CoRL 2022 Review

Mini Conference on Robot Learning

2023.01.06

Conference Venues

CONFERENCE ON ROBOT LEARNING
AUCKLAND, DEC 14-18, 2022

Colour meaning:

E	Owen G. Glen Building Entrance 12 Grafton Rd
E1	Engineering Building 1st Entrance 5 Grafton Rd
E2	Engineering Building 2nd Entrance 12 Symonds St
OGG 260	Building 260
405	Building 405
401	Building 401

Addresses: 5 Grafton Rd, 12 Grafton Rd, 12 Symonds St

Contents

- Robot Demonstrations
- Intersting Papers
  - Navigation
  - Open-World Detection
  - Autonomous Vehicle
  - Manipulation
  - Inverse Reinforcement Learning
  - ETC
- Take-home Messages

Robot Demonstrations

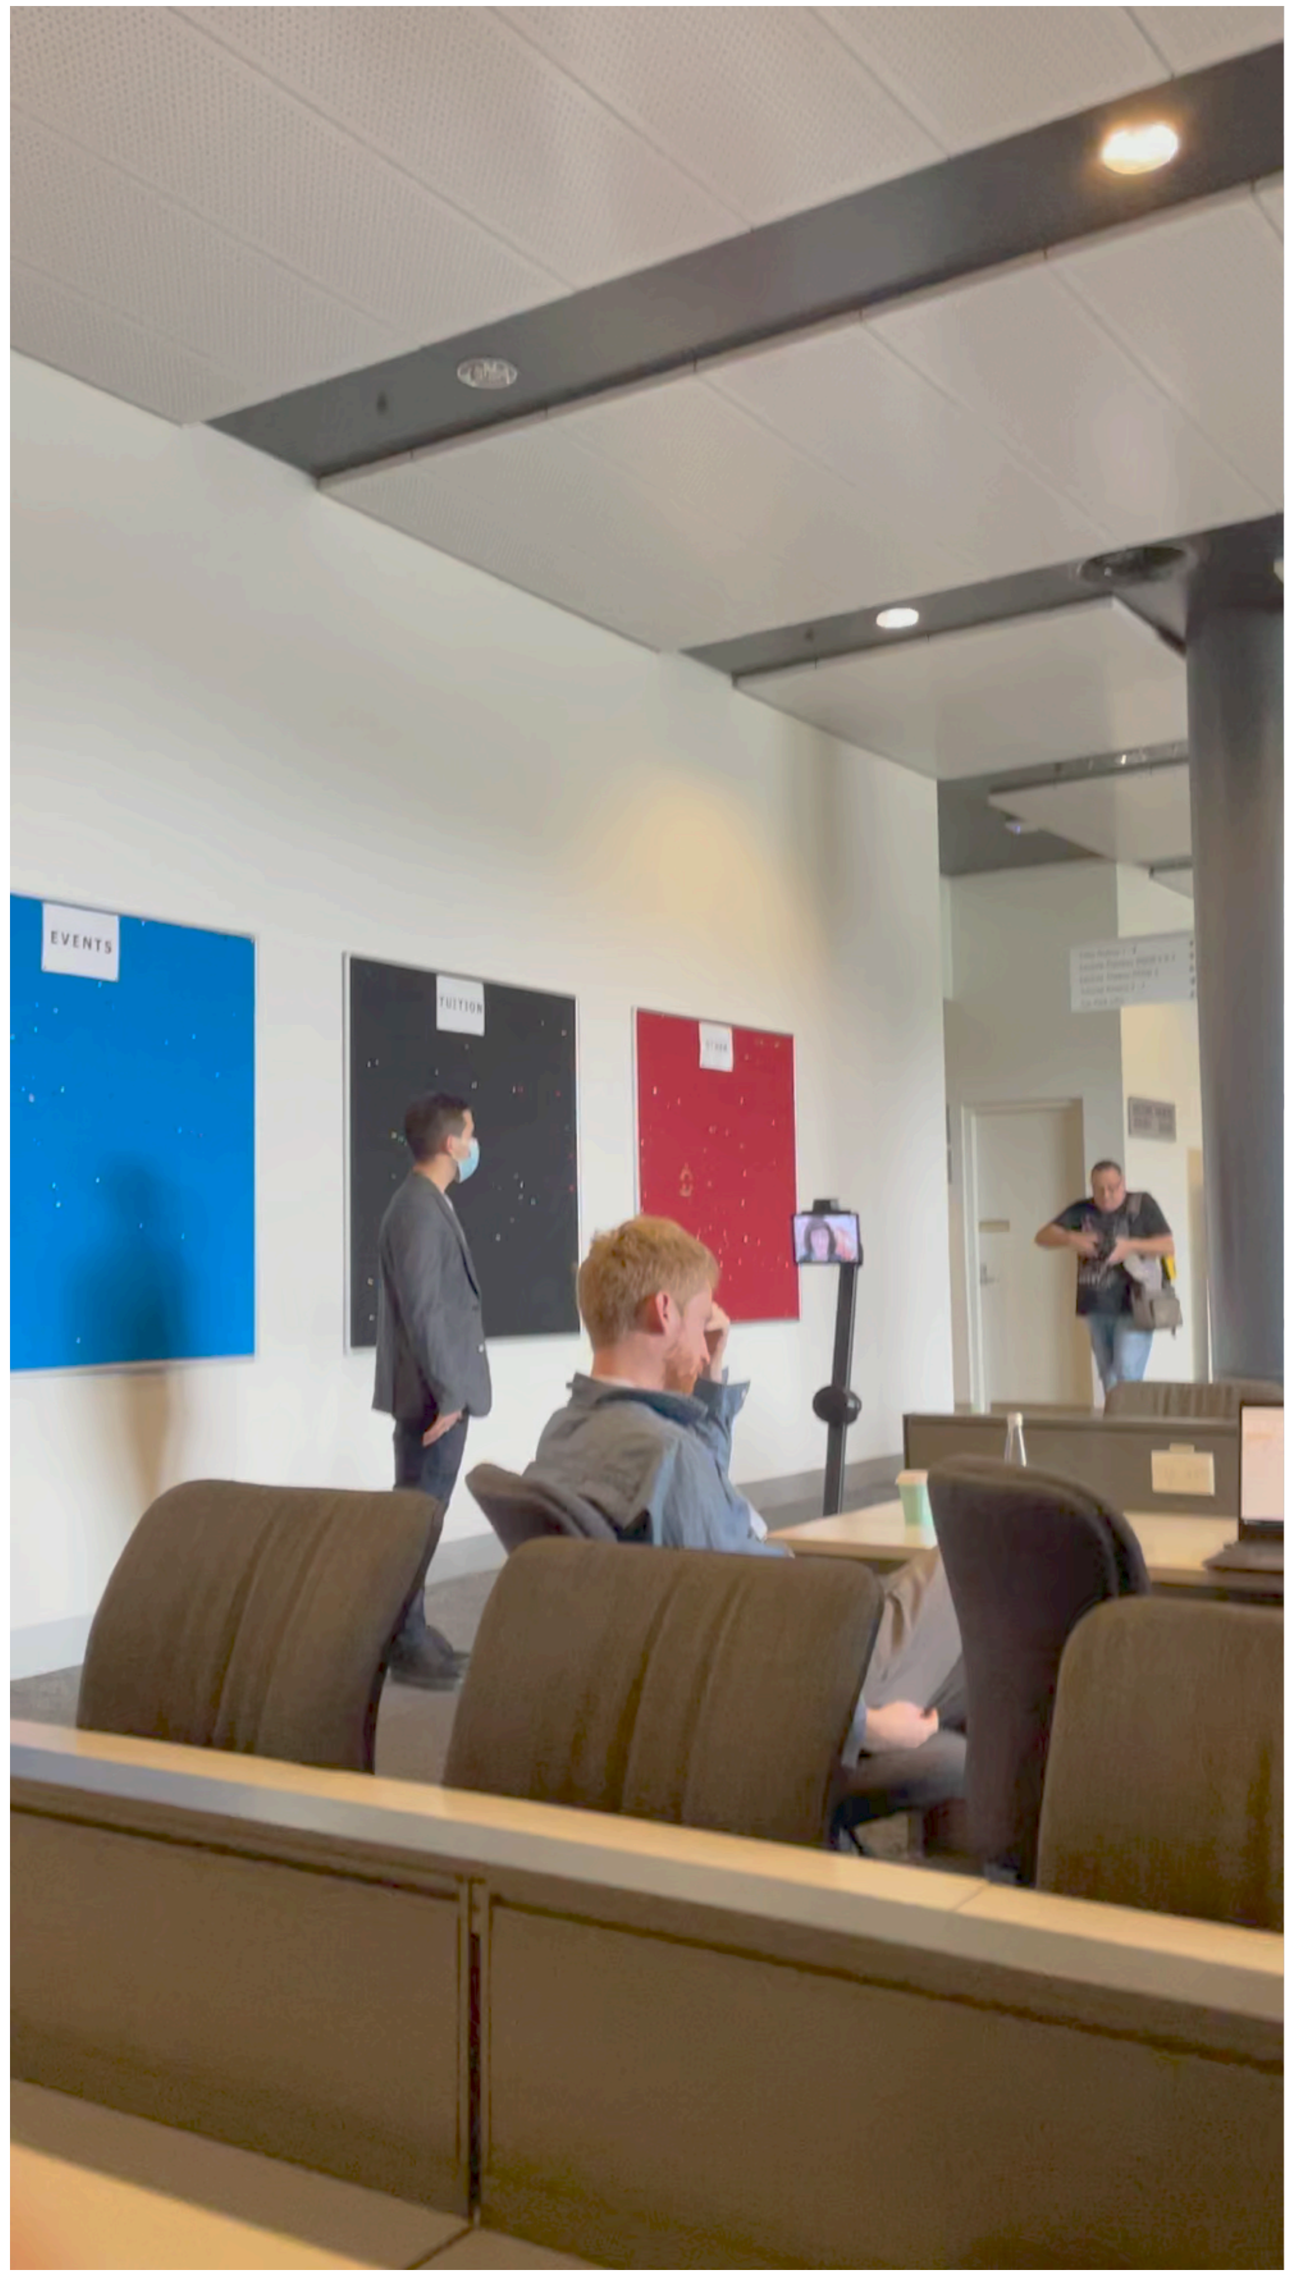

Robot Demonstrations

RLLAB
Robot Learning Laboratory

Navigation

- BEHAVIOR-1K: A Benchmark for Embodied AI with 1,000 Everyday Activities and Realistic Simulation ***Oral Presentation***
- LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision and Action
- Offline Reinforcement Learning for Customizable Visual Navigation ***Oral Presentation***
- Do As I Can, Not As I Say: Grounding Language in Robotic Affordances ***Best Systems Paper Award Finalist***

BEHAVIOR-1K: A Benchmark for Embodied AI with 1,000 Everyday Activities and Realistic Simulation

RLLAB
Robot Learning Laboratory

LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision and Action

RLLAB
Robot Learning Laboratory

LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision, and Action

Dhruv Shah⚓🐻, Błażej Osiński⚓🐻♦, Brian Ichter♣, Sergey Levine🐻

🐻UC Berkeley, ♦University of Warsaw, ♣Robotics at Google

https://sites.google.com/view/lmnav

BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

⚓ *Equal Contribution*

Offline Reinforcement Learning for Customizable Visual Navigation

RLLAB
Robot Learning Laboratory

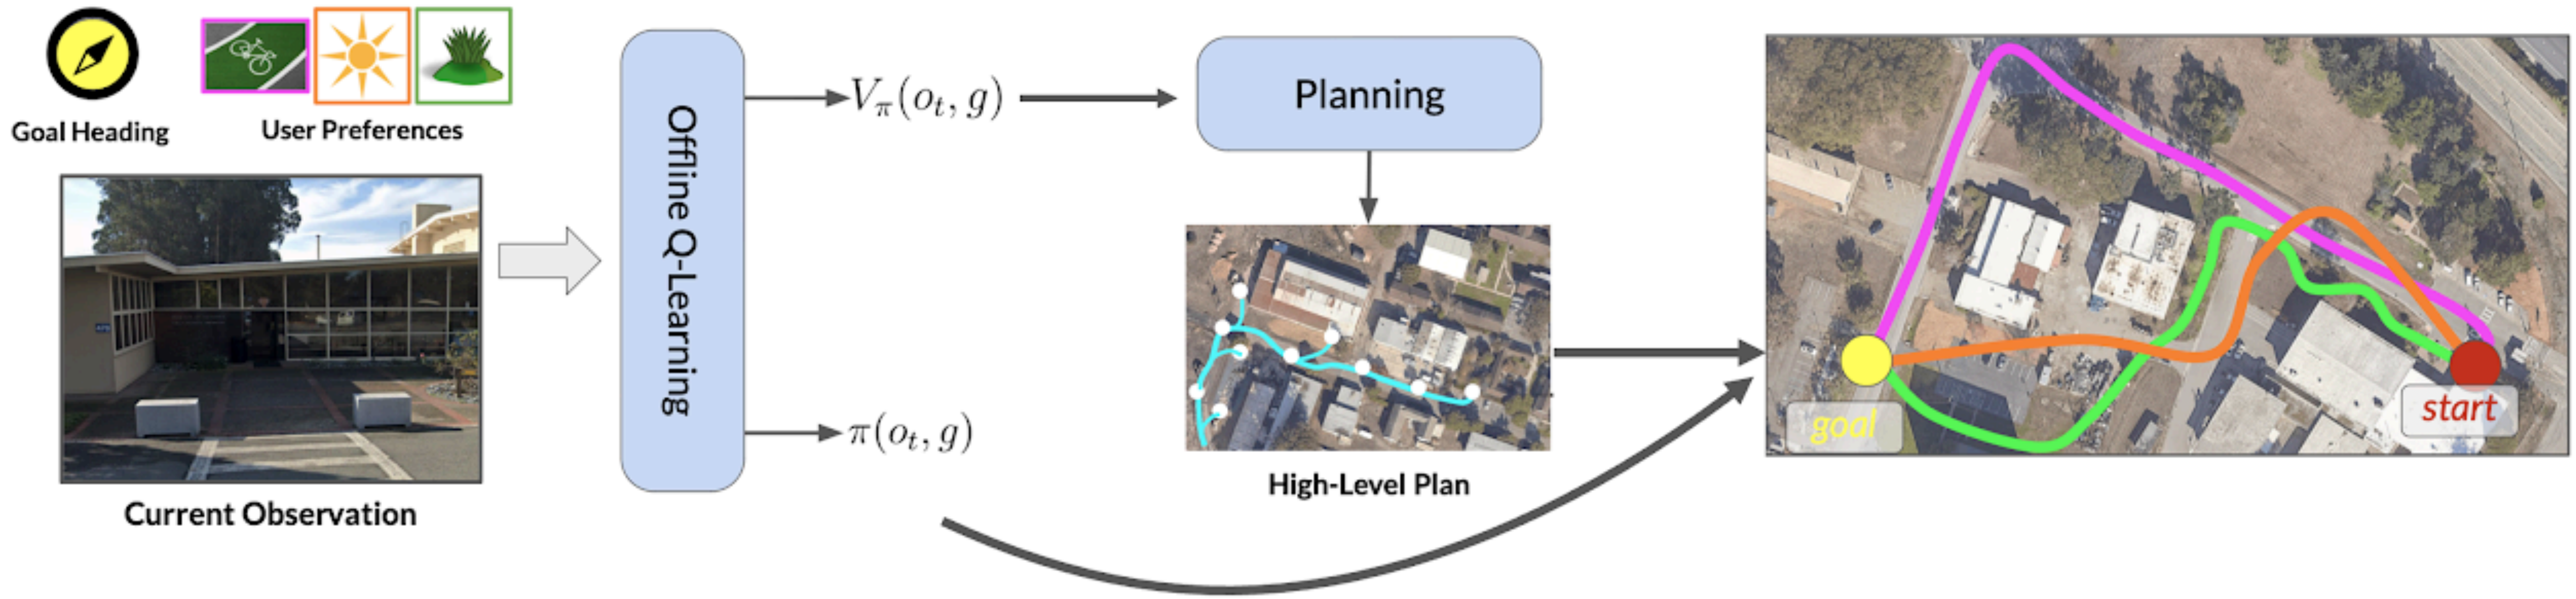

- Goal is not merely to reach a particular destination without collision, but to do so while maximizing some desired utility measure.
  - obeying the rules of the road, staying in a bike lane or off the lawn, maintaining safety
- It used a topological graph as a nonparametric memory
- The node in the graph represents the state and the edges represents the cumulative reward the robot will accumulate as it travels from one node to another.
- Then, finds the optimal path using Dijkstra algorithm, which maximizes the value
- Extracted meta labels, such as “sunny” and “grassy”, from RGB observations and use it as additive rewards.

Offline Reinforcement Learning for Visual Navigation

Dhruv Shah[⚓], Arjun Bhorkar[⚓], Hrish Leen,
Ilya Kostrikov, Nicholas Rhinehart, Sergey Levine

Conference on Robot Learning (CoRL) 2022
Auckland, New Zealand

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

[⚓] *Equal Contribution*

Do As I Can, Not As I Say: Grounding Language in Robotic Affordances

RLLAB
Robot Learning Laboratory

I spilled my drink, can you help?

GPT3	You could try using a vacuum cleaner.
LaMDA	Do you want me to find a cleaner?
FLAN	I'm sorry, I didn't mean to spill it.

I spilled my drink, can you help?

LLM

"find a cleaner"
"find a sponge"
"go to the trash can"
"pick up the sponge"
"try using the vacuum"

Value Functions

"find a cleaner"
"find a sponge"
"go to the trash can"
"pick up the sponge"
"try using the vacuum"

SayCan

"find a cleaner"
"find a sponge"
"go to the trash can"
"pick up the sponge"
"try using the vacuum"

I would:
1. find a sponge
2. pick up the sponge
3. come to you
4. put down the sponge
5. done

Supplementary Video for
"Do As I Can, Not As I Say:
Grounding Language in Robotics Affordances"

Robotics at Google and Everyday Robots

Robot Learning Laboratory

Open-World Detection

- Semantic Abstraction: Open-World 3D Scene Understanding from 2D Vision-Language Models

Semantic Abstraction: Open-World 3D Scene Understanding from 2D Vision-Language Models

Semantic Abstraction: Open-World 3D Scene Understanding from 2D Vision-Language Models

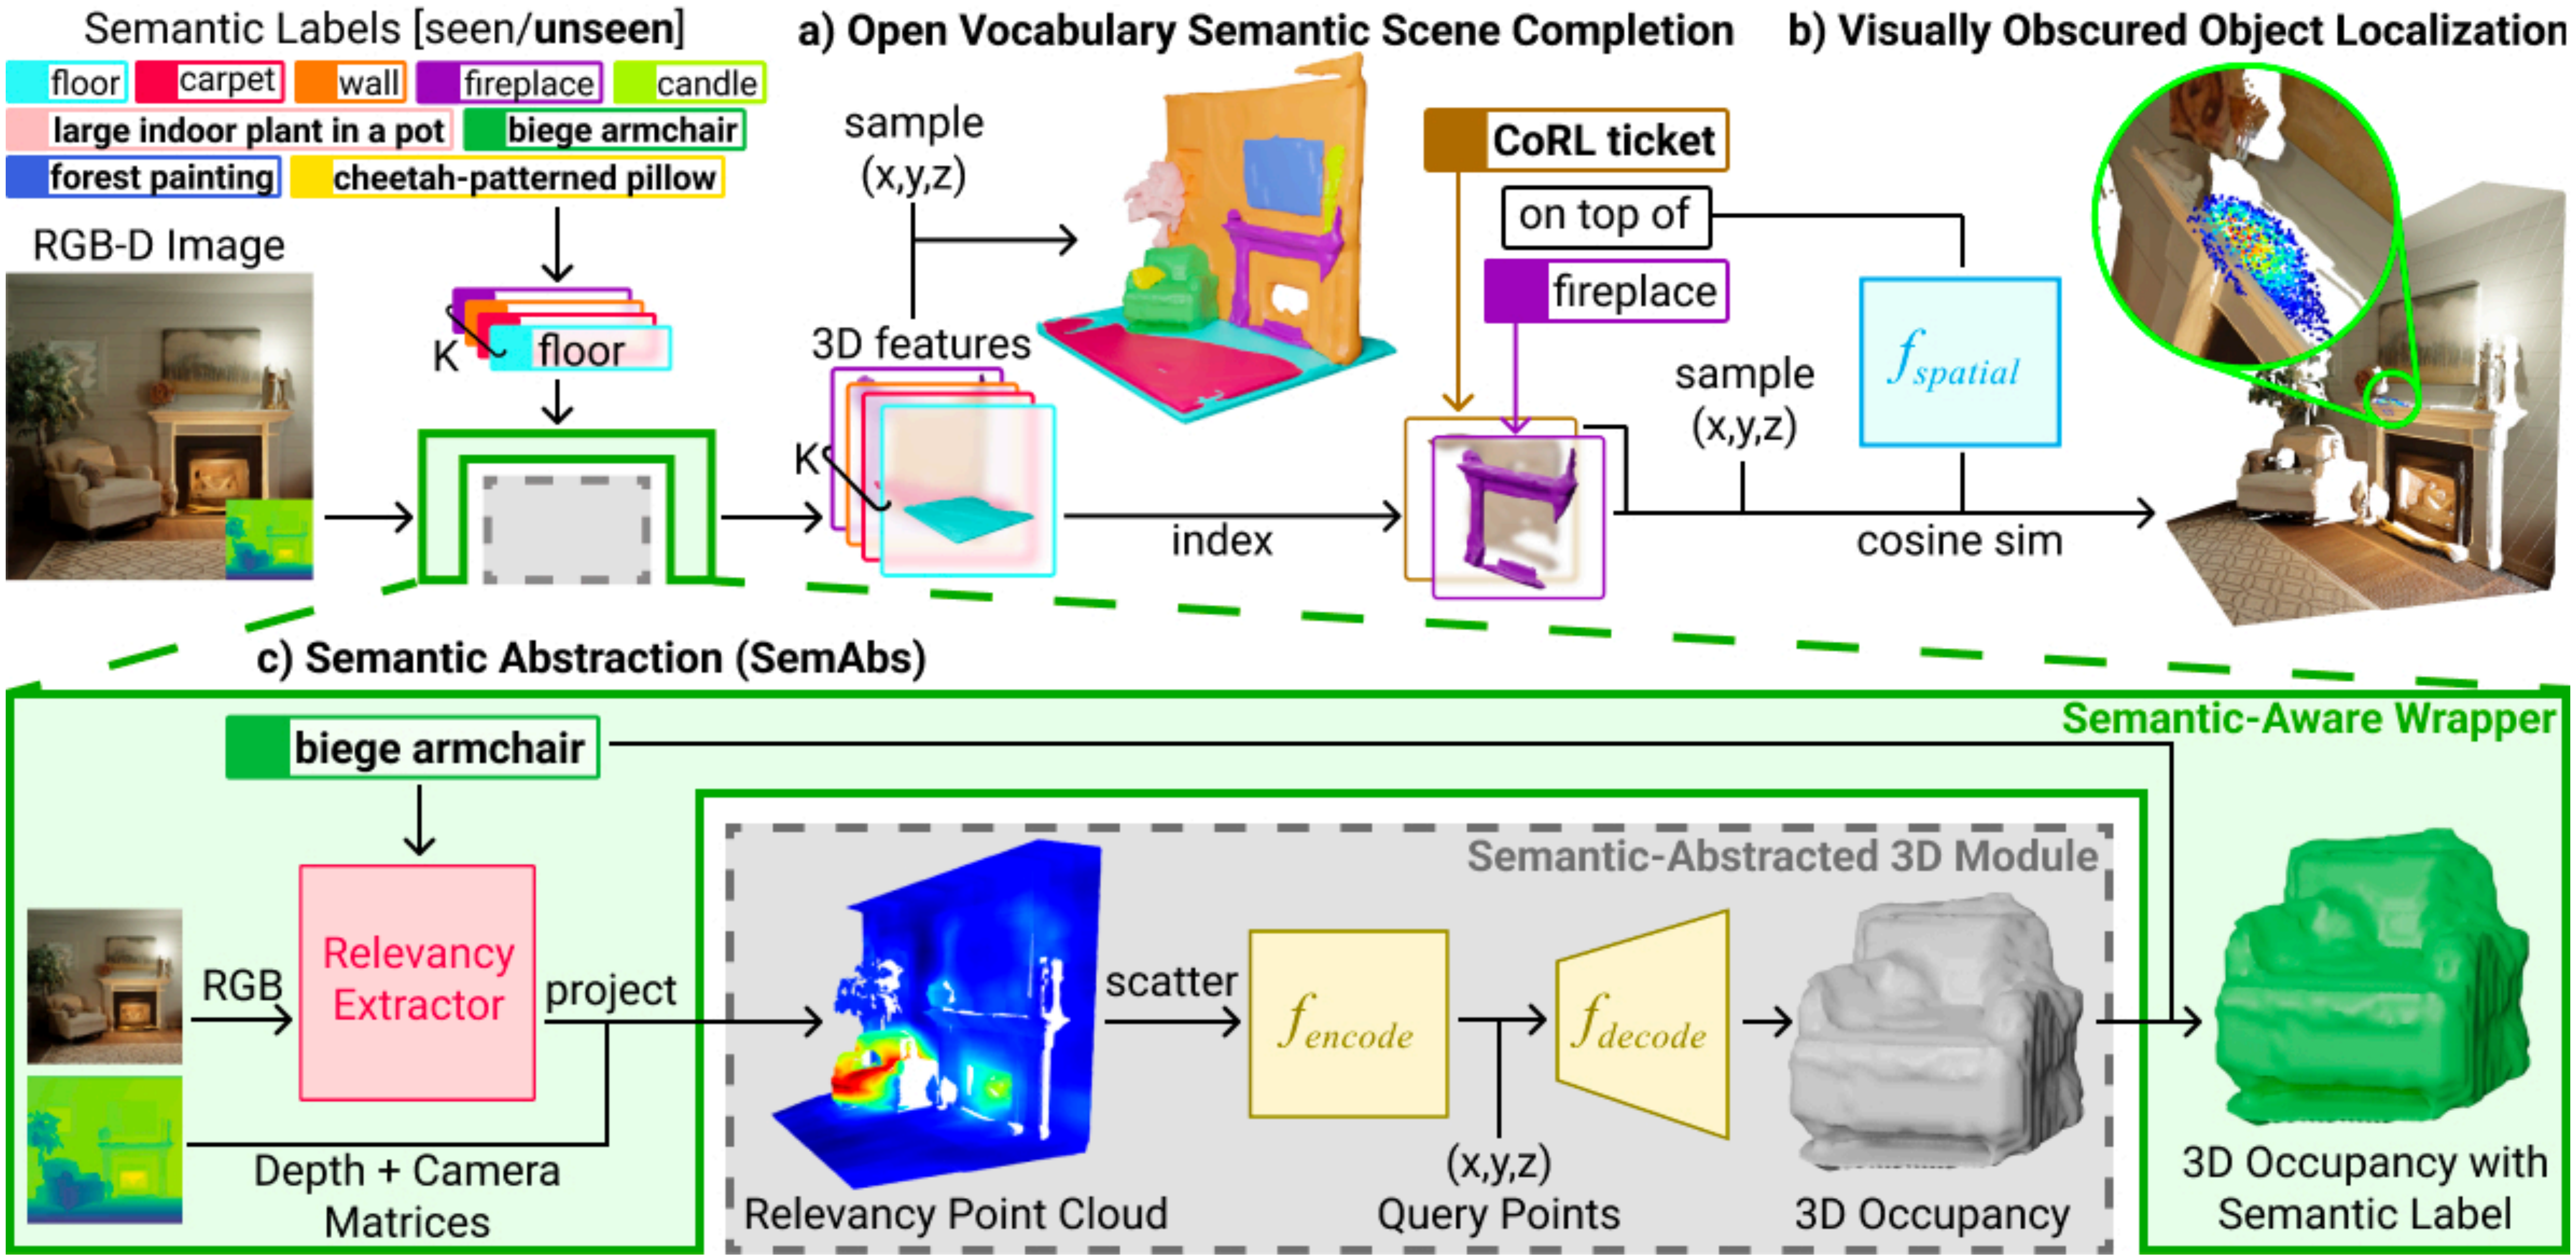

Open-Vocabulary Scene Completion

Text Input

Scene Completion

Semantic labels

RGB-D Input

Open-Vocabulary Scene Completion

Text Input

RGB-D Input

Scene Completion

What about objects that are really hard to see (i.e, *visually-obscured*) ?

e.g., How can we tell the robot to take the

lemon on the shelf

Autonomous Vehicle

- LEADER: Learning Attention over Driving Behaviors ***Best Paper Award Finalist***

LEADER: Learning Attention over Driving Behaviors

LEADER:
LEARNING ATTENTION
OVER DRIVING BEHAVIORS
FOR PLANNING UNDER UNCERTAINTY

Mohamad H. Danesh, Panpan Cai, David Hsu

Oregon State University
上海交通大学
SHANGHAI JIAO TONG UNIVERSITY
NUS
National University of Singapore

Manipulation

- Evo-NeRF: Evolving NeRF for Sequential Robot Grasping ***Oral Presentation***
- SE(2)-Equivariant Pushing Dynamics Models for Tabletop Object Manipulations ***Oral Presentation***

Evo-NeRF: Evolving NeRF for Sequential Robot Grasping

RLLAB
Robot Learning Laboratory

(A) Early Stopped Image Capture (B) Incremental NeRF Optimization (C) Grasp

Evo-NeRF: Evolving NeRF for Sequential Robot Grasping

RLLAB
Robot Learning Laboratory

(b) Grasp

(a) Full Capture Trajectory (c) Update Trajectory (d) Consecutive NeRF Updates

SE(2)-Equivariant Pushing Dynamics Models for Tabletop Object Manipulations

- The pushing dynamics model needs to be equivariant to the SE(2) transformation.

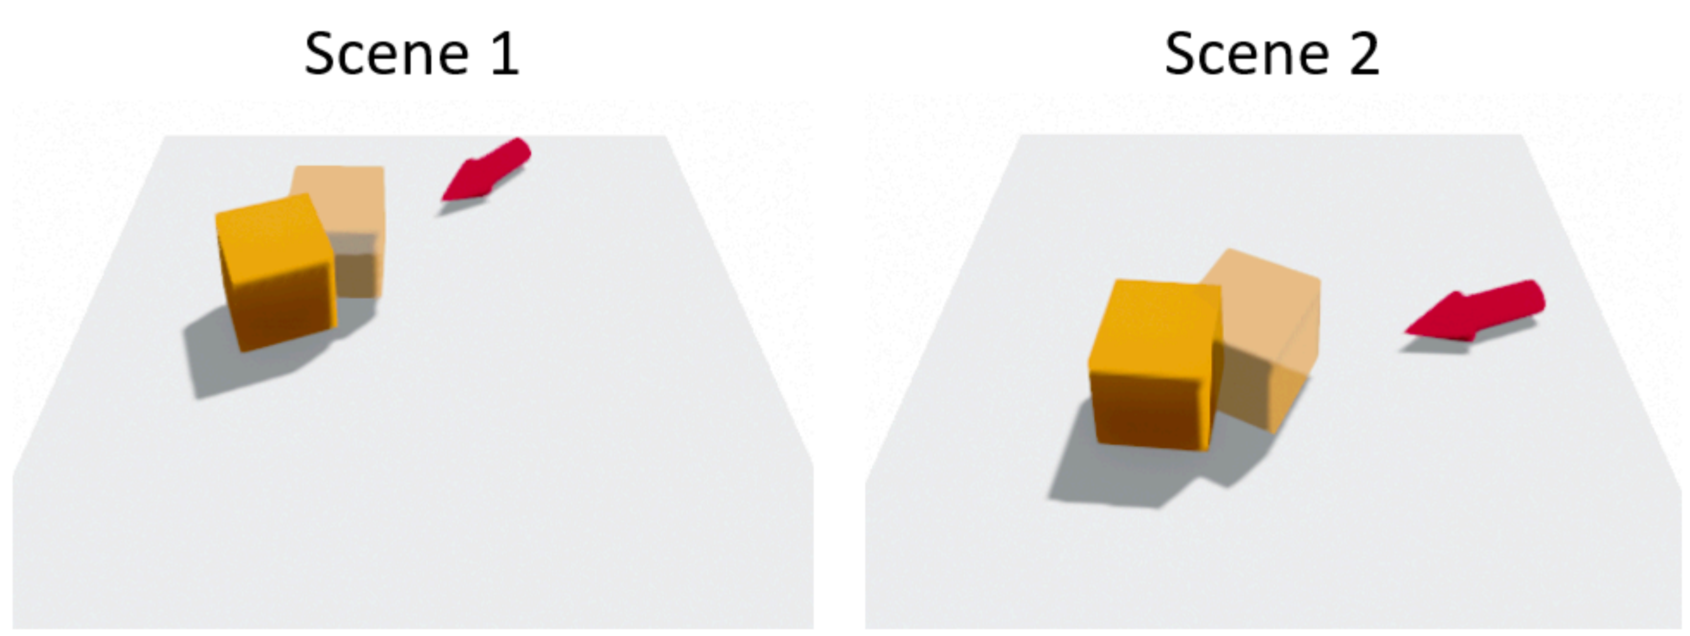

Definition 1 *A pushing dynamics model f is* SE(2)*-equivariant if*

$$\{\mathbf{C}\mathbf{T}'_i\}_{i=1}^N = f(\{(\mathbf{C}\mathbf{T}_i, \mathbf{q}_i)\}_{i=1}^N, (\mathbf{Rot}(\hat{\mathbf{z}}, \theta)\mathbf{p} + \mathbf{t_{xy}}, \mathbf{Rot}(\hat{\mathbf{z}}, \theta)\mathbf{v})) \tag{1}$$

for all object numbers $N \leq M$ and rigid-body transformations $\mathbf{C}$ that have the following form

$$\mathbf{C} = \begin{bmatrix} \mathbf{Rot}(\hat{\mathbf{z}}, \theta) & \mathbf{t_{xy}} \\ 0 & 1 \end{bmatrix}, \tag{2}$$

where $\mathbf{Rot}(\hat{\mathbf{z}}, \theta)$ is a 3×3 rotation matrix for rotations around z-axis and $\mathbf{t_{xy}} = (t_x, t_y, 0) \in \mathbb{R}^3$.

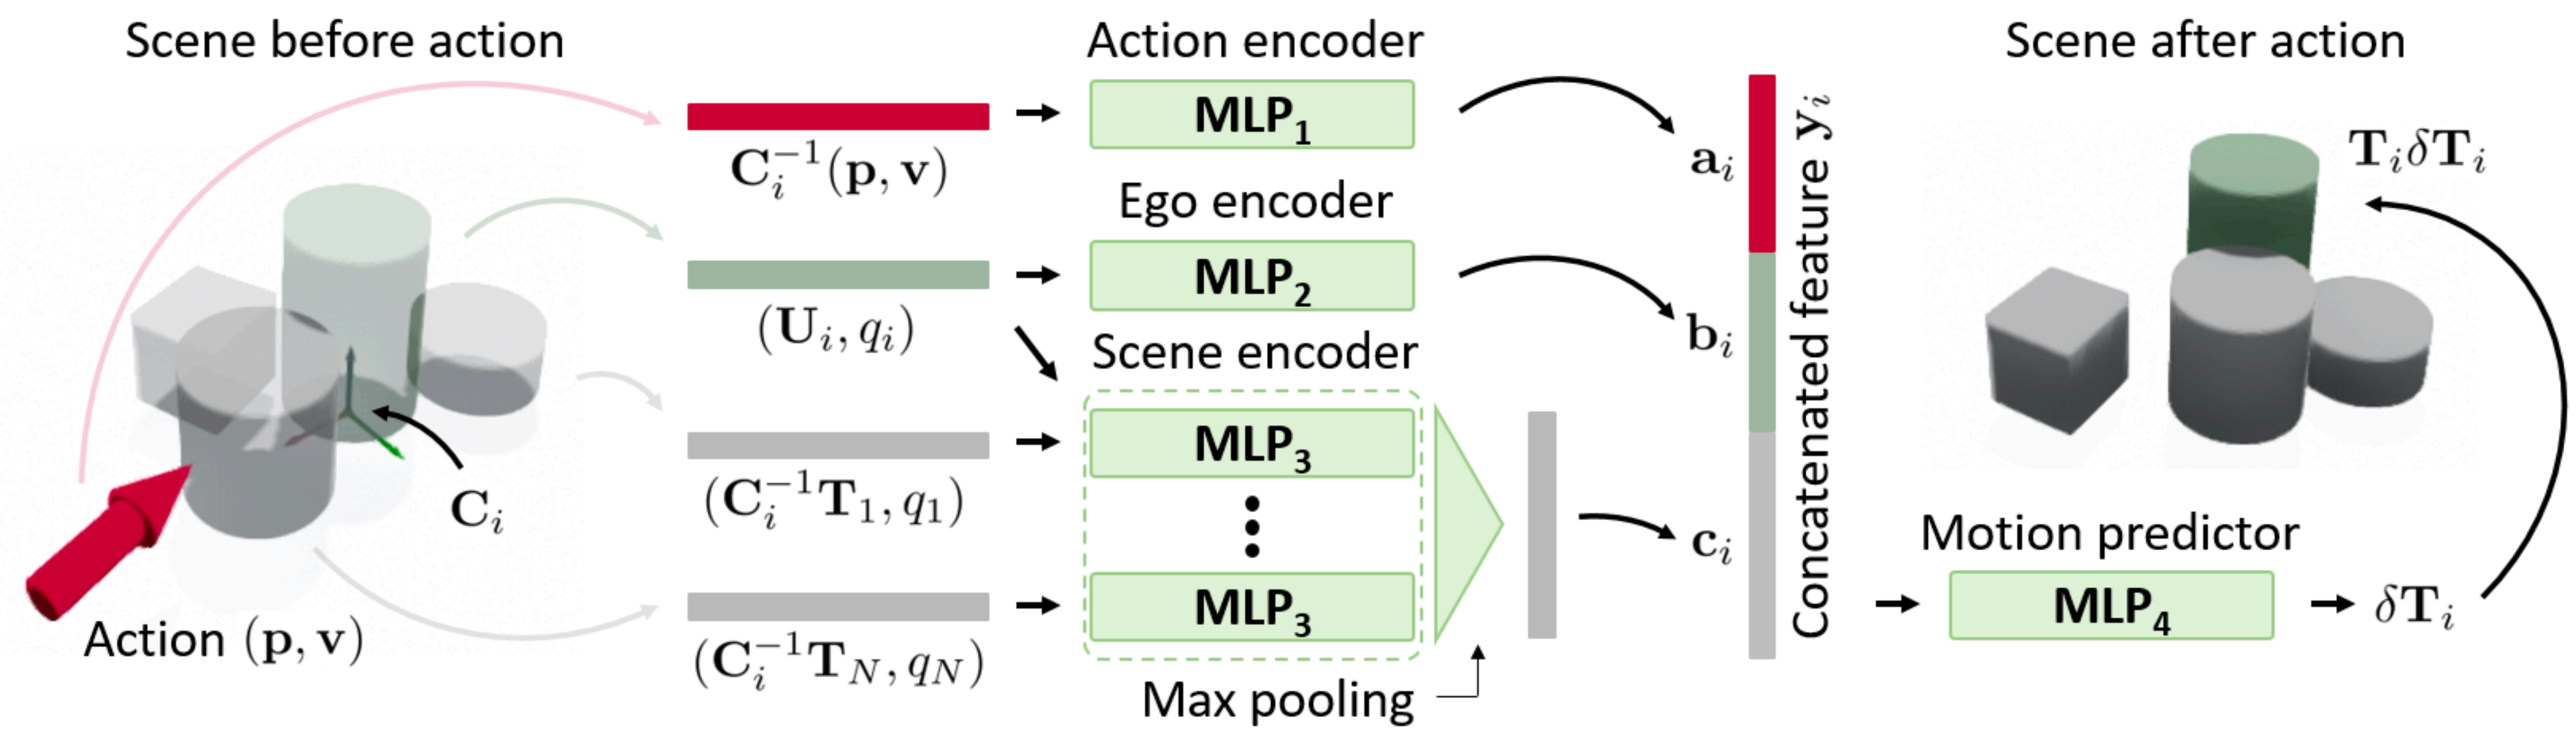

SE(2)-Equivariant Pushing Dynamics Models for Tabletop Object Manipulations

RLLAB
Robot Learning Laboratory

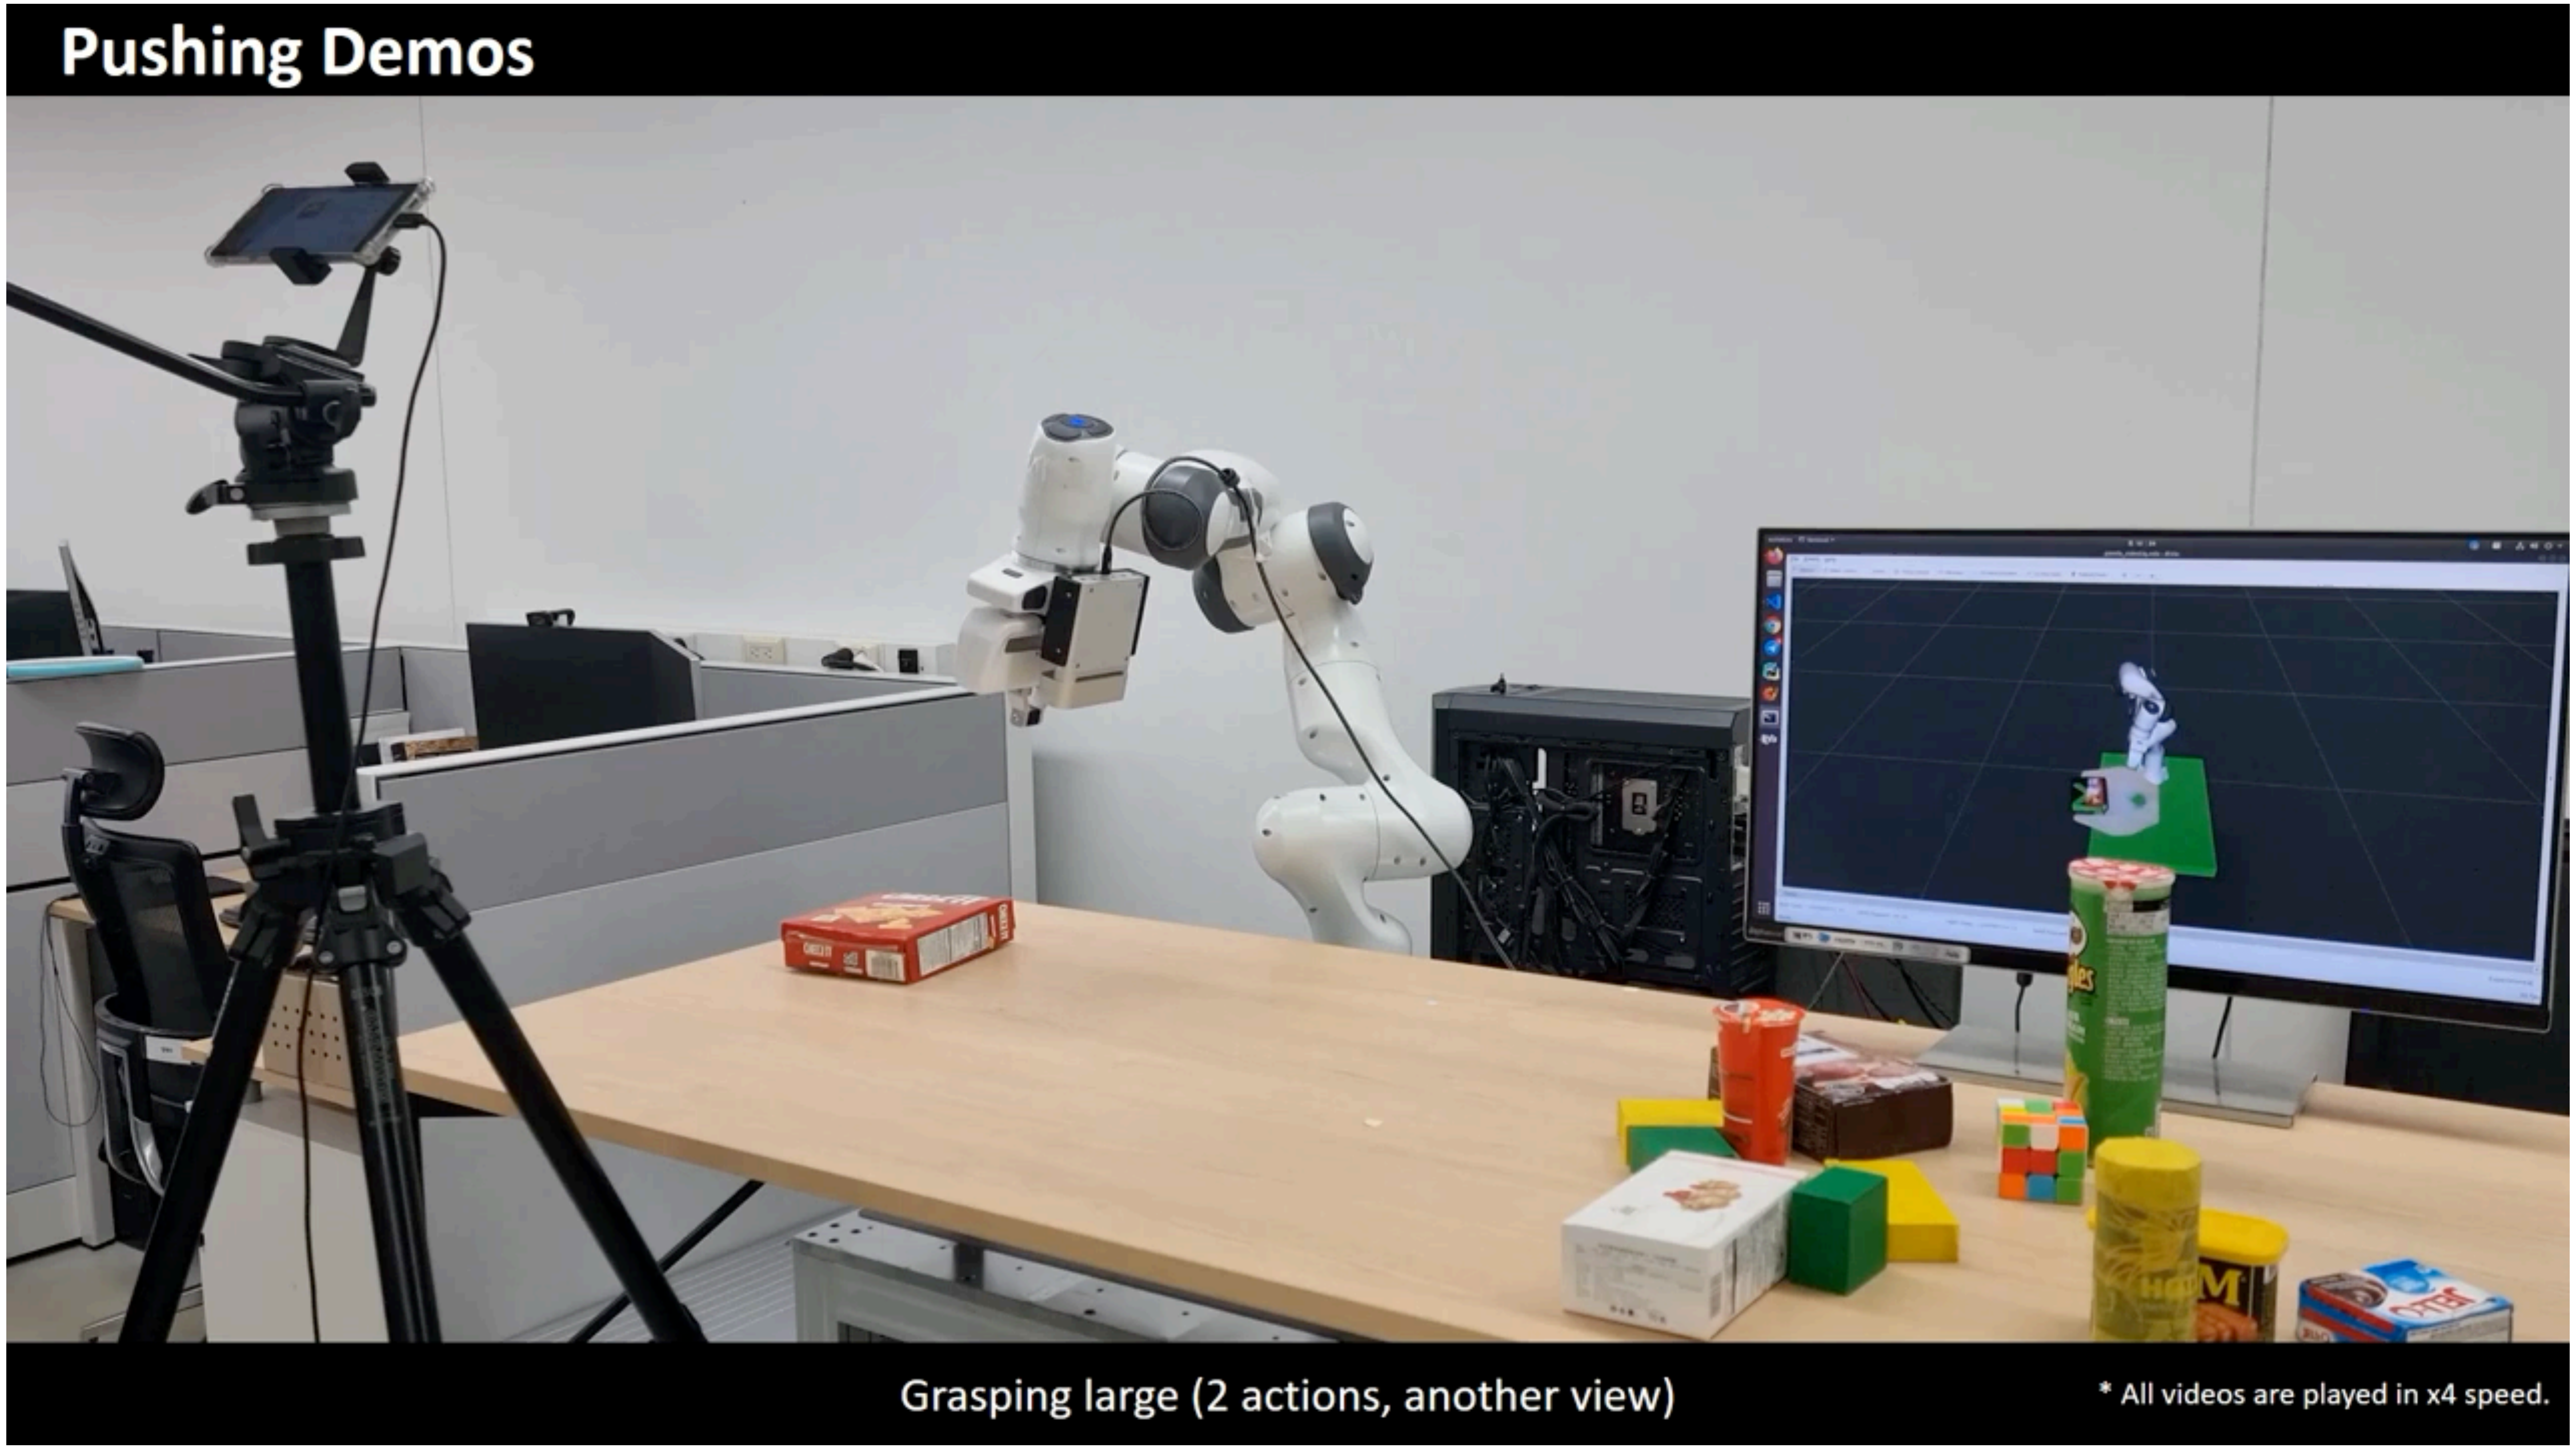

RLLAB
Robot Learning Laboratory

Inverse Reinforcement Learning

- Training Robots to Evaluate Robots: Example-Based Interactive Reward Functions for Policy Learning
- Learning Agile Skills via Adversarial Imitation of Rough Partial Demonstrations ***Best Paper Award Finalist***

Training Robots to Evaluate Robots: Example-Based Interactive Reward Functions for Policy Learning

RLLAB
Robot Learning Laboratory

- Learning an interactive reward function from successful / failure cases

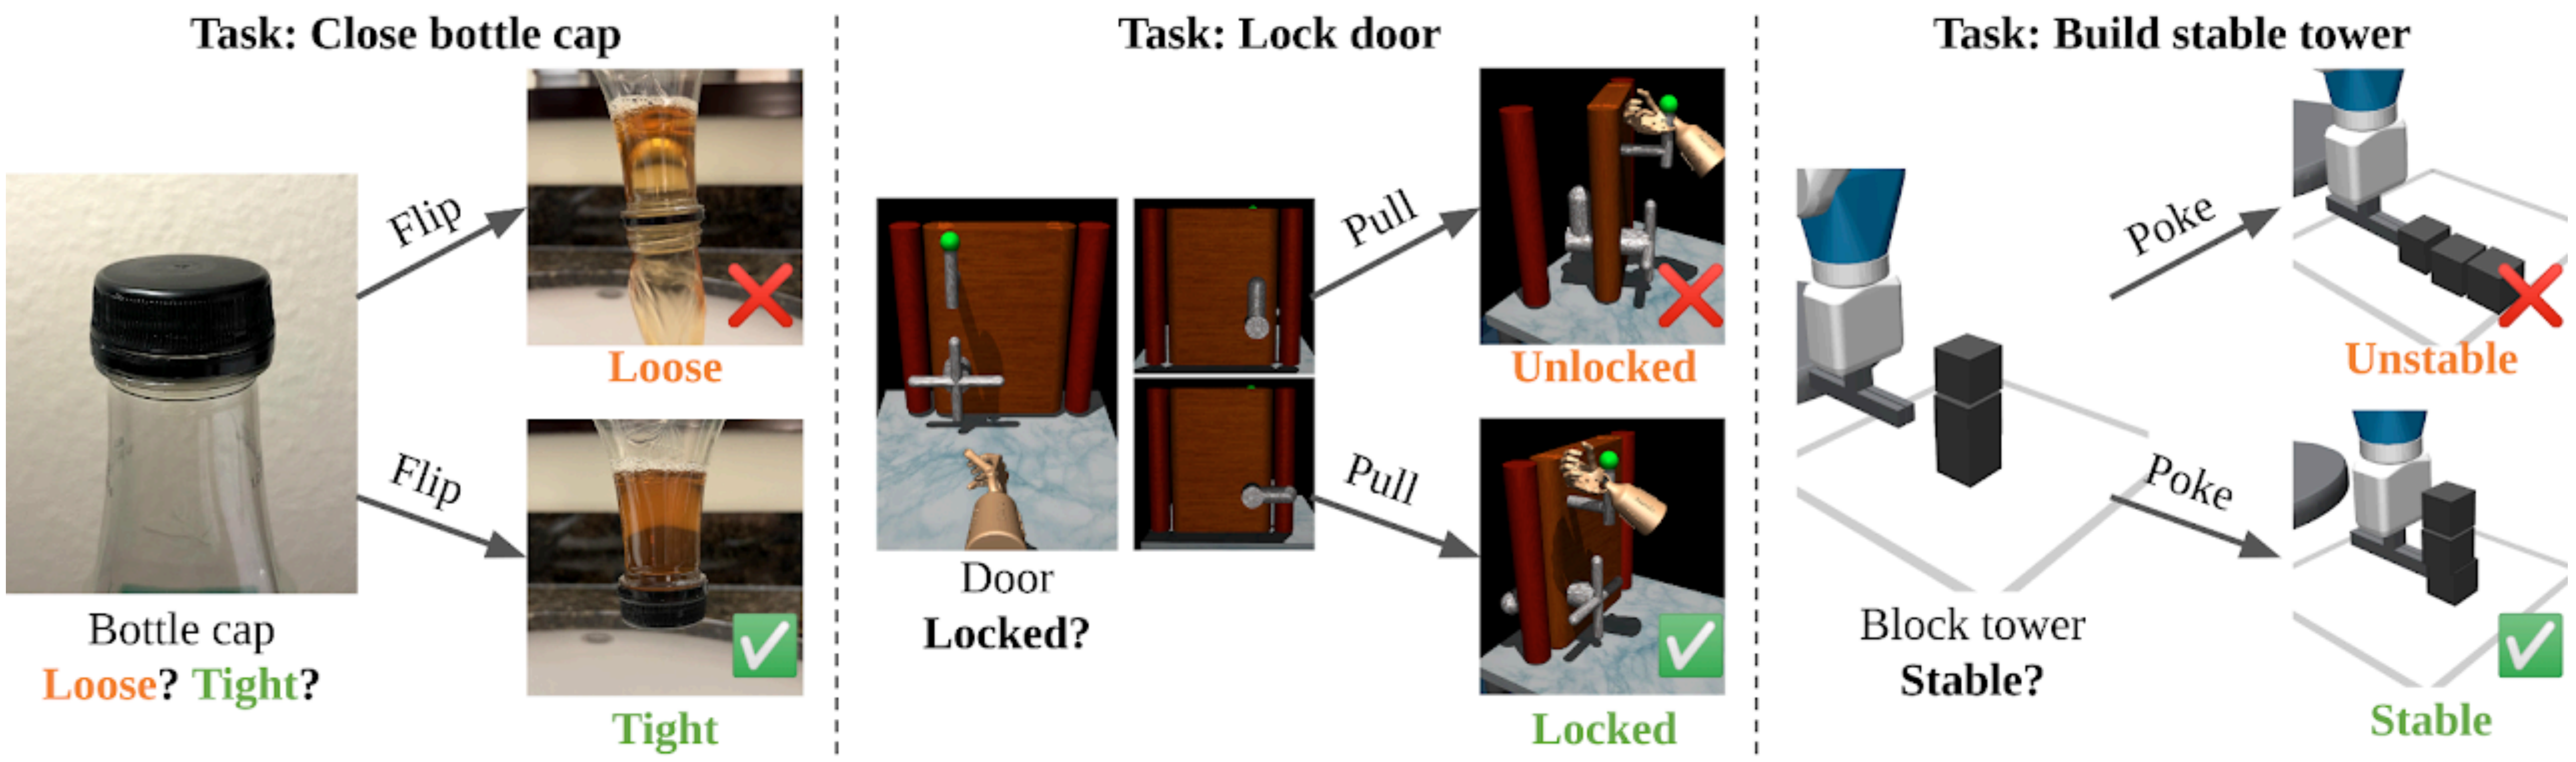

Training Robots to Evaluate Robots: Example-Based Interactive Reward Functions for Policy Learning

RLLAB
Robot Learning Laboratory

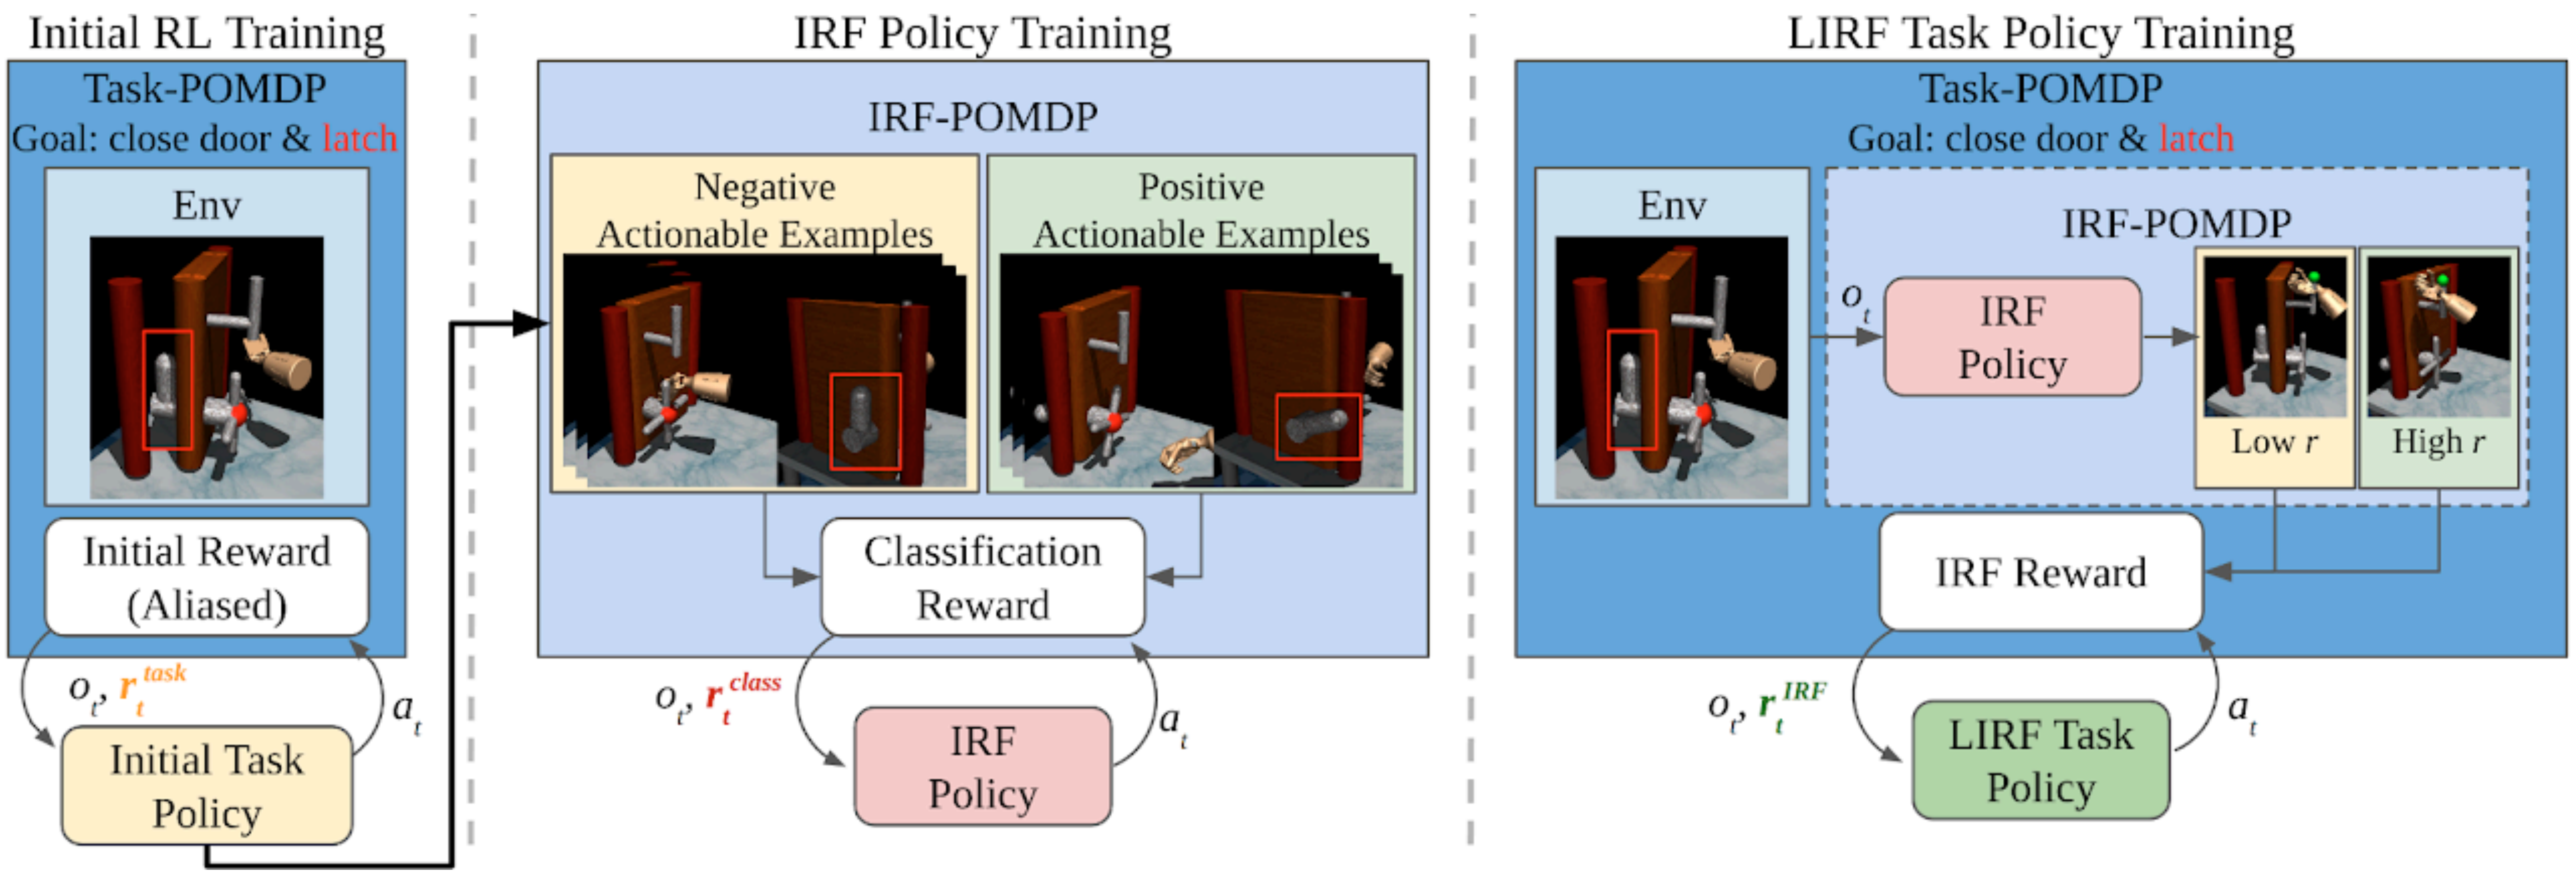


MAX PLANCK INSTITUTE
FOR INTELLIGENT SYSTEMS
ETH zürich

RLLAB
Robot Learning Laboratory

ETC

- DayDreamer: World Models for Physical Robot Learning
- Watch and Match: Supercharging Imitation with Regularized Optimal Transport ***Best Paper Award Finalist***
- Proactive Robot Assistance via SpatioTemporal Object Modeling
- Real-time Mapping of Physical Scene Properties with an Autonomous Robot Experimenter ***Oral Presentation***
- Temporal Logic Imitation: Learning Plan-Satisficing Motion Policies from Demonstrations ***Oral Presentation***

DayDreamer: World Models for Physical Robot Learning

A1 Quadruped Walking

UR5 Multi-Object Visual Pick Place

XArm Visual Pick and Place

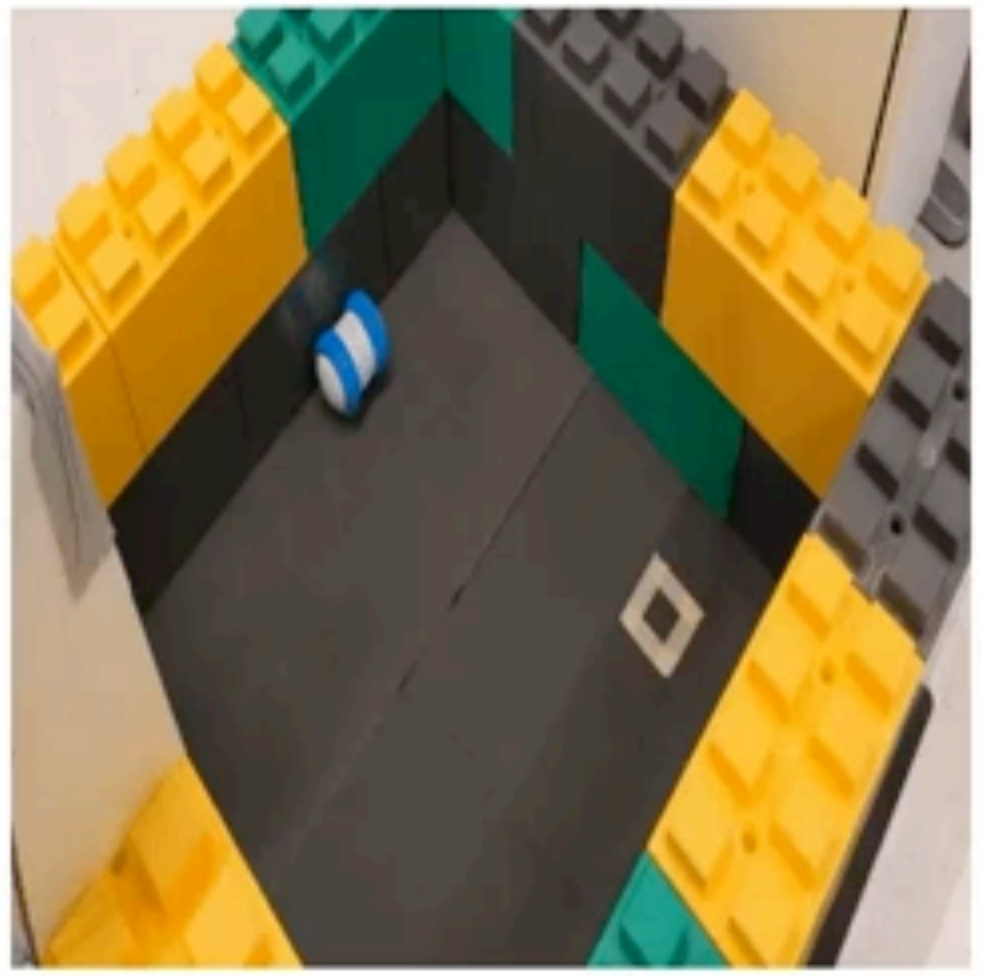

Sphero Ollie Visual Navigation

Learned from Scratch in the Real World

Watch and Match: Supercharging Imitation with Regularized Optimal Transport

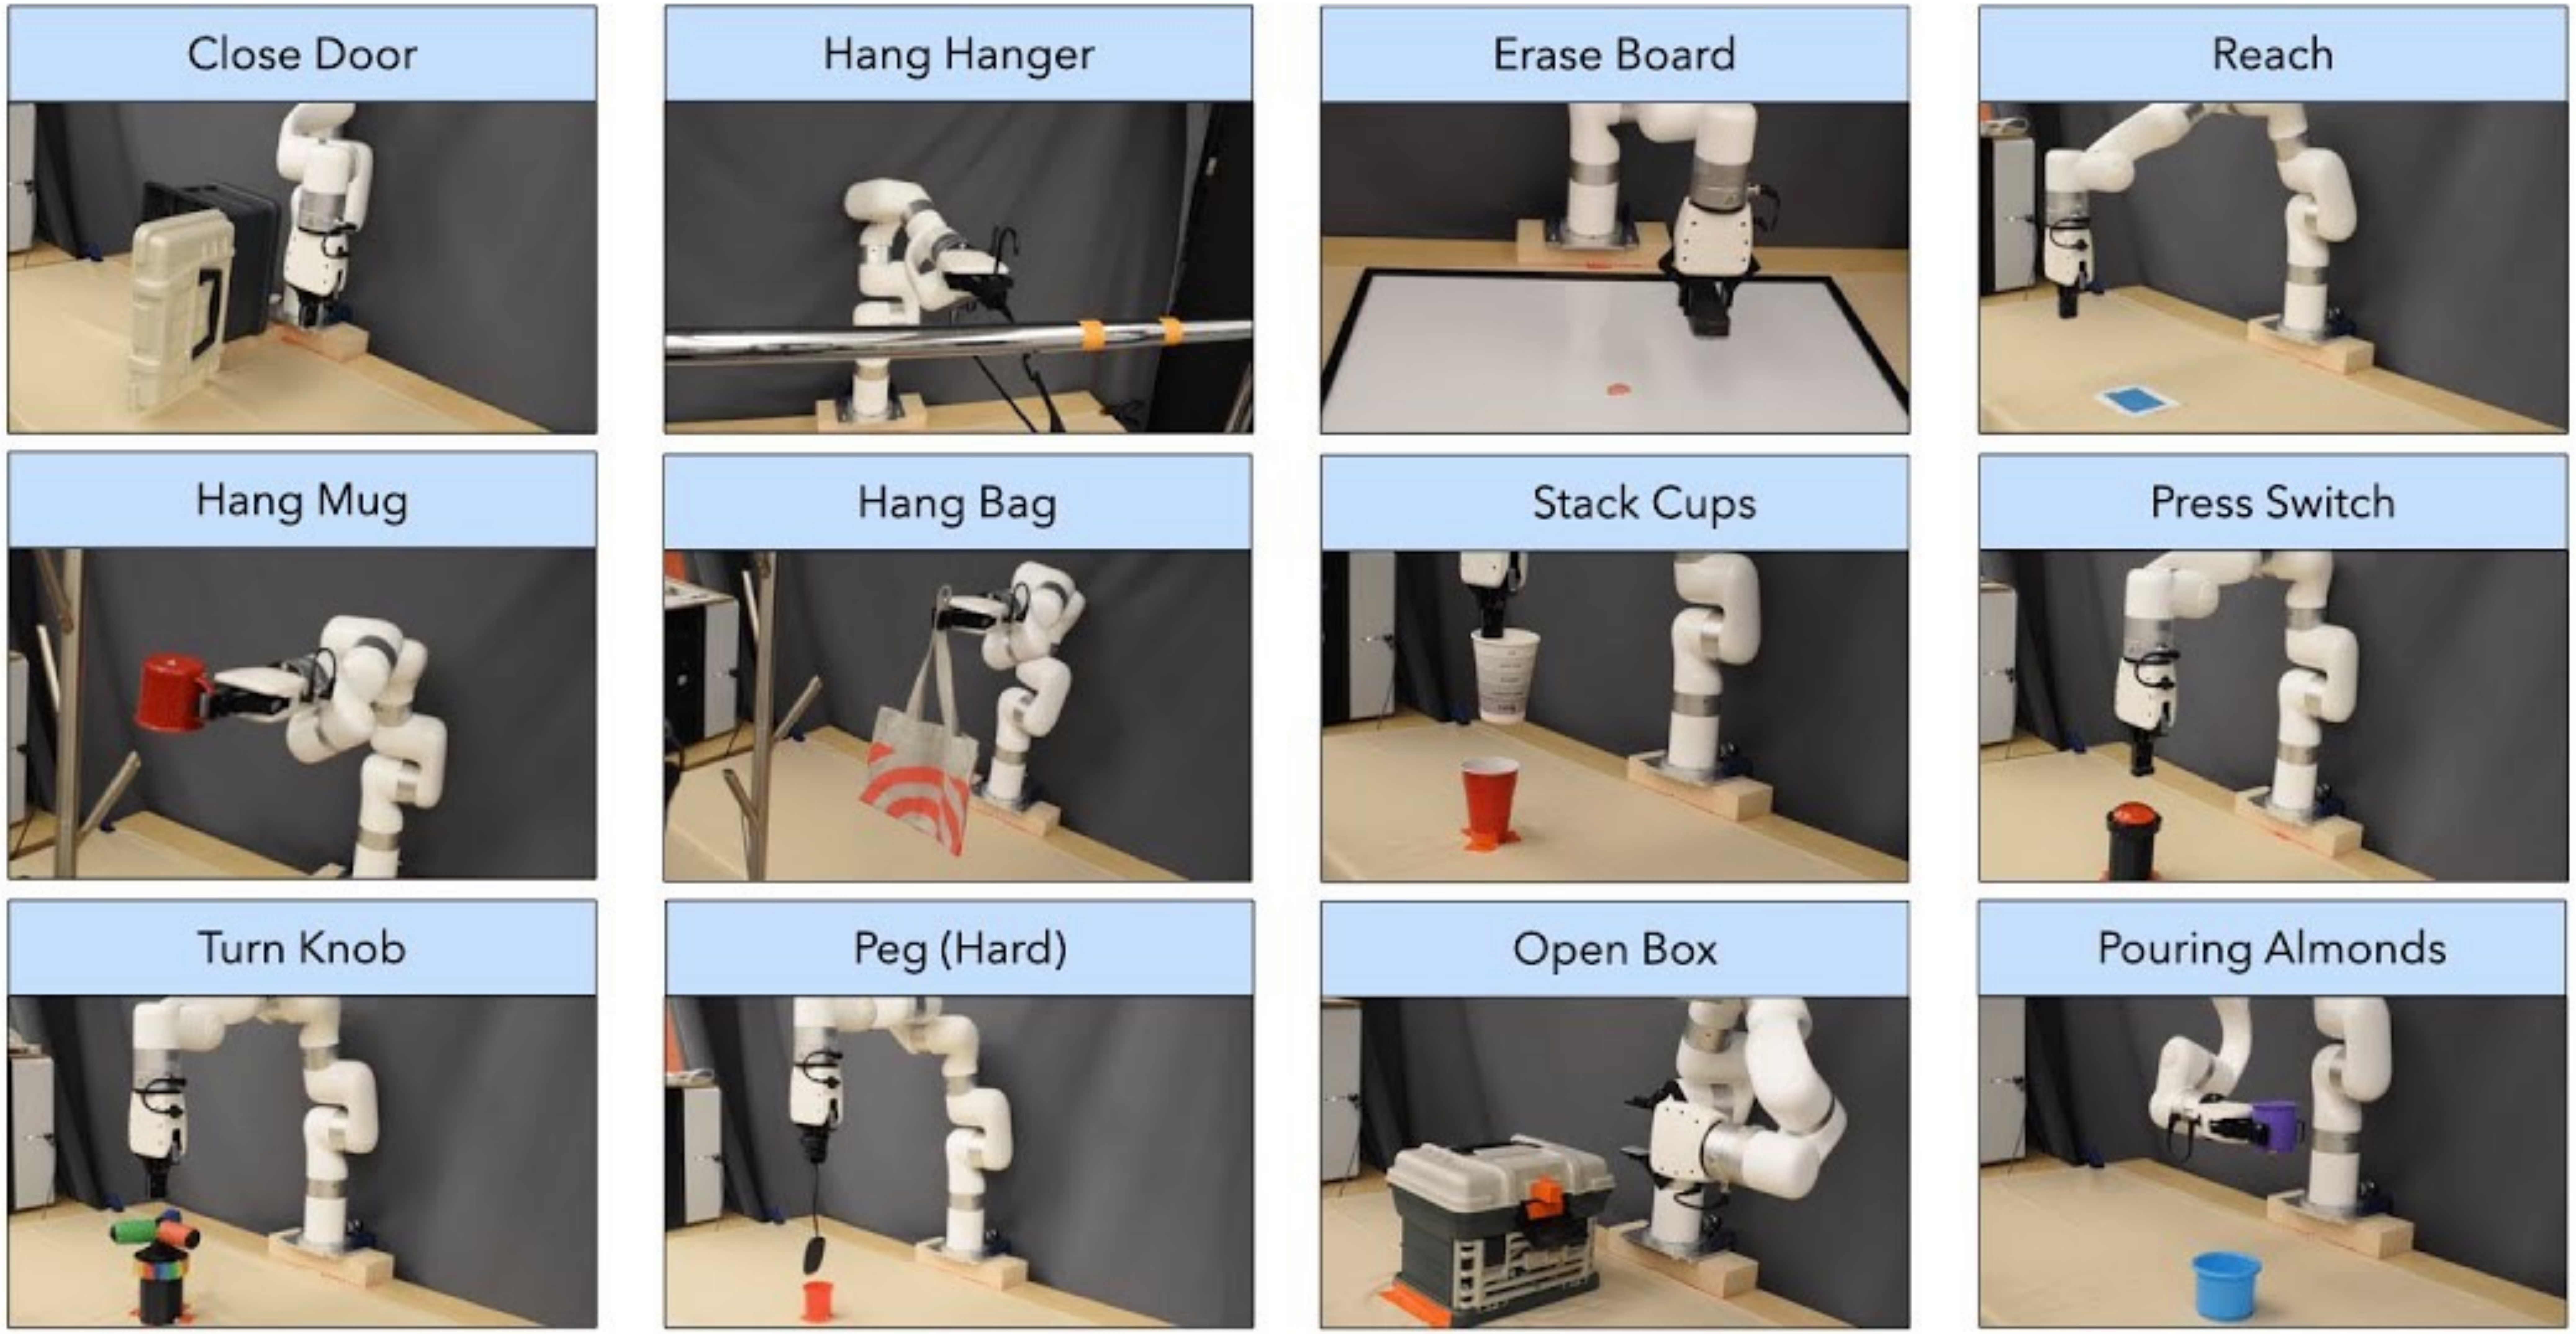


Real-time Mapping of
Physical Scene Properties
with an
Autonomous Robot Experimenter
Iain Haughton, Edgar Sucar, Andre Mouton,
Edward Johns, Andrew J. Davison
Dyson Technology Ltd.
Imperial College

Temporal Logic Imitation: Learning Plan-Satisficing Motion Policies from Demonstrations

RILAB
Robot Learning Laboratory

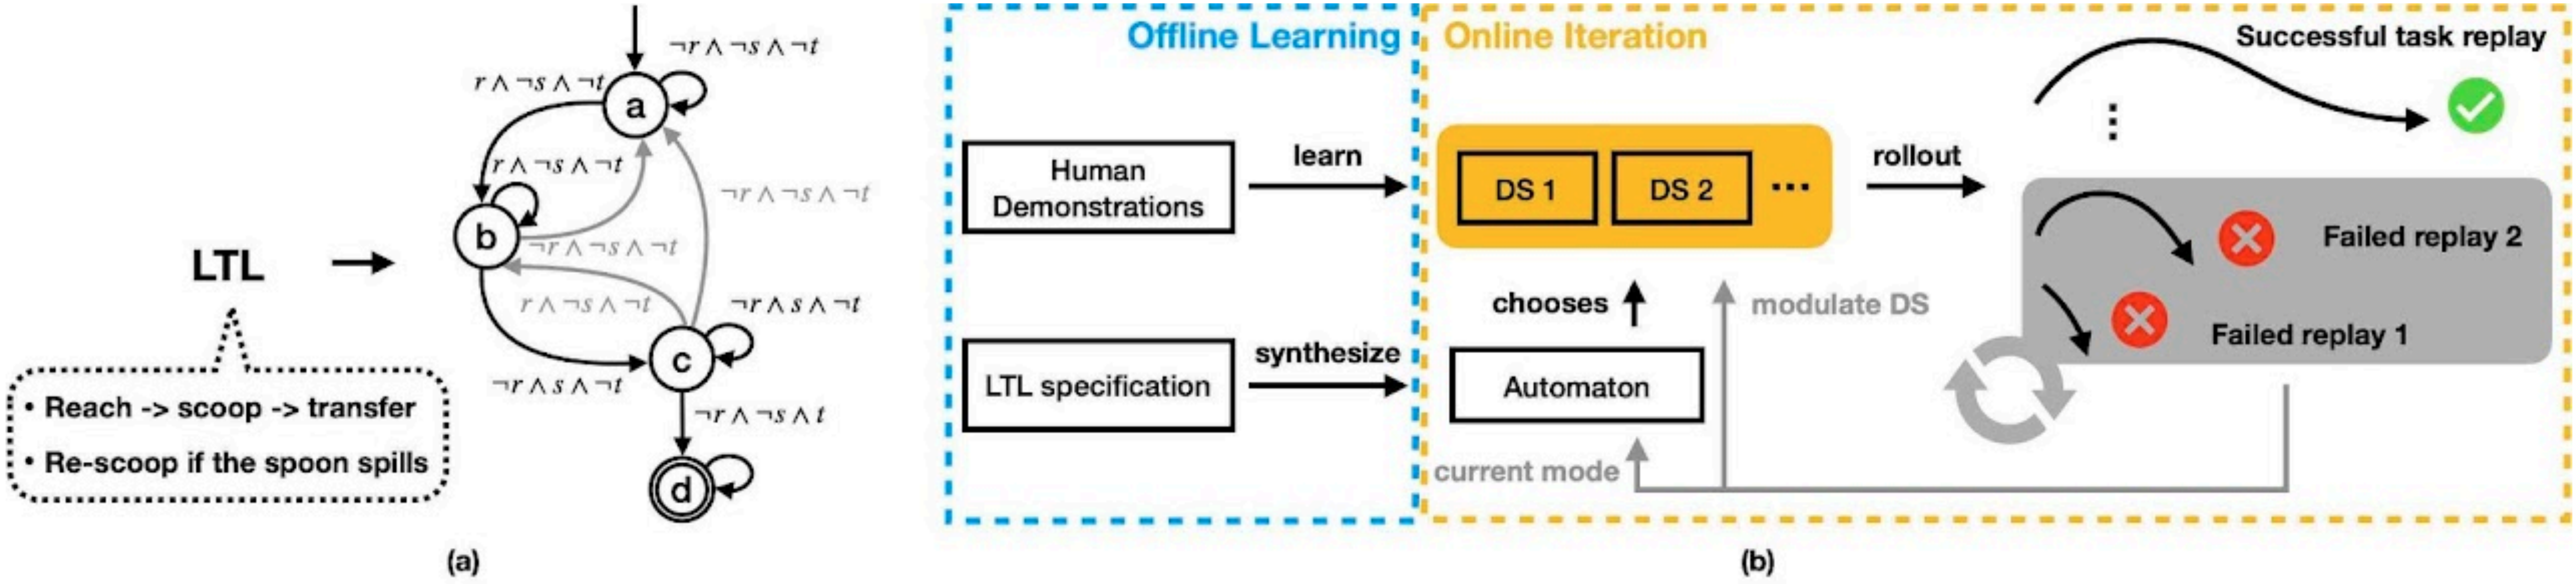

Collecting Demonstrations for Scooping Task

Take-home Messages

Robot Learning Laboratory

- CoRL: the best conference for robot learning.
- Large language model is an inevitable technology.

Oral & Poster Presentation

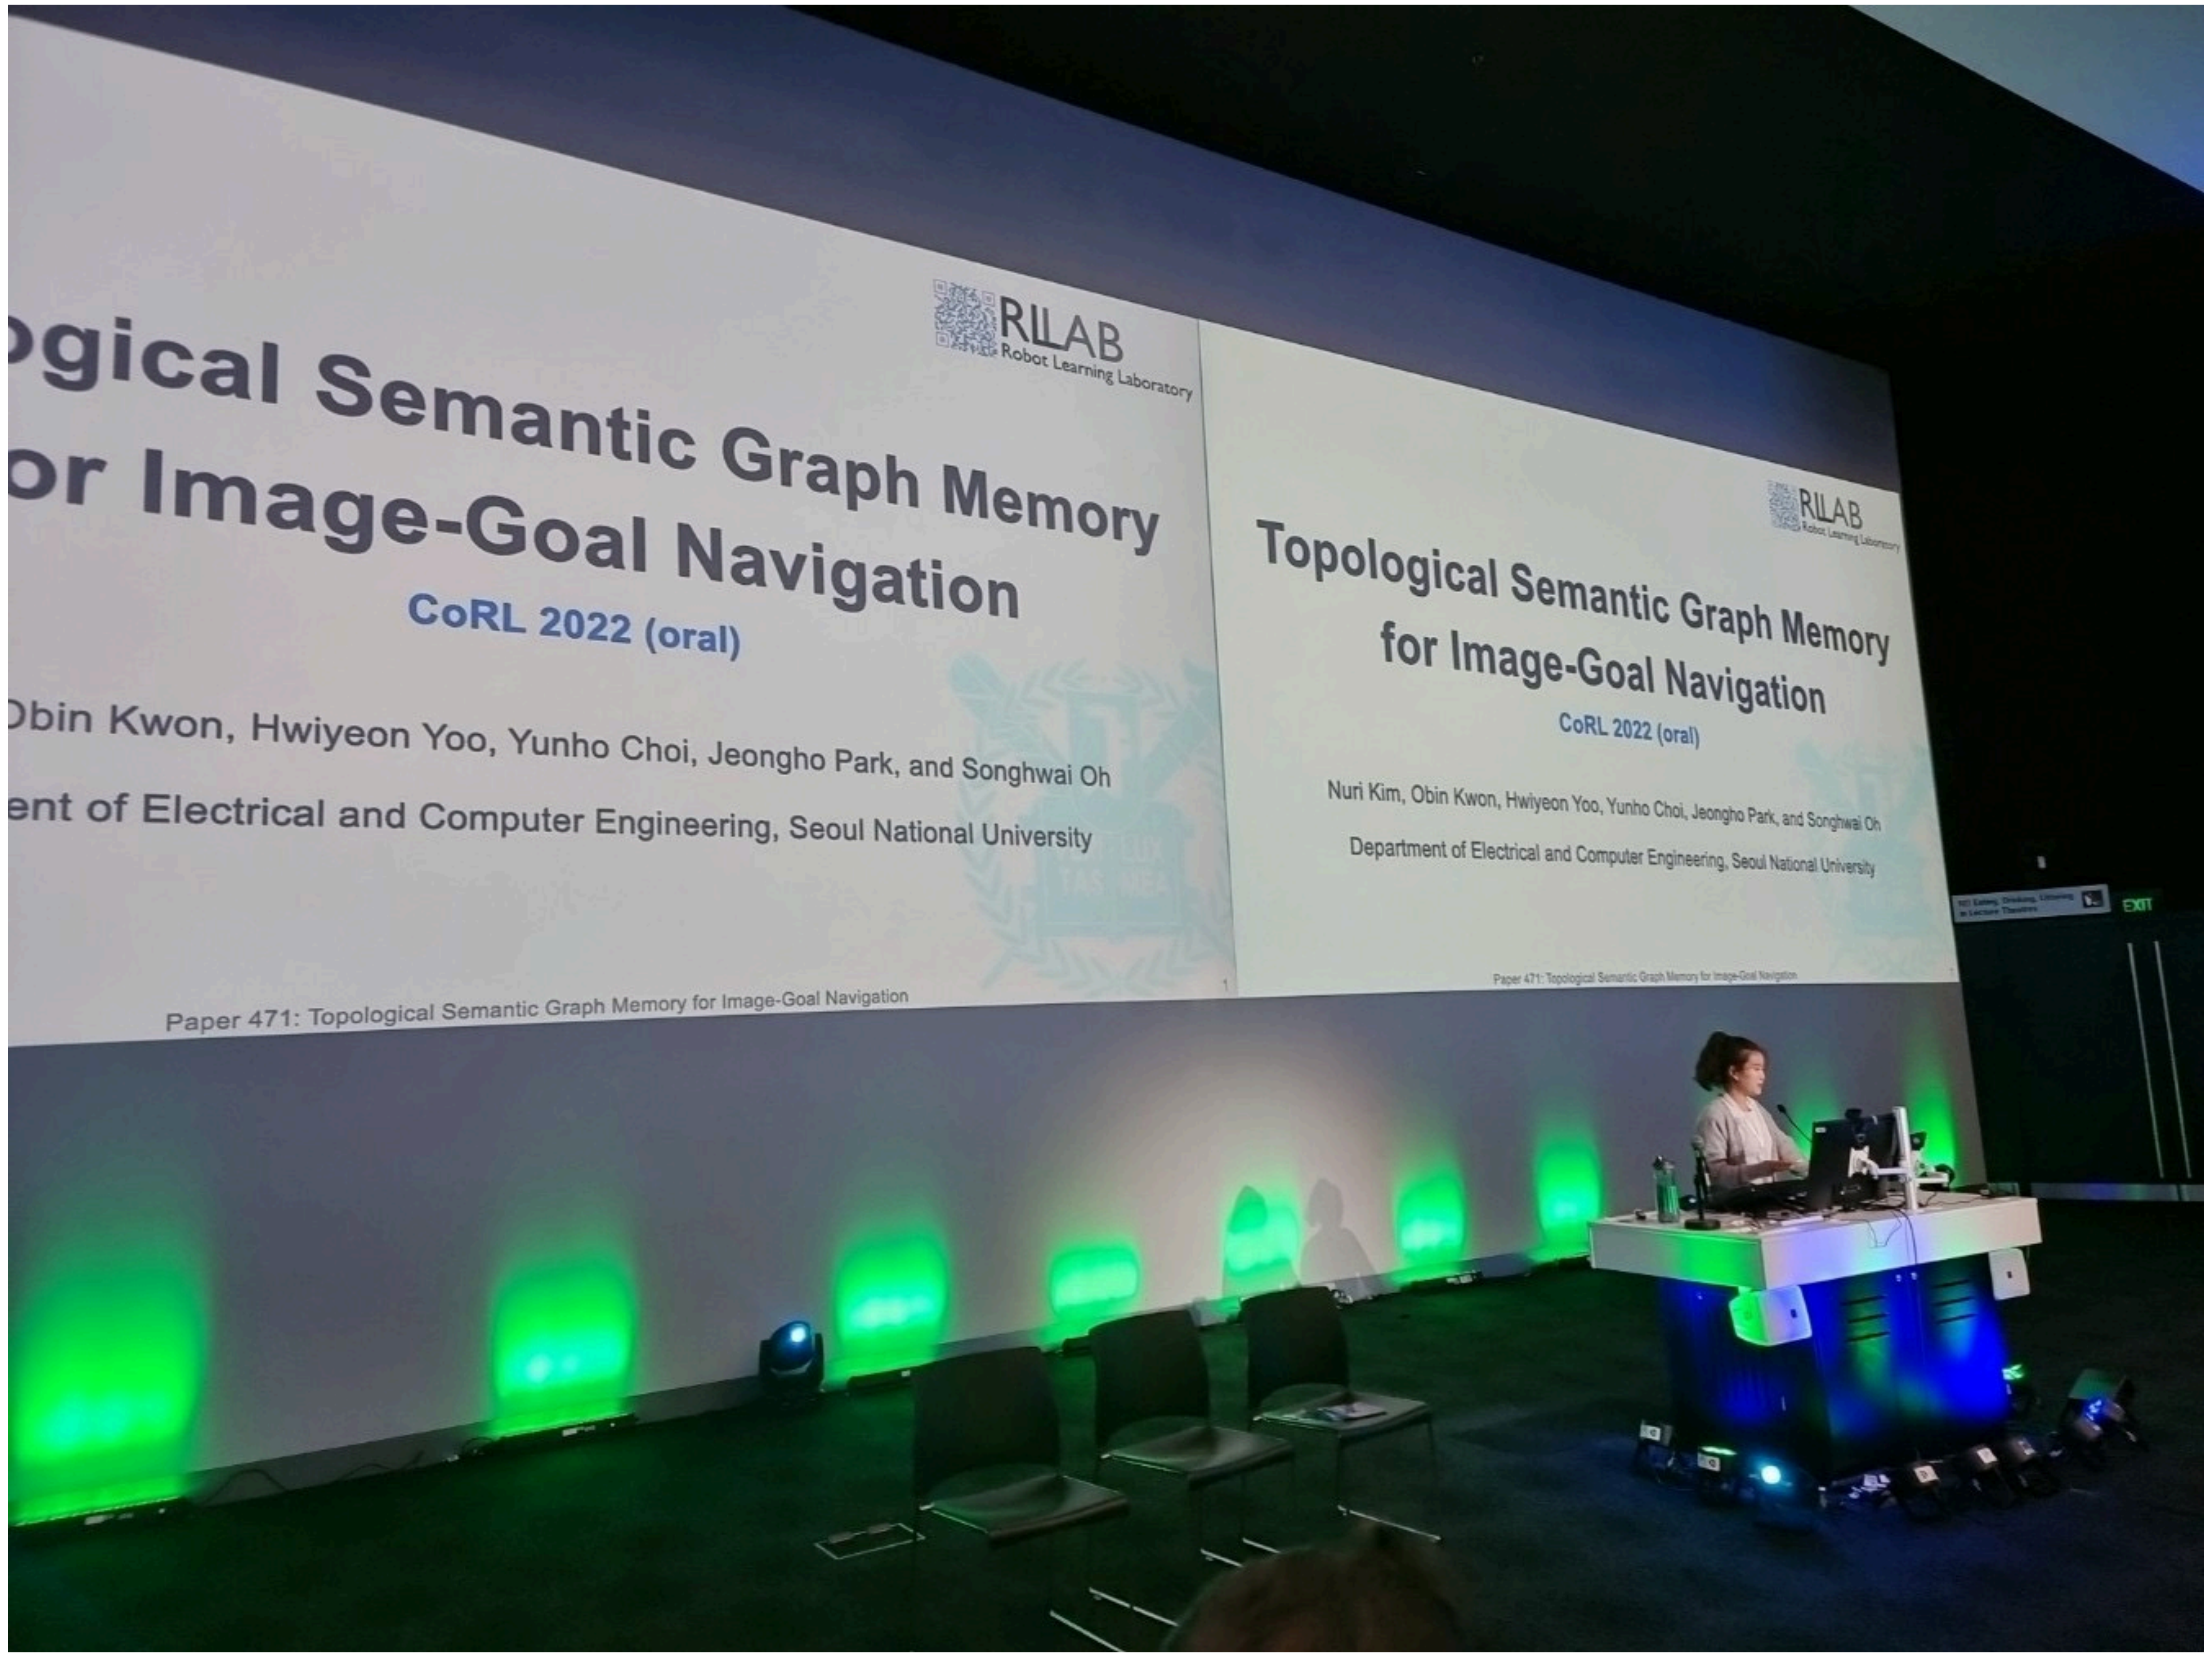

Thank you for your attention