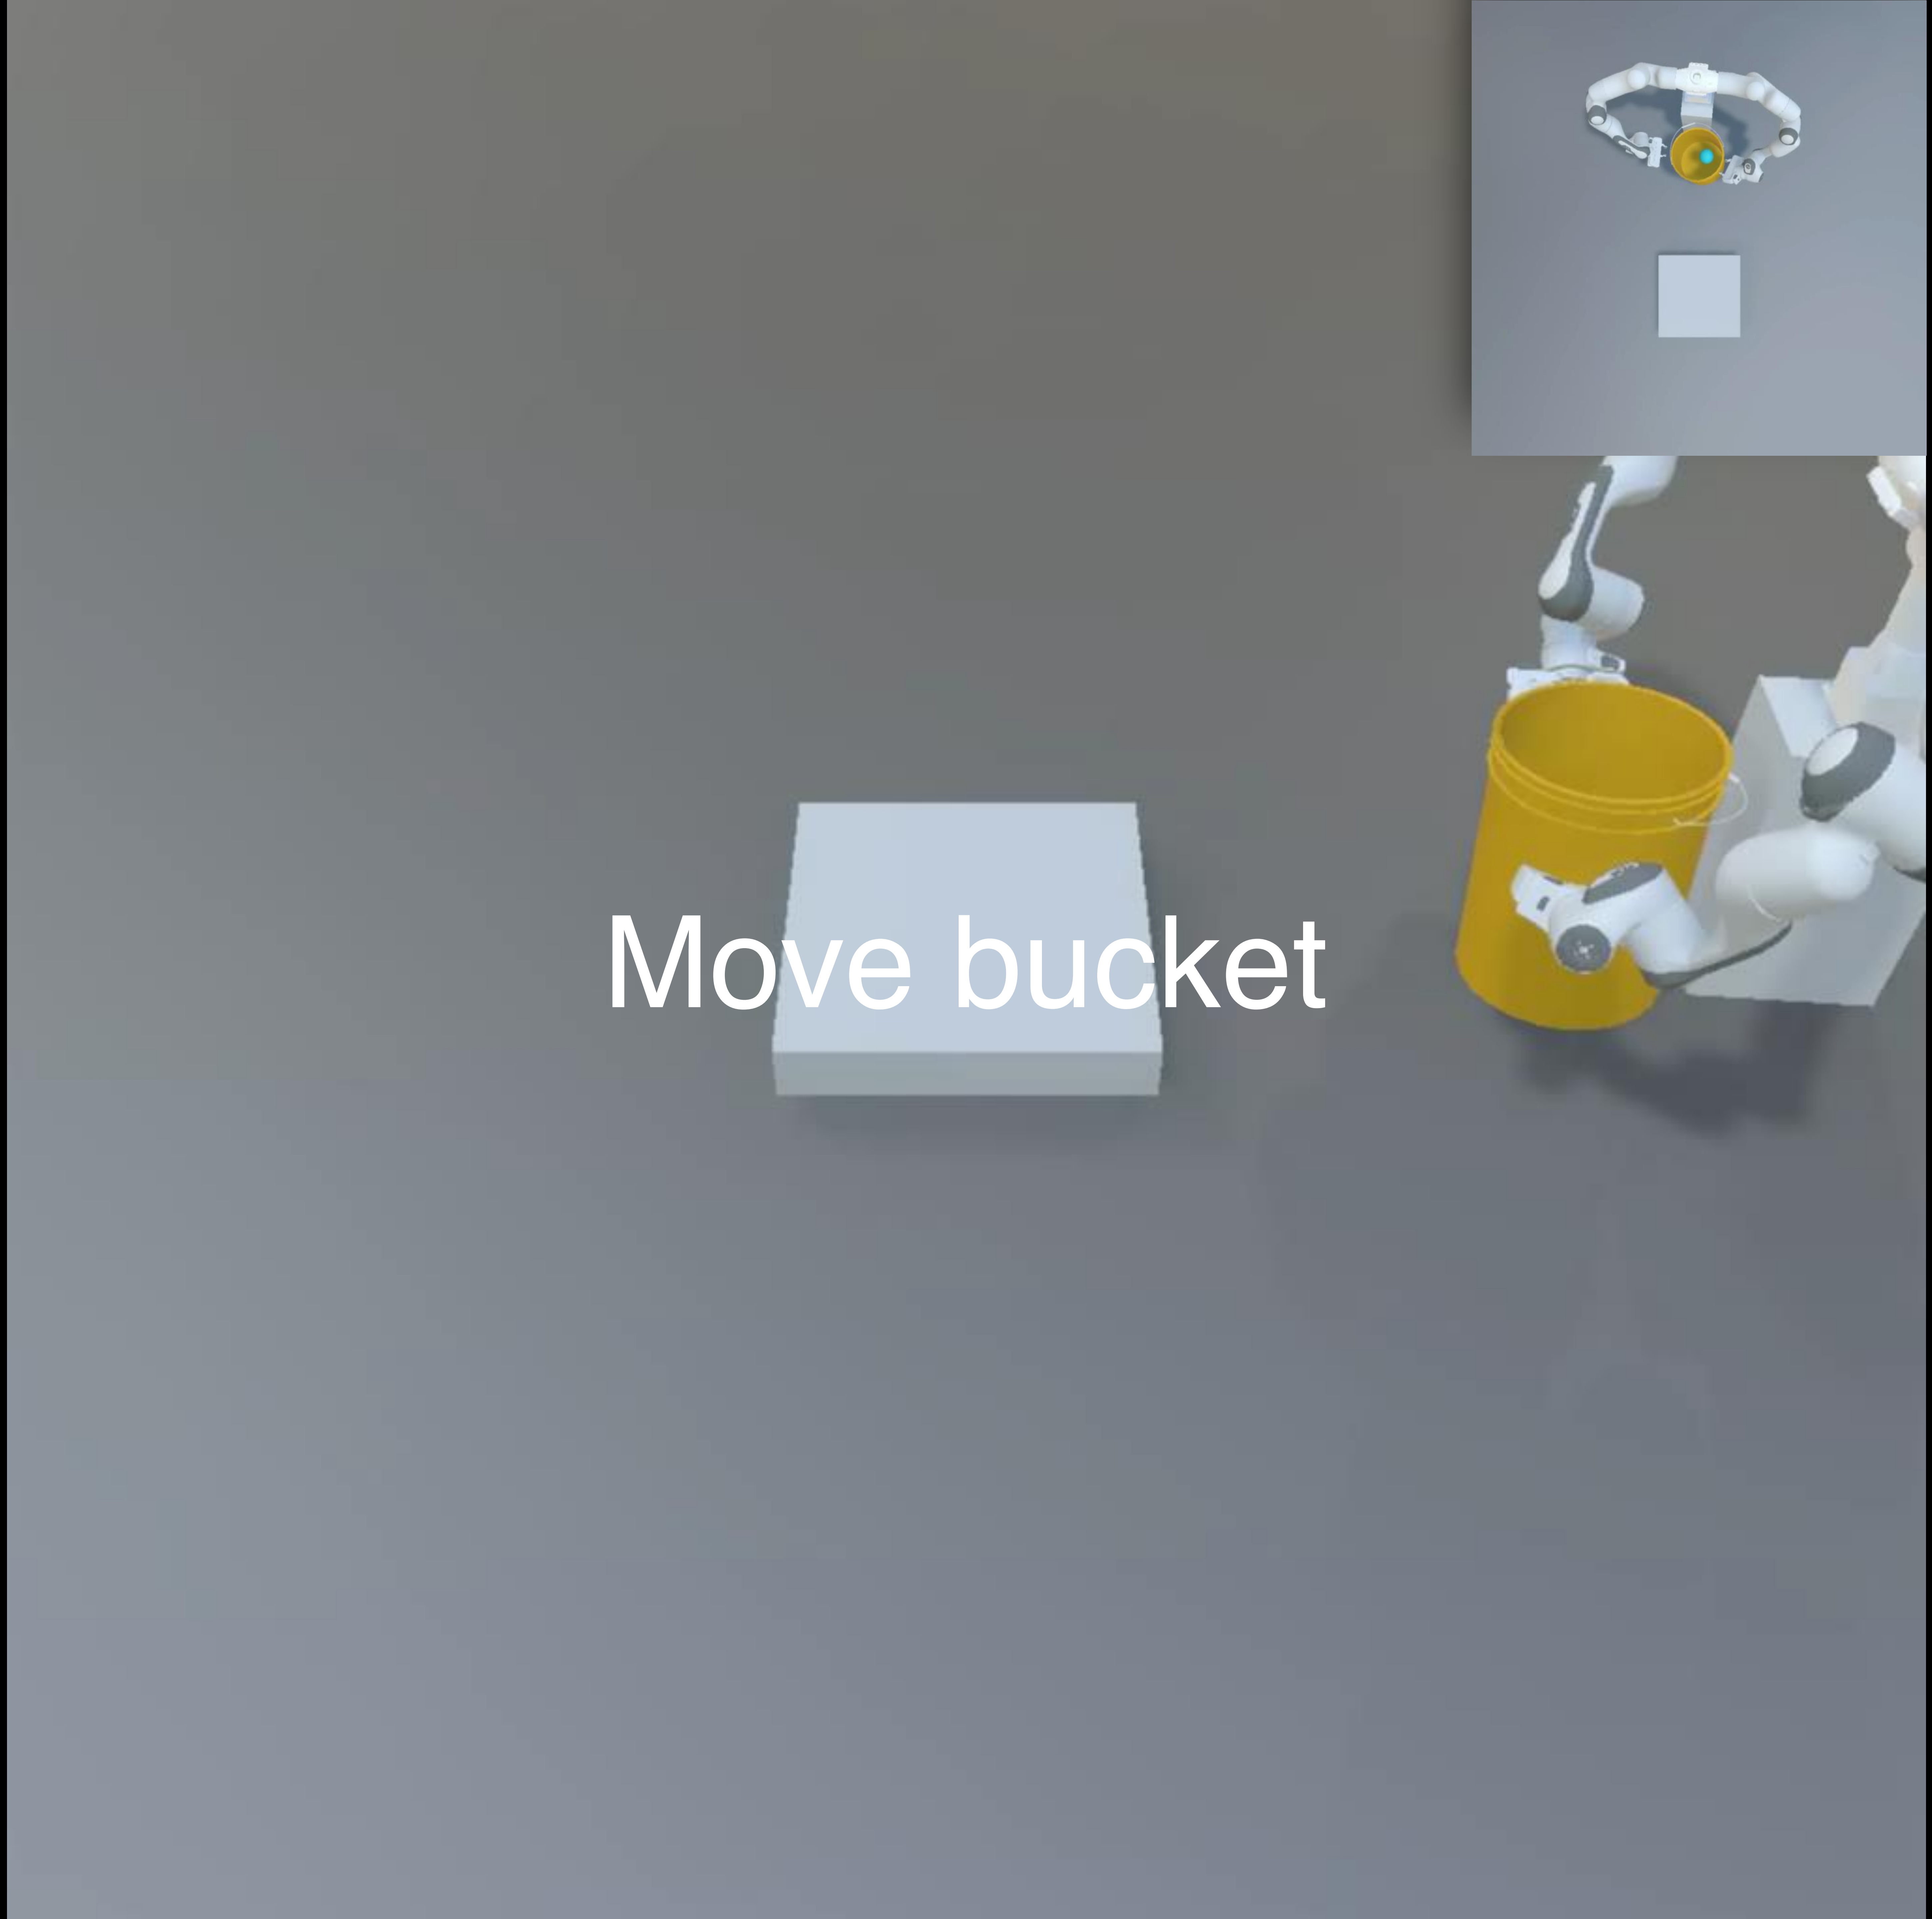
Move bucket

Statistics

- We have provided
  - 241 objects for the 4 tasks in total
  - 309 parts to be manipulated (e.g., a cabinet may have multiple doors)
- For each part to be manipulated, we will provide 300 demonstration trajectories

Feature III: RGBD/Point Cloud Input

Feature IV: Easy to Start With

- Demonstrations provided, and 3 tracks to choose from:
  - **No interaction track:** Learn from demonstration only, no interaction with the environment. (welcome, 3DV experts!)
  - **No additional annotation track:** Interaction allowed on top of demonstrations, no additional data and environment annotations. (welcome, RL experts!)
  - **No restriction track:** Do whatever you want to solve the problem. (welcome, robotics experts!)

ManiSkill Challenge Awards

- Each track
  - 1st place: **3,000 USD**
  - 2nd place: **2,000 USD**
  - 3rd place: **1,000 USD**

- Special award based on both team diversity and ranking
  - Up to 4 teams, **2,000 USD** in total

- Sponsored by Qualcomm AI research

Outline

- Low-cost interactive environments
  - SAPIEN++: Realistic Simulator for Manipulation Research
  - ManiSkill: SAPIEN Manipulation Challenge
- **Frameworks of Learning-from-Demonstrations**

Agent-Environment Interface

- At each step t the agent
  - Executes action A_t
  - Receives state S_t
  - Receives scalar reward R_t
- The environment
  - Receives action A_t
  - Emits state S_{t+1}
  - Emits scalar reward R_{t+1}

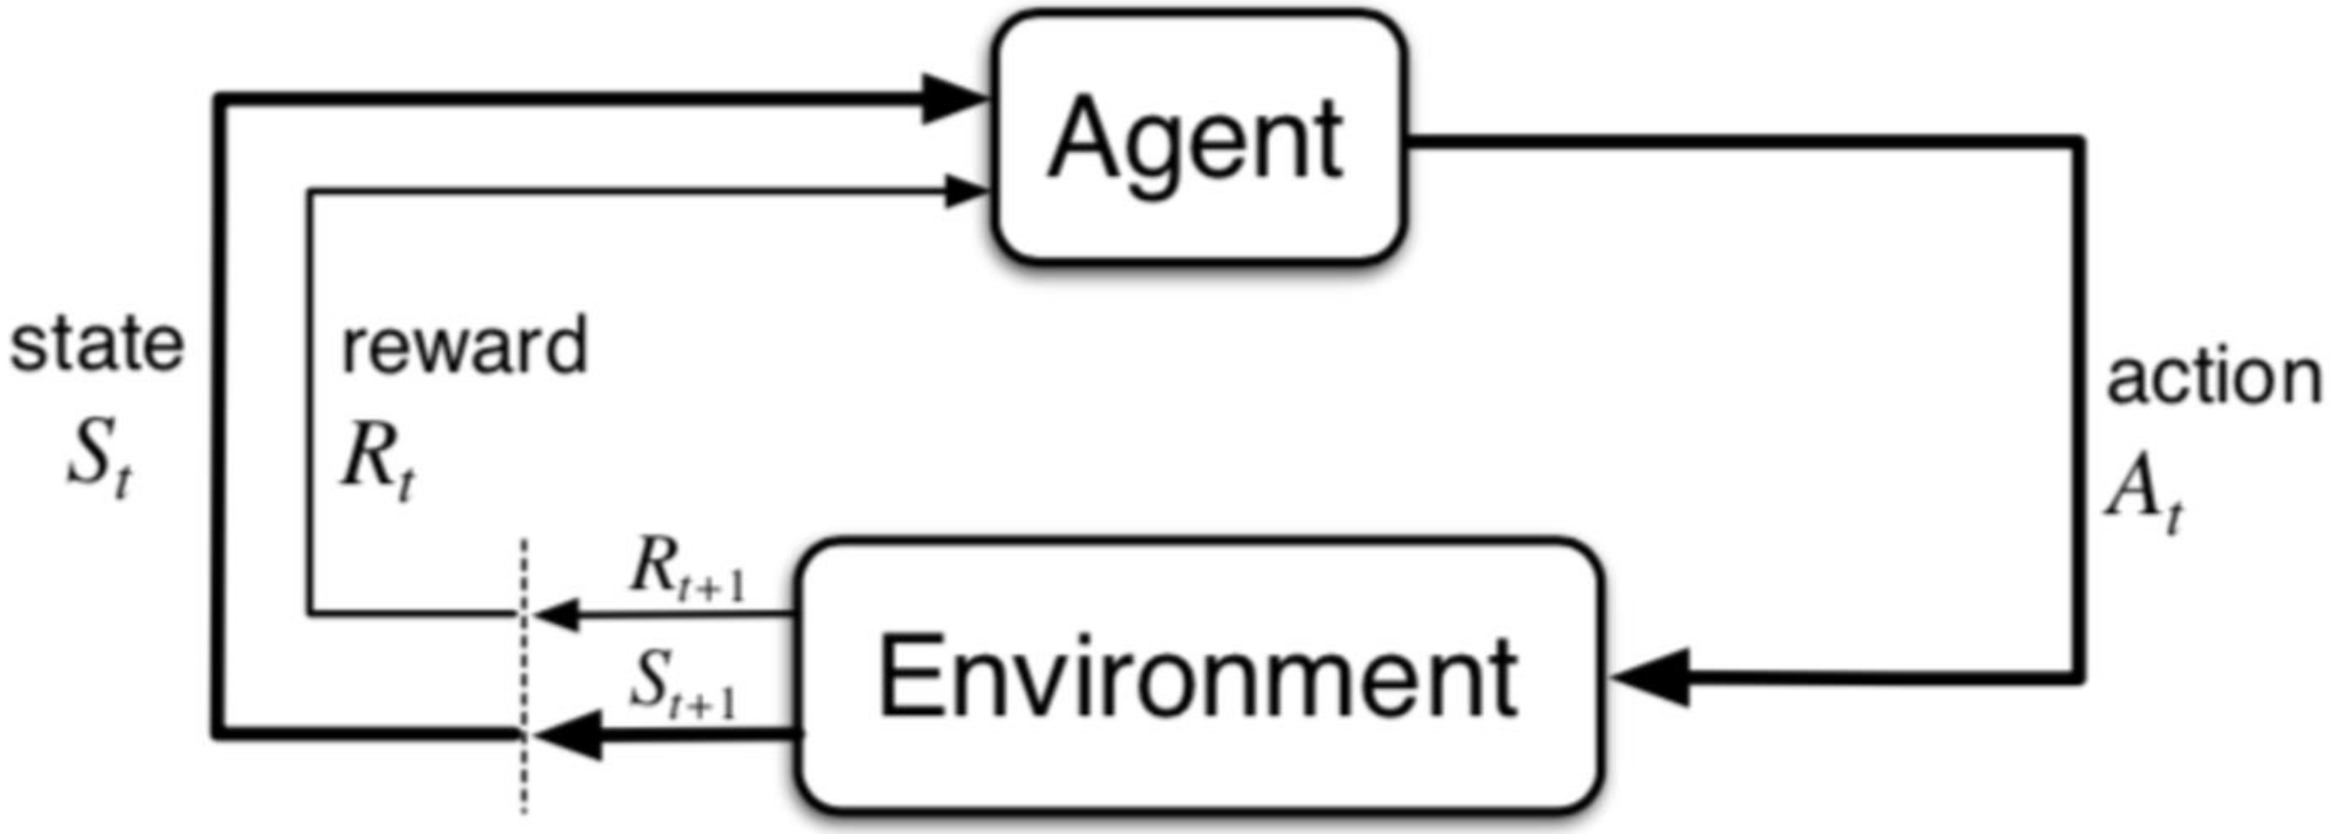

Figure 3.1: The agent–environment interaction in a Markov decision process.

Source: https://haosulab.github.io/ml-for-robotics/SP21/lectures/html/L12_RL_Framework_I.html#/step-6

Example: Take a Step in ManiSkill

joint 1	+0.3
joint 2	+0.8
joint 3	-0.5

+ +0.5

Current observation | Action | Next observation | Reward

$$(\; s_t, \; a_t, \; s_{t+1}, \; r_t \;)$$

s_t: visual input + robot state

s_{t+1}: visual input + robot state

Policy

- A policy is the agent's behavior

- It is a map from state to action, e.g.,
  - Deterministic policy: $a = \pi(s)$
  - Stochastic policy: $\pi(a|s) = \Pr(A_t = a|S_t = s)$

Demonstrations

- We execute a policy in the environment, until we reach the maximum step number n, or have completed the task

- This will generate a trajectory: $\{(s_t, a_t, s_{t+1}, r_t)\}_{t=1}^{n}$

- In ManiSkill, we provide 300 demonstrations for each part to be manipulated, as trajectories of up to 200 steps.

Basic Frameworks of Learning from Demonstrations

- Imitation learning

- Offline reinforcement learning

Imitation Learning

- Behavior cloning (BC). A straight-forward but very effective way:

$$\underset{\theta}{\text{minimize}} \quad \sum_i \|a_i - \boxed{\pi_\theta(s_i)}\|^2$$

neural network,
e.g., point cloud network

- The network needs to generalize across states that are
  - in the same environment (e.g., different gripper starting poses for manipulating the same object)
  - across environments for different objects

Influence of Network on ManiSkill (I)

- Vanilla PointNet

Point Cloud

XYZ = (N×3)
RGB = (N×3)
SEG = (N×K)

For example, for cabinet door opening, we label

- the door to be opened
- the handle on the door
- robot
- background points

Concatenate to each point

Robot State

*SEG means segmentation mask
e.g., the cabinet door needs to be opened

Influence of Network on ManiSkill (I)

- Vanilla PointNet

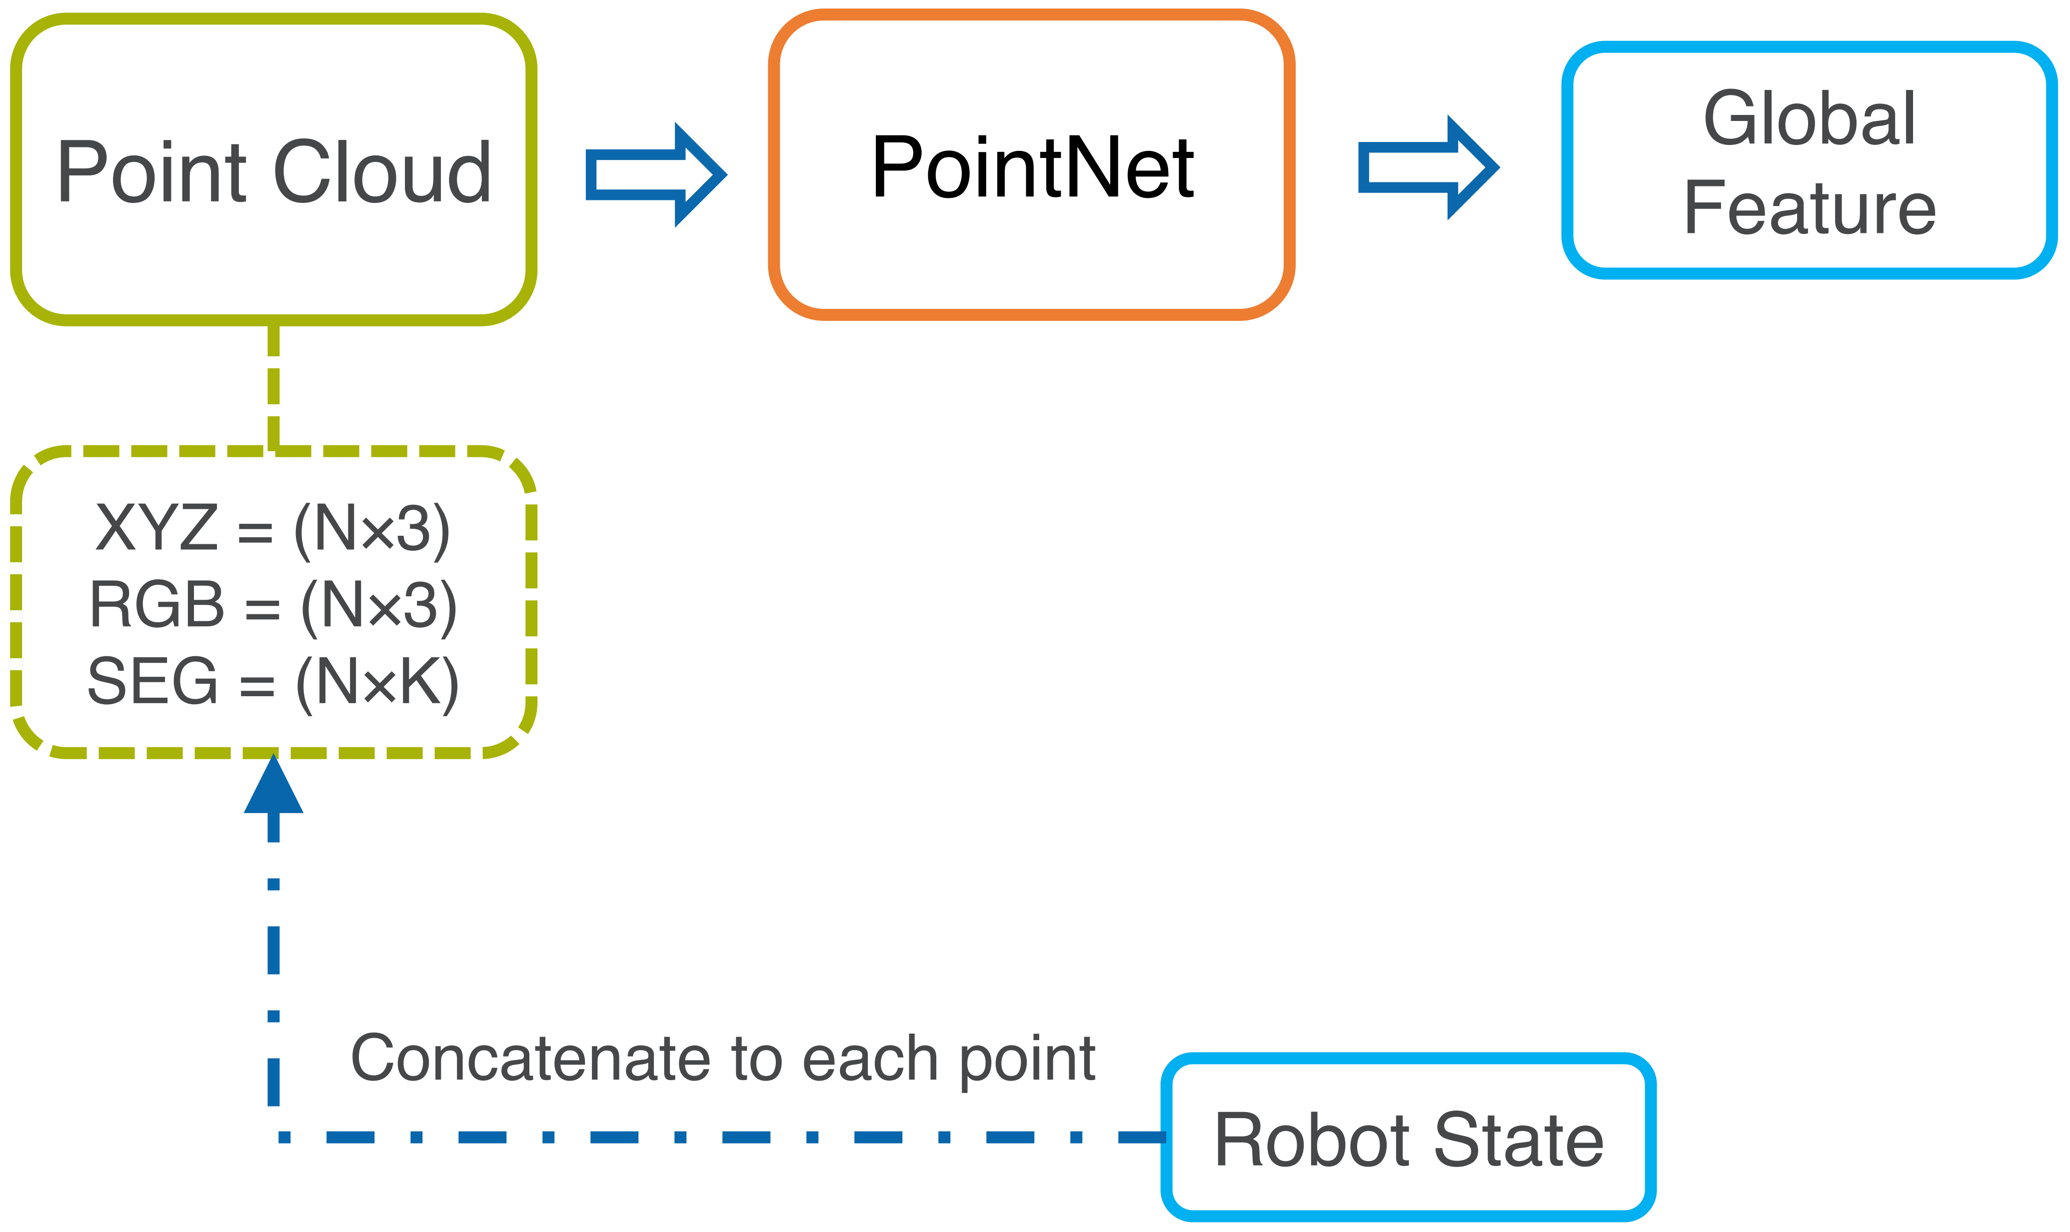

*SEG means segmentation mask
e.g., the cabinet door needs to be opened

Influence of Network on ManiSkill

- Vanilla PointNet

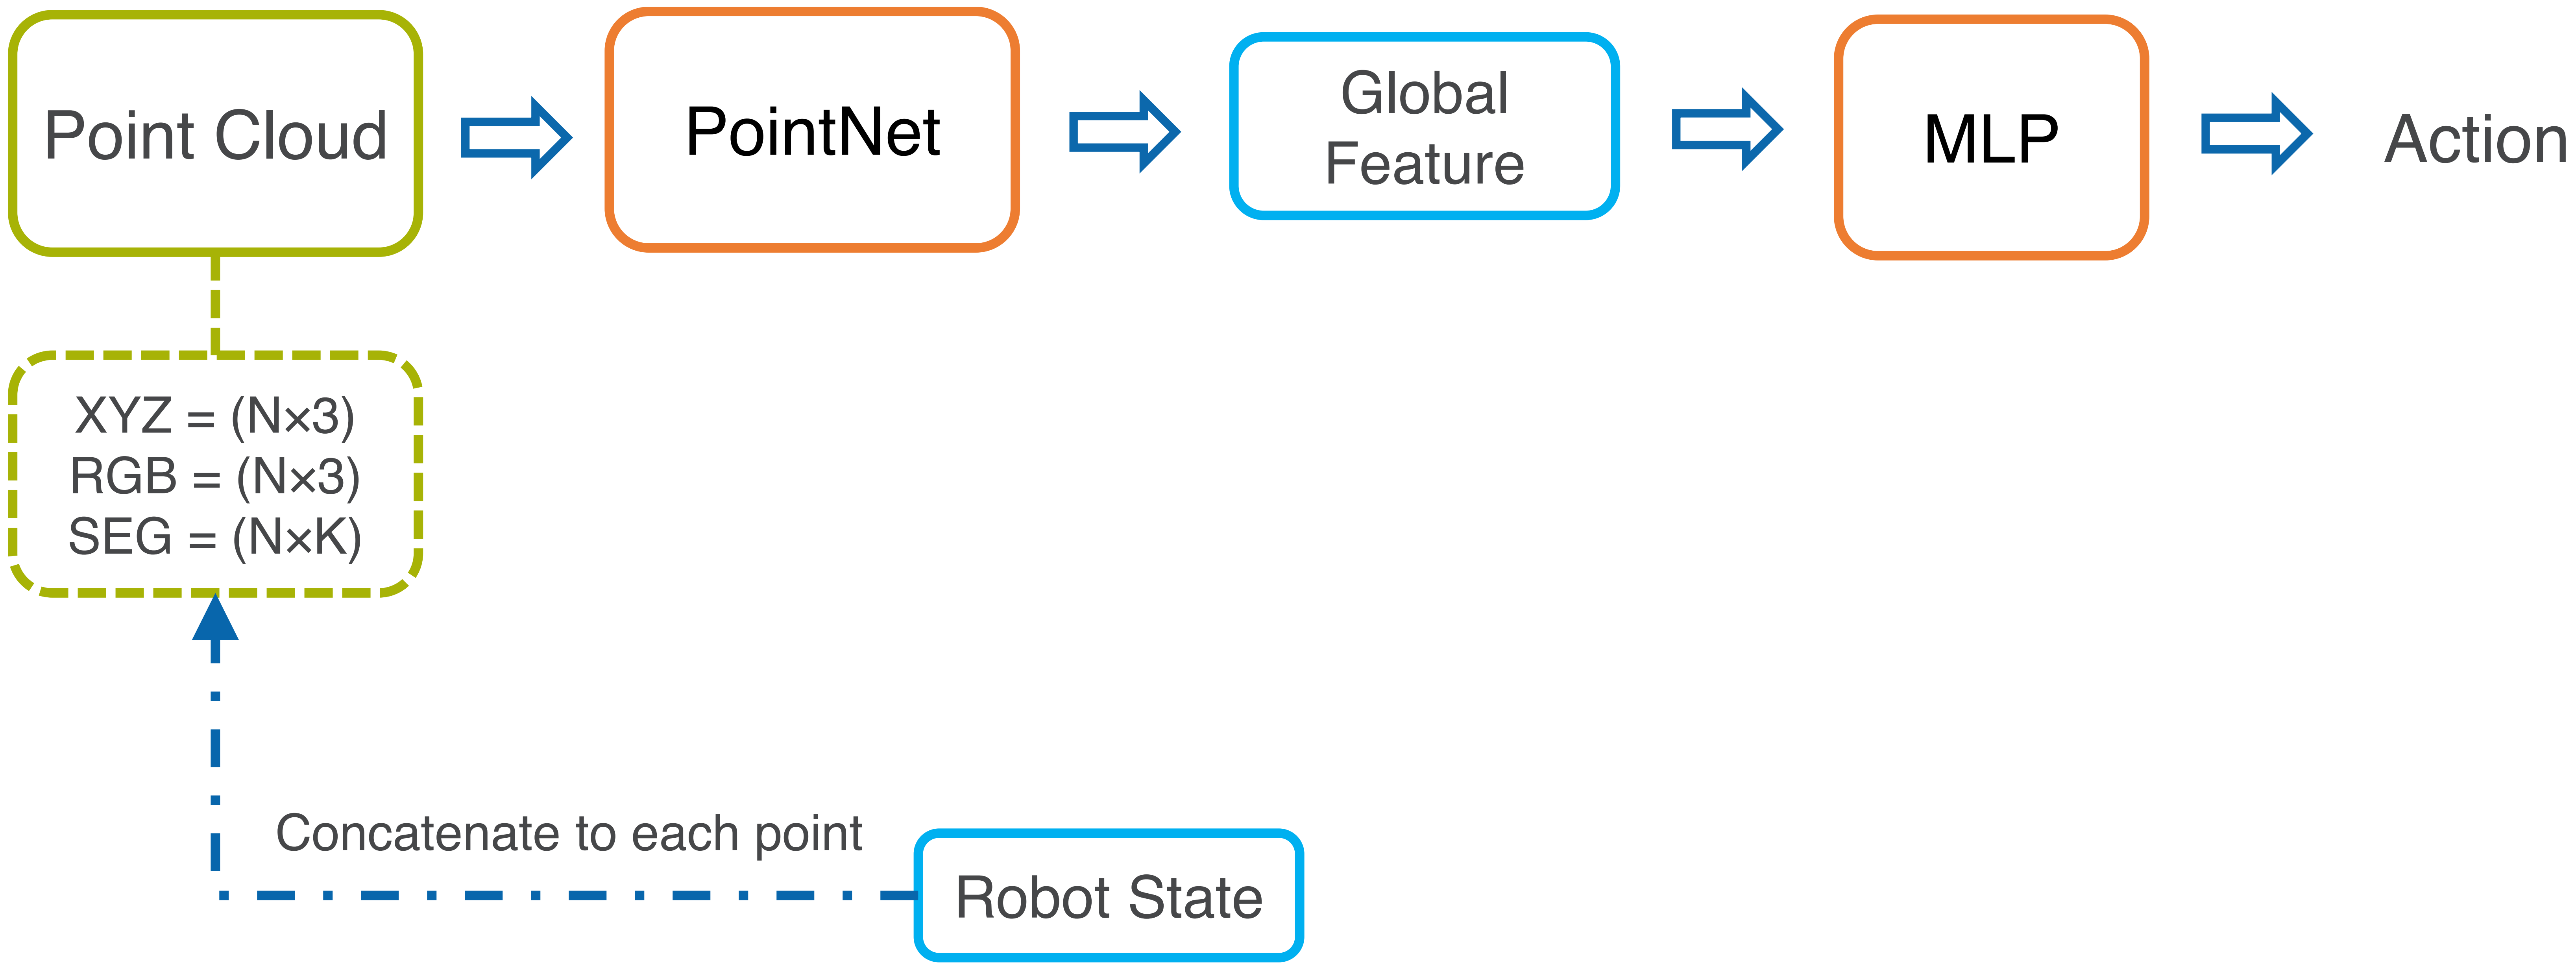

*SEG means segmentation mask
e.g., the cabinet door needs to be opened

Influence of Network on ManiSkill

- PointNet + Transformer

Input:
Point Cloud + Masks

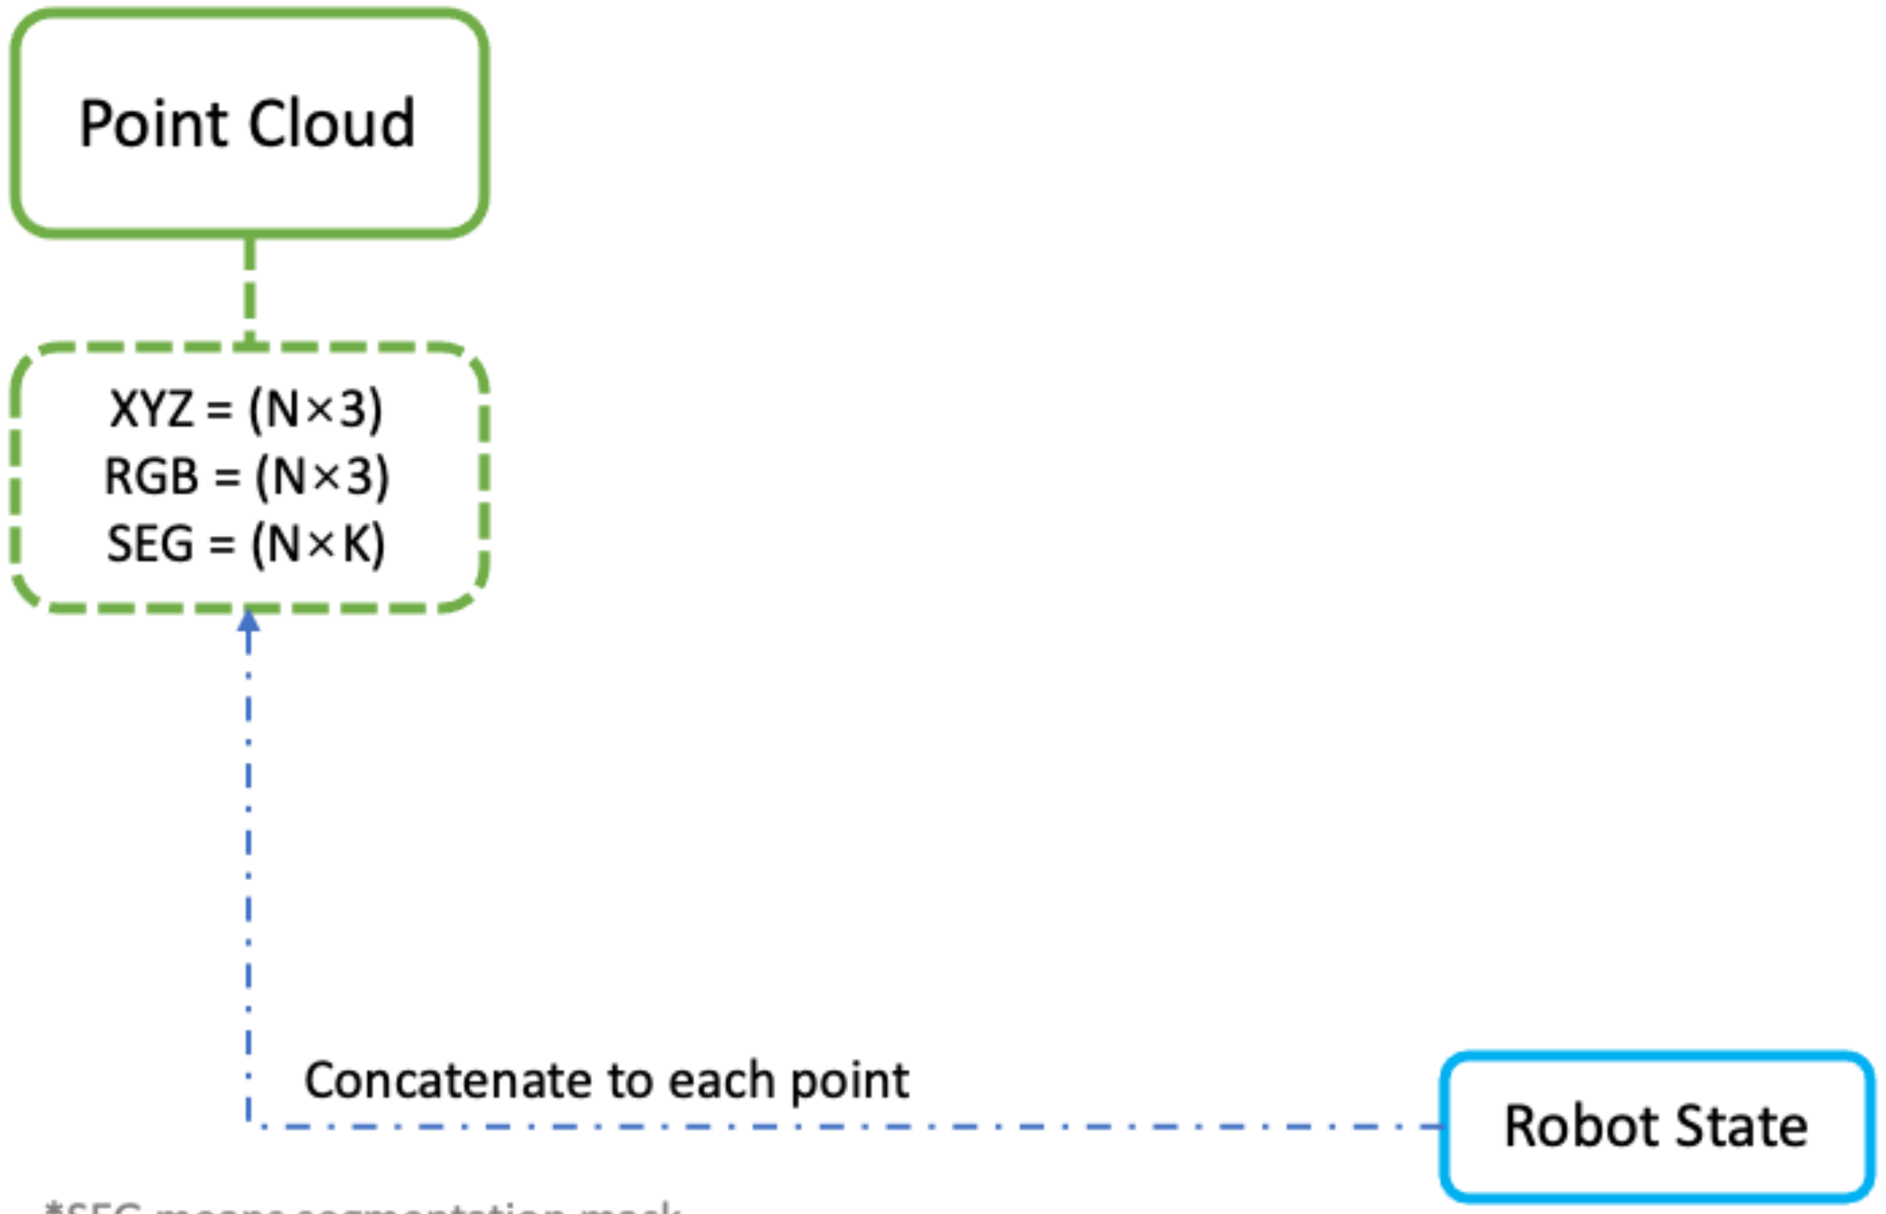

*SEG means segmentation mask

Influence of Network on ManiSkill

- PointNet + Transformer

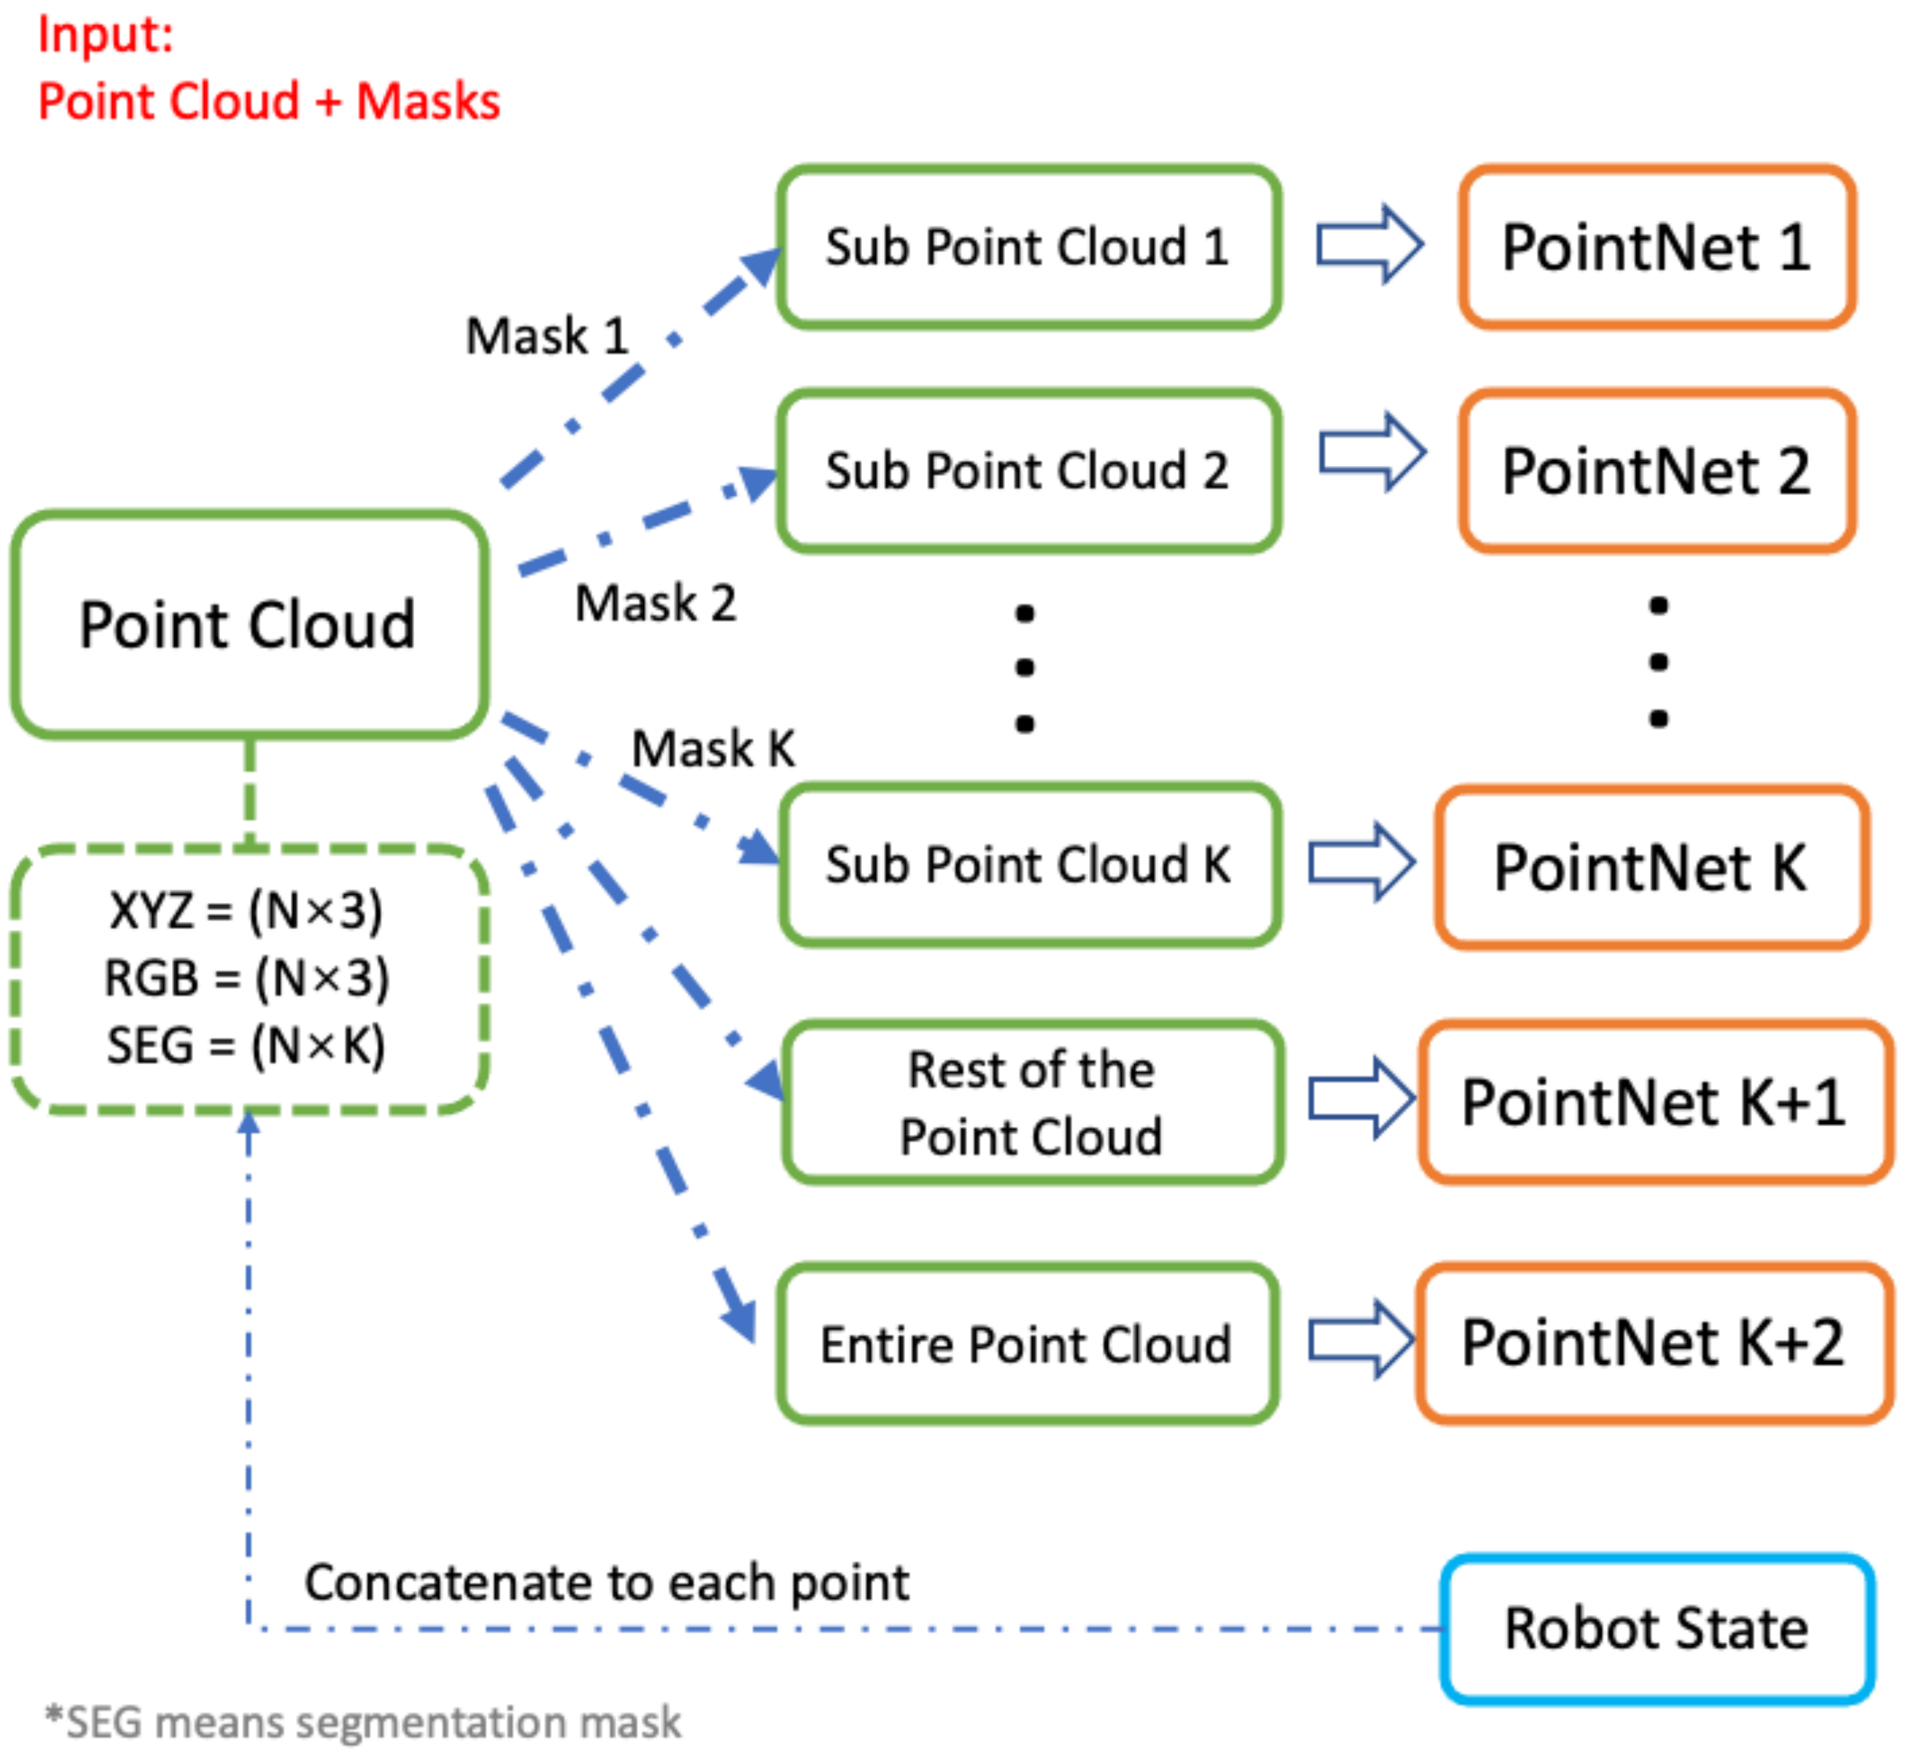

Influence of Network on ManiSkill

- PointNet + Transformer

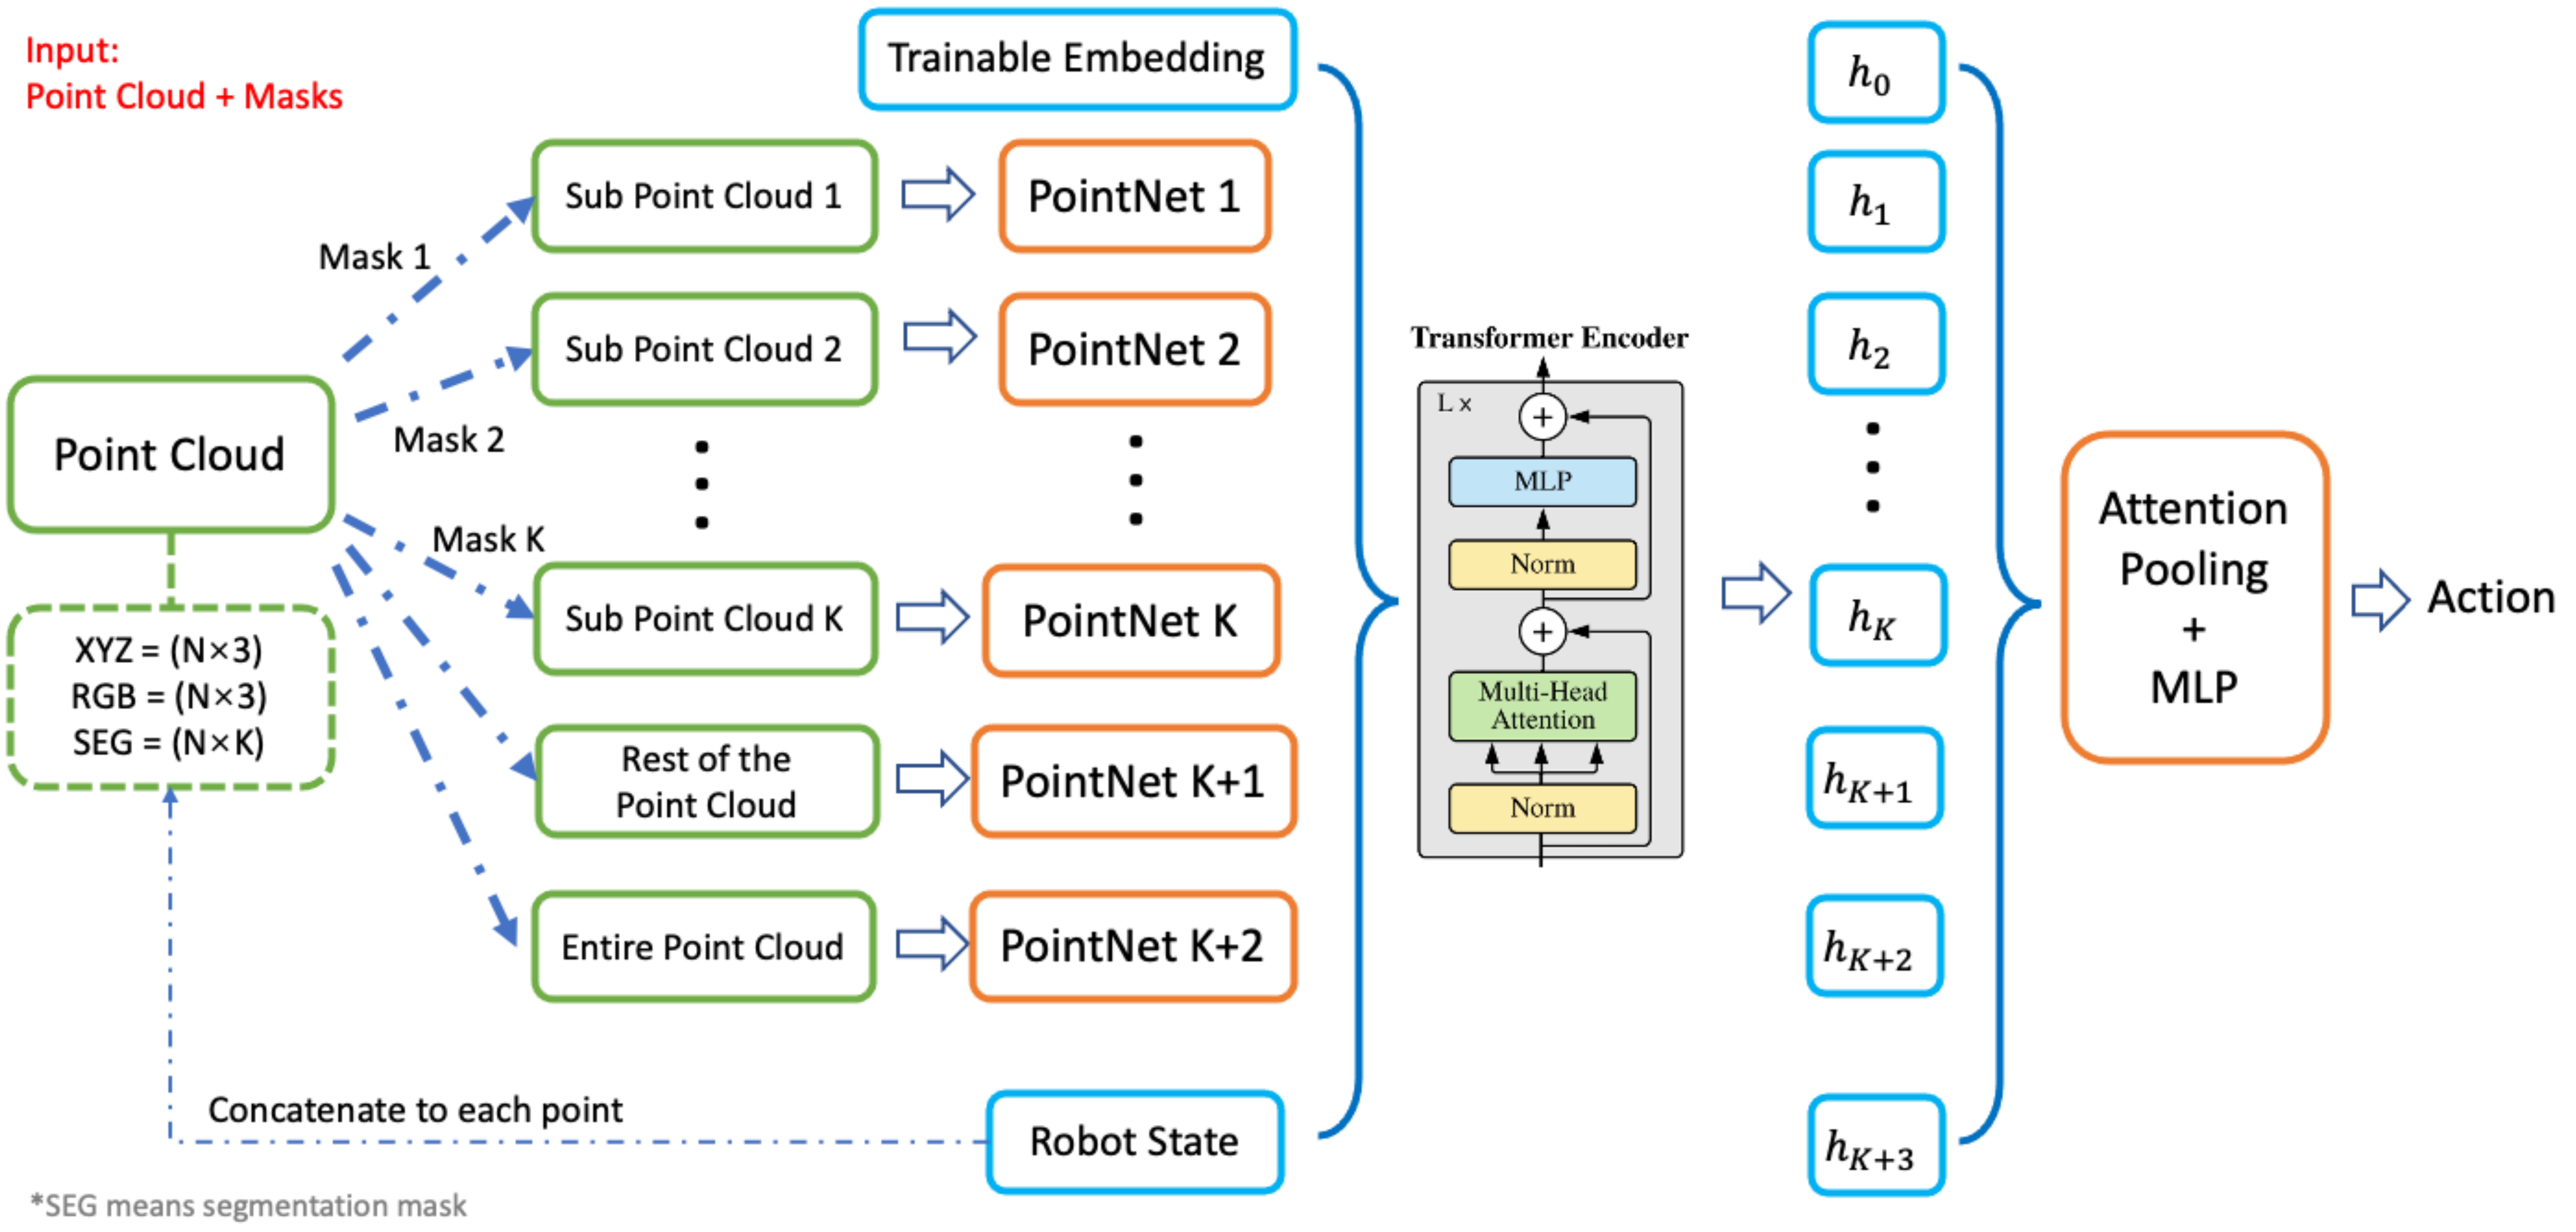

Influence of Network on ManiSkill

- On Behavior Cloning, we achieve better performance on PointNet + Transformer compared to vanilla PointNet.

Algorithm	BC			
Architecture	PointNet		PointNet + Transformer	
Split	Training	Test	Training	Test
OpenCabinetDoor	0.19	0.04	0.27	0.30
OpenCabinetDrawer	0.28	0.09	0.47	0.44
PushChair	0.12	0.07	0.19	0.07
MoveBucket	0.05	0.02	0.18	0.08

Offline Reinforcement Learning

- We can also conduct reinforcement learning on the demonstrations.

- Basic idea:
  - Demonstrations are past experiences
  - The agent should review and exploit them as much as possible
  - e.g., replay the experiences in its mind, and use dynamic programming to update the policy

Offline Reinforcement Learning

- Common pipeline in many offline RL algorithms, such as Batch Constrained Q-Learning (BCQ, Fujimoto et al., 2018) and TD3+BC (Fujimoto et al., 2021)

$$\underset{\pi}{\text{minimize}} \quad L_{\text{reward}} + L_{\text{constraint}}$$

- L_{reward}: learn a policy to maximize the cumulative reward according to the demo set (as in usual RL)
- $L_{\text{constraint}}$: restrict the policy not to deviate too much from the demo policy

Offline Reinforcement Learning

- Offline RL can be combined with online RL to further improve the policy
- A hot topic in ML community today
- We have not found an effective way to do offline RL in ManiSkill
- We leave further explorations for future work

Many Exciting Progress not Covered

- State-based imitation learning
- Offline RL + online RL
- Hierarchical learning

Summary of ManiSkill Challenge

- Learning to manipulate unseen objects with point cloud inputs
- Three tracks designed for researchers working on
  - CV only
  - RL
  - robotics
- 20,000 USD awards!

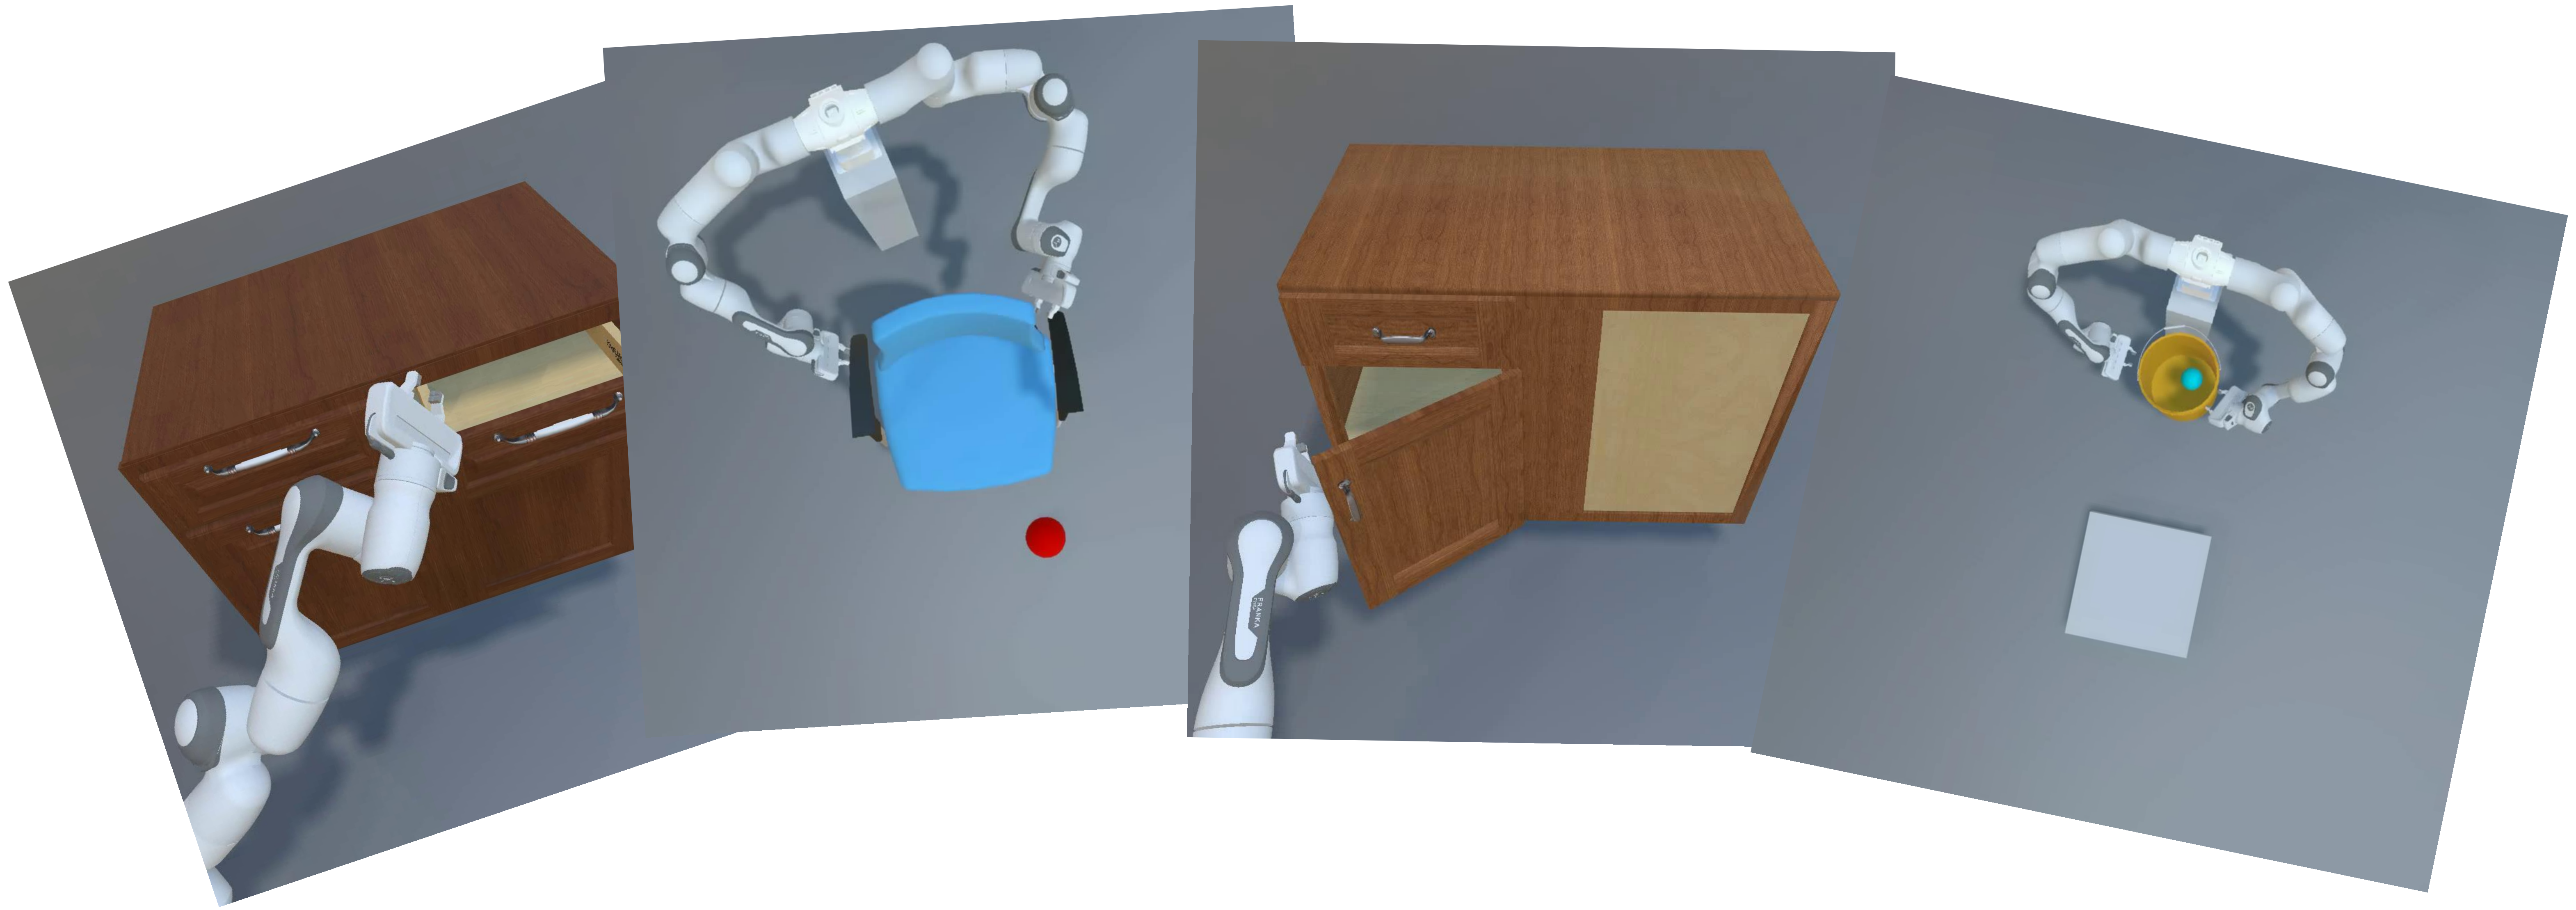

Timelines

- SAPIEN++
  - https://sapien.ucsd.edu
  - Release today (real-time ray-tracing patch in late August)
- ManiSkill Challenge
  - Starts now
    - Register at https://sapien.ucsd.edu/challenges/maniskill2021/
  - Ends: Dec 10, 2021
- Follow our twitter for important announcements
  - https://twitter.com/HaoSuLabUCSD

Website

Twitter

Thank you! Q&A