CS283: **Robotics Spring 2025: DL & Ethics**

Sören Schwertfeger / **师泽仁**

ShanghaiTech University

AI for Robotics vs Embodied AI

- AI for Robotics:
  - Improve Robot Capabilities by employing various AI techniques
    - E.g.: https://www.nvidia.com/en-us/industries/robotics/
- Embodied AI:
  - Improve AI Capabilities by giving it a body
    - Increase intelligence through interaction with the physical world:
    - See; Talk; Listen; Act; Reason
    - https://embodied-ai.org/

Previously in Robotics

- Everything hand-coded
  - Computer vision (e.g. SIFT, SURF, …)
  - Control (e.g. PID, MPC, …)
  - Voice Recognition, Dialogue Systems
  - Planning
  - Navigation
  - SLAM
  - …
- Needs to be carefully tuned
  - E.g. DARPA Robotics Challenge 2015

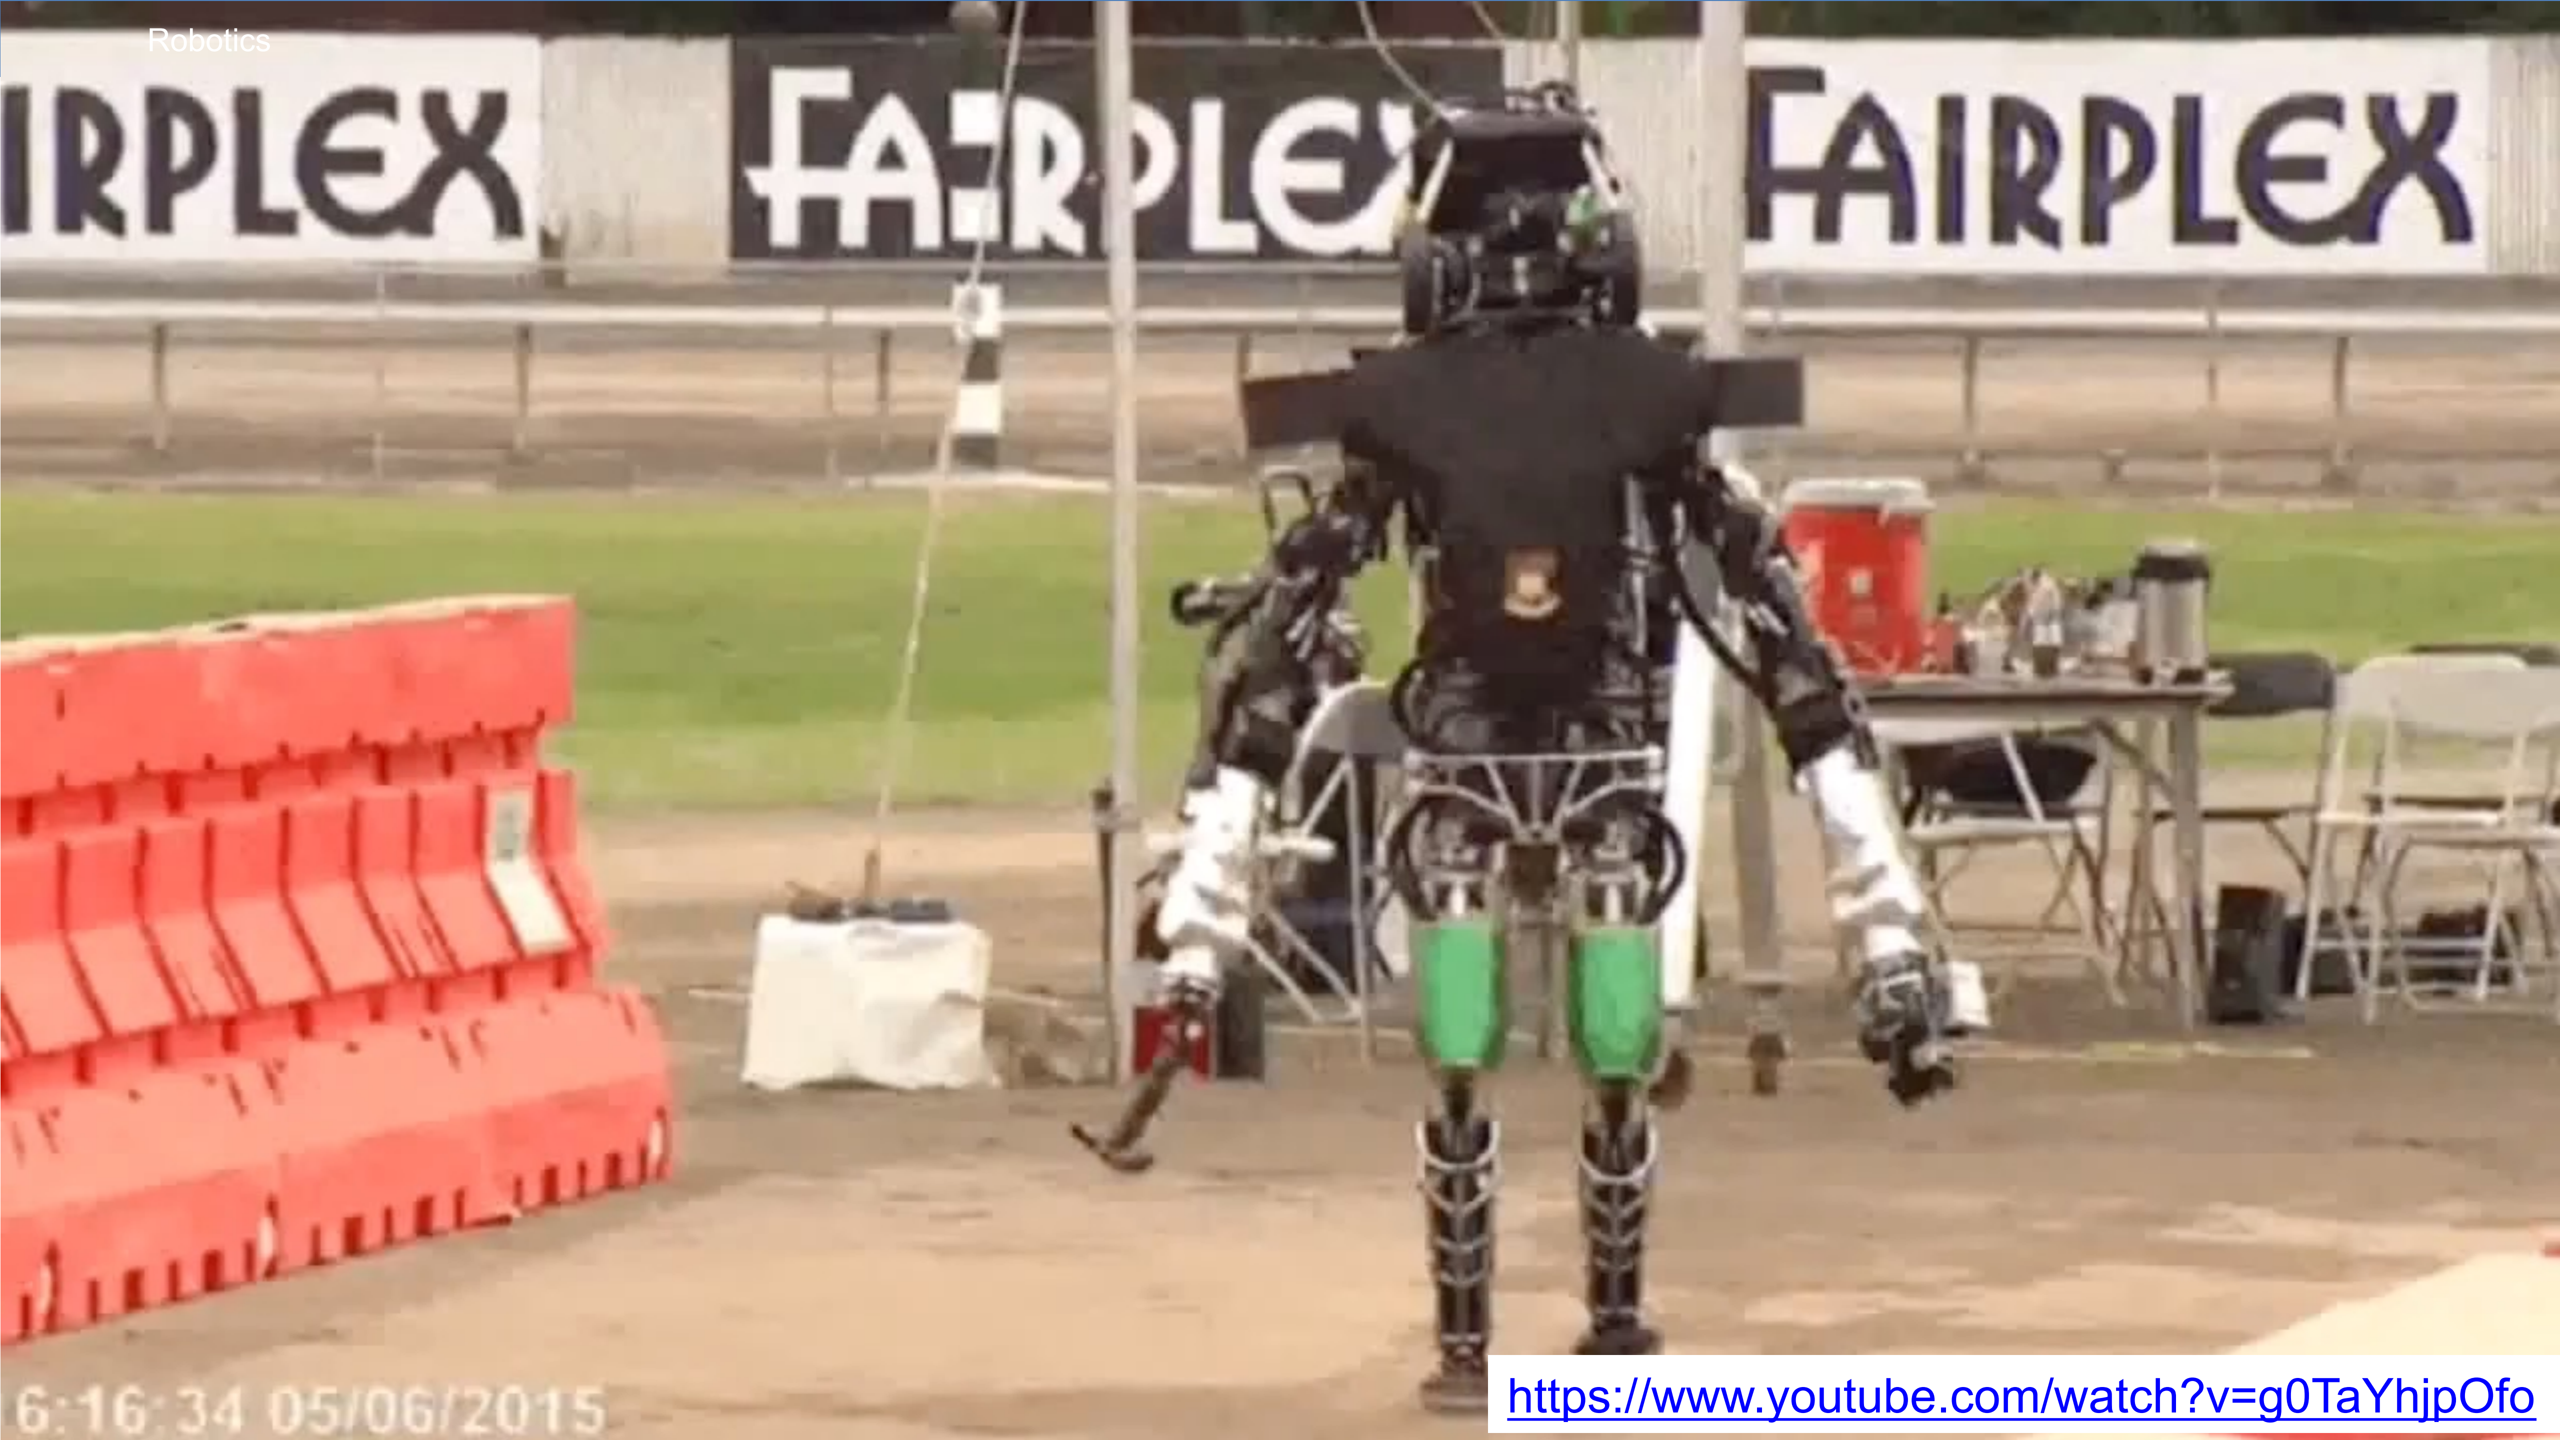

https://www.youtube.com/watch?v=g0TaYhjpOfo

Modern Robotics with the help of Learning/ AI

- Learning for Robotics: Train a Neural Network to do the things that are too hard to program by hand, e.g.:

- DL based computer vision and recognition (e.g. also 3D point clouds)
- Reinforcement Learning for control
- LLM for human robot interaction and intelligence
- …

Overview of Learning Approaches in Robotics

- **Goal:** To explore various learning paradigms that enable robots to perform tasks autonomously.
- **Categories:**
  - Model-Free vs. Model-Based Learning
  - Supervised vs. Unsupervised Learning
  - Passive vs. Active Learning
  - Reinforcement Learning (RL)
  - Imitation Learning
  - End-to-End Deep Learning
  - Actor-Critic Learning
  - Evolutionary Algorithms
  - Transfer Learning
  - Self-Supervised Learning
  - Few-Shot and Zero-Shot Learning
  - Multi-Agent Learning
  - Curriculum Learning
  - LLM
  - Foundation Models
  - Other types of “learning”

ROBOTIC LEARNING

Model-Based vs Model-Free Learning

Model-Based Learning:

- Involves learning a model of the environment or dynamics (e.g., using physics or system dynamics).
- Robot can plan and predict actions based on this model.
- **Example:** Planning with a learned dynamics model in robotic control tasks.

Model-Free Learning:

- Directly learns a mapping from states to actions or rewards without modeling the environment.
- **Example**: Q-learning or policy gradient methods in Reinforcement Learning.

Supervised vs. Unsupervised Learning

Supervised Learning:

- Learning from labeled data (input-output pairs).
- Requires large datasets and human supervision.
- **Example:** Image classification for object detection in robotics, such as recognizing "graspable" objects in a scene.

Unsupervised Learning:

- Learning from unlabeled data to find hidden patterns (e.g., clustering or representation learning).
- **Example:** A robot exploring its environment autonomously to cluster sensory data (e.g., LIDAR or visual data) into distinct regions like walls, furniture, or open spaces. This clustering can later help the robot map its environment for navigation.

Passive vs Active Learning

Passive Learning:

- The robot learns from a fixed dataset (either labeled or unlabeled).
- **Example:** Supervised learning with a fixed dataset.

Active Learning:

- The robot queries the environment for more informative data based on its current knowledge or uncertainty.
- **Example:** A robot selects which objects to interact with in order to maximize learning.

End-to-End Deep Learning in Robotics

- **Definition:** Learning a direct mapping from raw input (e.g., images, sensory data) to the output (e.g., control commands).
- **Example:** A robot controlling a gripper using only camera images.
- **Advantages:** Simplifies the pipeline by learning a direct mapping.
- **Challenges:** Requires large amounts of labeled data.

Reinforcement Learning (RL)

- **Definition:** An agent learns to take actions in an environment to maximize cumulative reward over time.
- **Key Components:** States, actions, rewards, policy.
- **Example:** Training a robot to navigate using trial-and-error.
- **Types:**
  - **Model-Free:** Methods like Q-learning, policy gradients.
  - **Model-Based:** Use of learned models to simulate and plan actions.

Imitation Learning

- **Definition:** Robots learn by observing and imitating human demonstrations or expert behaviors.

- **Approaches:**
  - **Behavior Cloning:** Supervised learning from demonstrations.
  - **Inverse Reinforcement Learning (IRL):** Learning the underlying reward function from expert demonstrations.

- **Example:** Teaching a robot to grasp objects by mimicking human actions.

LLM for Robotics

- **Definition:** LLMs are AI systems trained on massive text corpora to process, understand, and generate human-like text.
- **Key Capabilities in Robotics:**
  - **Natural Language Understanding:** Interpreting commands and queries.
  - **Knowledge Integration:** Retrieving and applying knowledge to tasks (e.g., assembly instructions).
  - **Reasoning and Task Decomposition:** Breaking down complex instructions into actionable steps.
- **Advantages:**
  - Provides high-level reasoning and task planning.
  - Reduces the need for detailed programming in language-based tasks.
  - Can handle diverse instructions using pre-trained knowledge

How LLMs Are Used in Robotics

- **Applications:**
  - **Human-Robot Interaction:** Robots can interpret and execute natural language instructions (e.g., “Bring me a cup of water”).
  - **Task Planning:** Combining linguistic reasoning with real-world task execution.
  - **Multi-Modal Integration:** Enhancing decision-making by linking text, vision, and sensory inputs.
- **Challenges:**
  - Ensuring grounding in physical environments (e.g., interpreting "left" in a spatial context).
  - Real-time response constraints due to the size of models.
  - Domain-specific fine-tuning for robotics applications.

Robotics Foundation Models

- **Definition:** Large-scale AI models pre-trained on diverse, multi-modal datasets (e.g., text, images, videos).
- **Core Characteristics:**
  - **Versatile Pre-training:** Serve as a base for fine-tuning on specific tasks.
  - **Multi-Modal Understanding:** Integrate text, vision, and other sensory inputs for broader applicability.
- **Key Advantages for Robotics:**
  - Generalize across multiple tasks with minimal retraining.
  - Simplify the training pipeline by leveraging shared representations.
  - Adaptable to new tasks without extensive data collection.

How Foundation Models Empower Robotics

- **Applications:**
  - **Perception:** Models like CLIP interpret visual data for scene understanding.
  - **Control:** Leveraging shared representations for motion planning and actuation.
  - **Task Generalization:** Performing varied tasks without task-specific training.
  - **Simulation-to-Real Transfer:** Reducing the gap between simulated and real-world performance.
- **Challenges:**
  - High computational costs for pre-training and fine-tuning.
  - Limited grounding in physical dynamics without additional modeling.
  - Potential biases from pre-training on non-robotic data.

RL Algorithms

• Finite Markov Decision Processes	MDP
• Temporal-Difference Learning	TD Learning
• State-Action-Reward-State-Action	SARSA TD Learning
• Q-learning: Off-policy TD Control	
• Deep Q-Networks	DQN

- Policy Gradient Methods
  - Actor-Critic Methods

- Asynchronous Reinforcement Learning

Some Examples

Robot Learning

Cognitive Intelligence

Athletic Intelligence

Manipulation

Locomotion/
Control

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

Customers want a robot that handles

all household tasks

and is **commanded by natural language**

Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware (ACT / ALOHA)

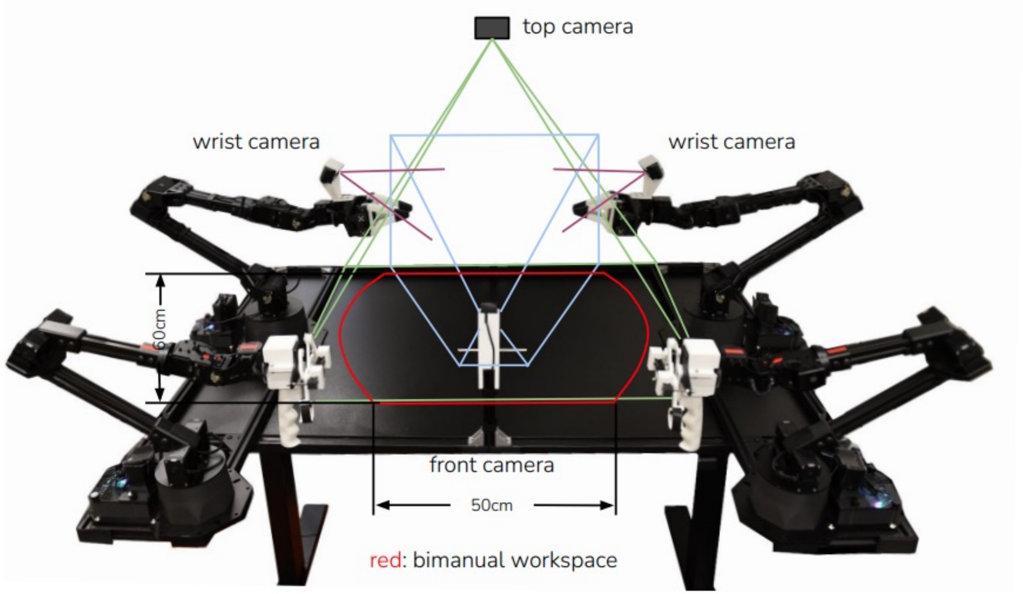

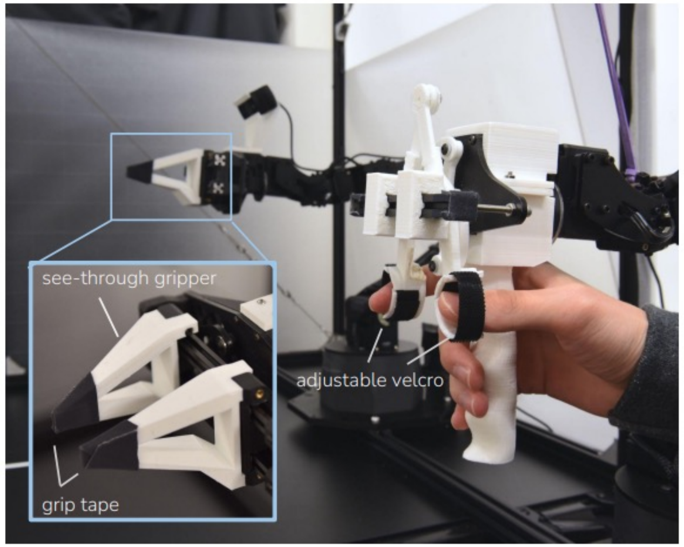

ViperX 6dof Arm (follower)

#Dofs	6+gripper
Reach	750mm
Span	1500mm
Repeatability	1mm
Accuracy	5-8mm
Working Payload	750g

Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware (ACT / ALOHA)

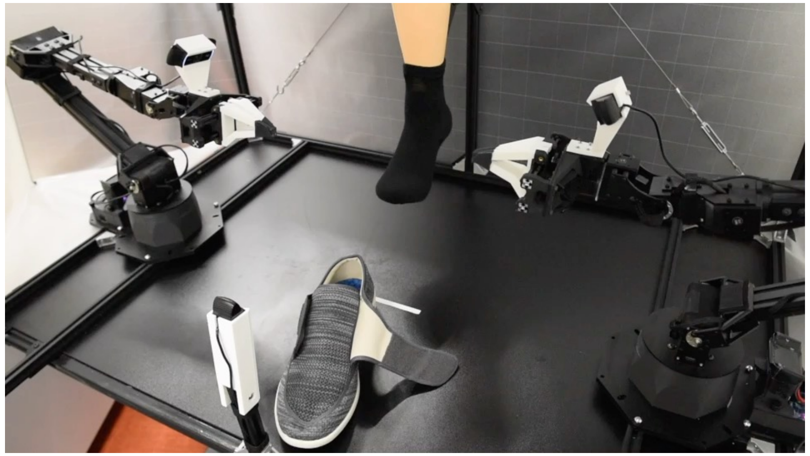

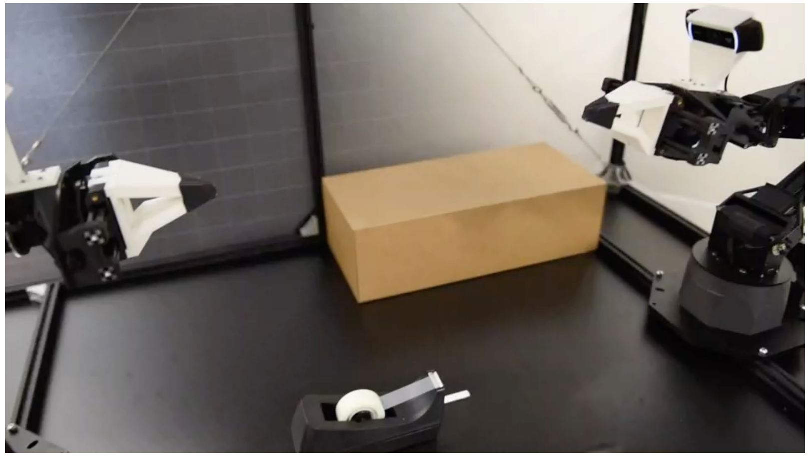

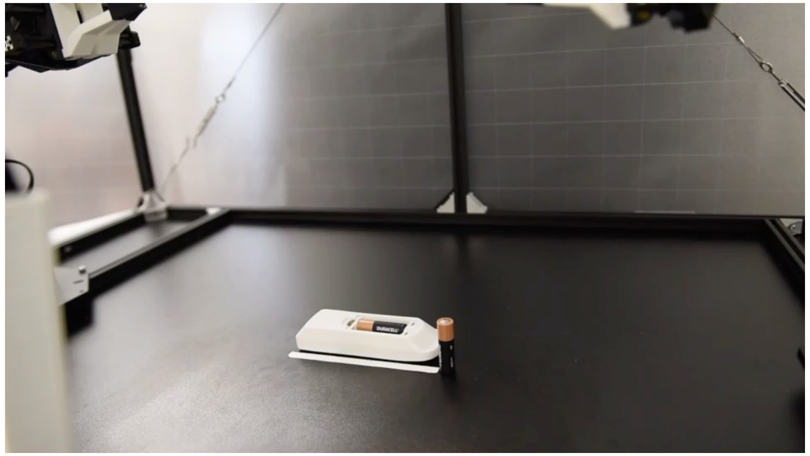

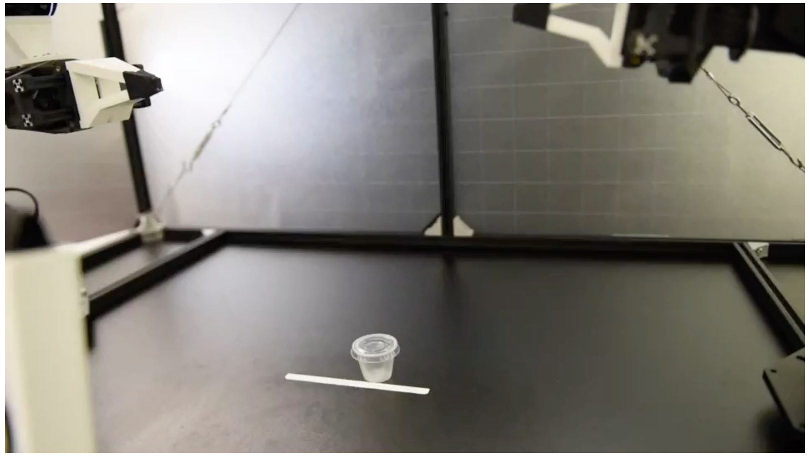

Athletic Intelligence (quadruped robot)

2020: RL is able to work on quadruped locomotion
2021: RL is simple enough to train quadruped robot
2022: quadruped robot can utilize vision to guide the gait
2023: quadruped robot outperforms all other mobile robots

Athletic Intelligence (bipedal robot)

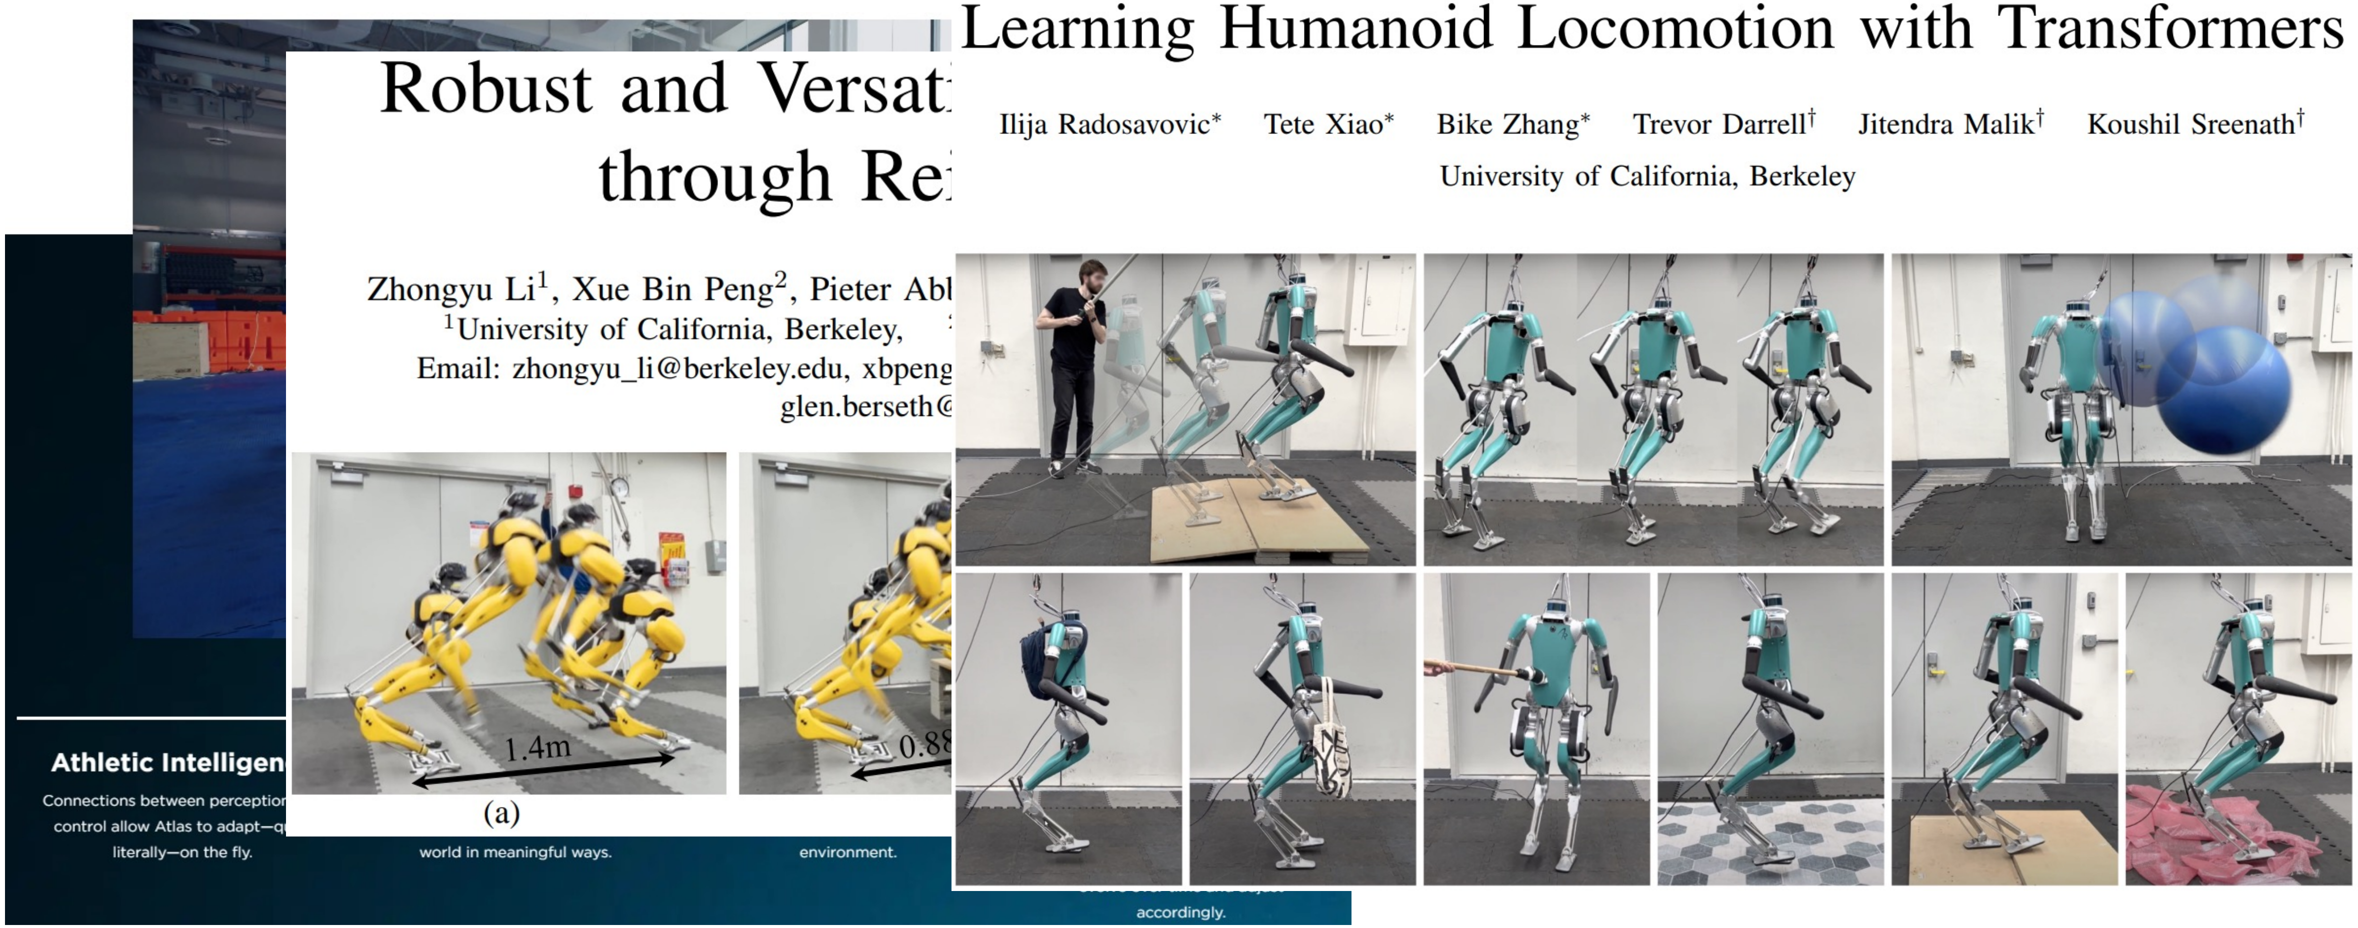

Robot Parkour Learning

- end-to-end vision-based parkour learning (depths images)
- RL pre-training with soft dynamics constraints +
- RL fine-tuning with hard dynamics constraints

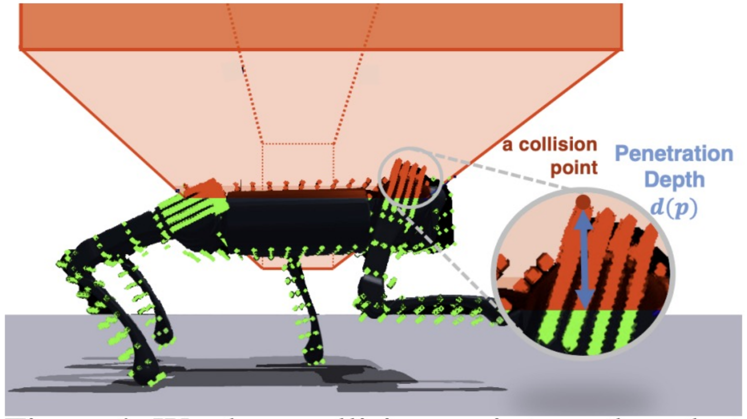

Figure 4: We show collisions points on the robot. Collision points that penetrate obstacles are in red.

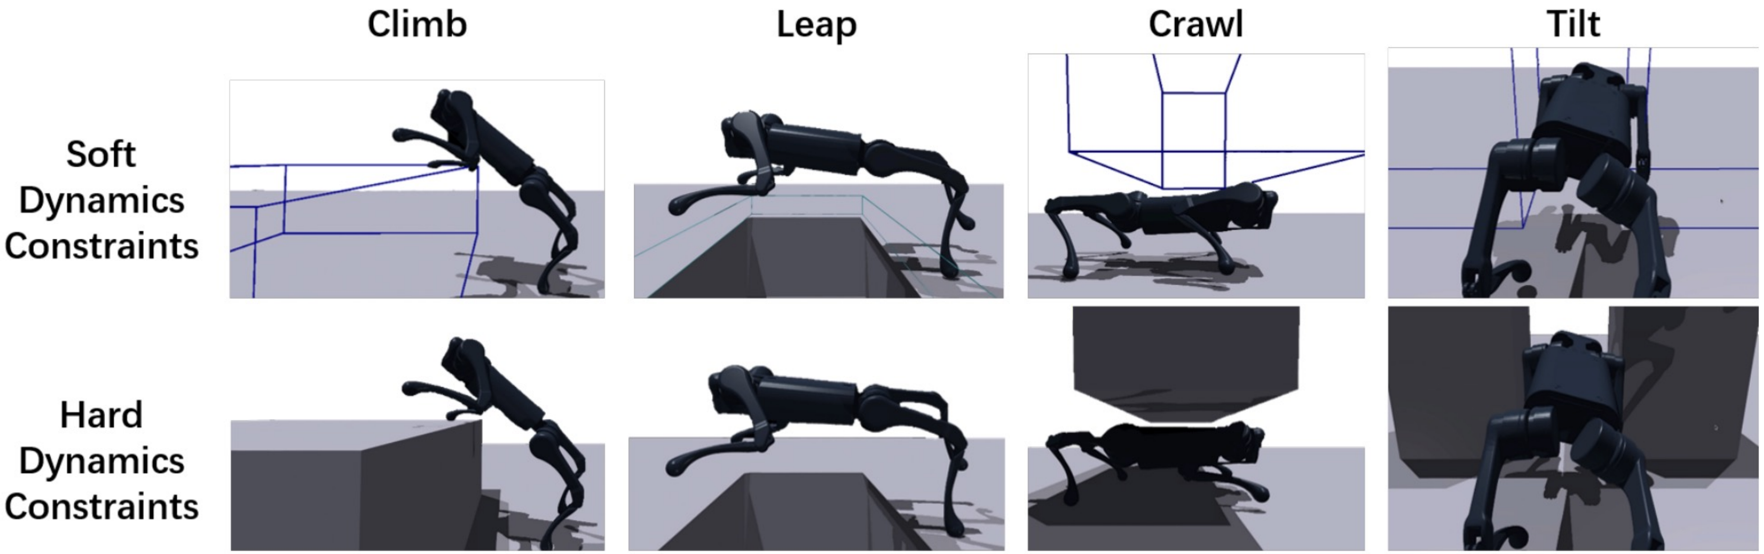

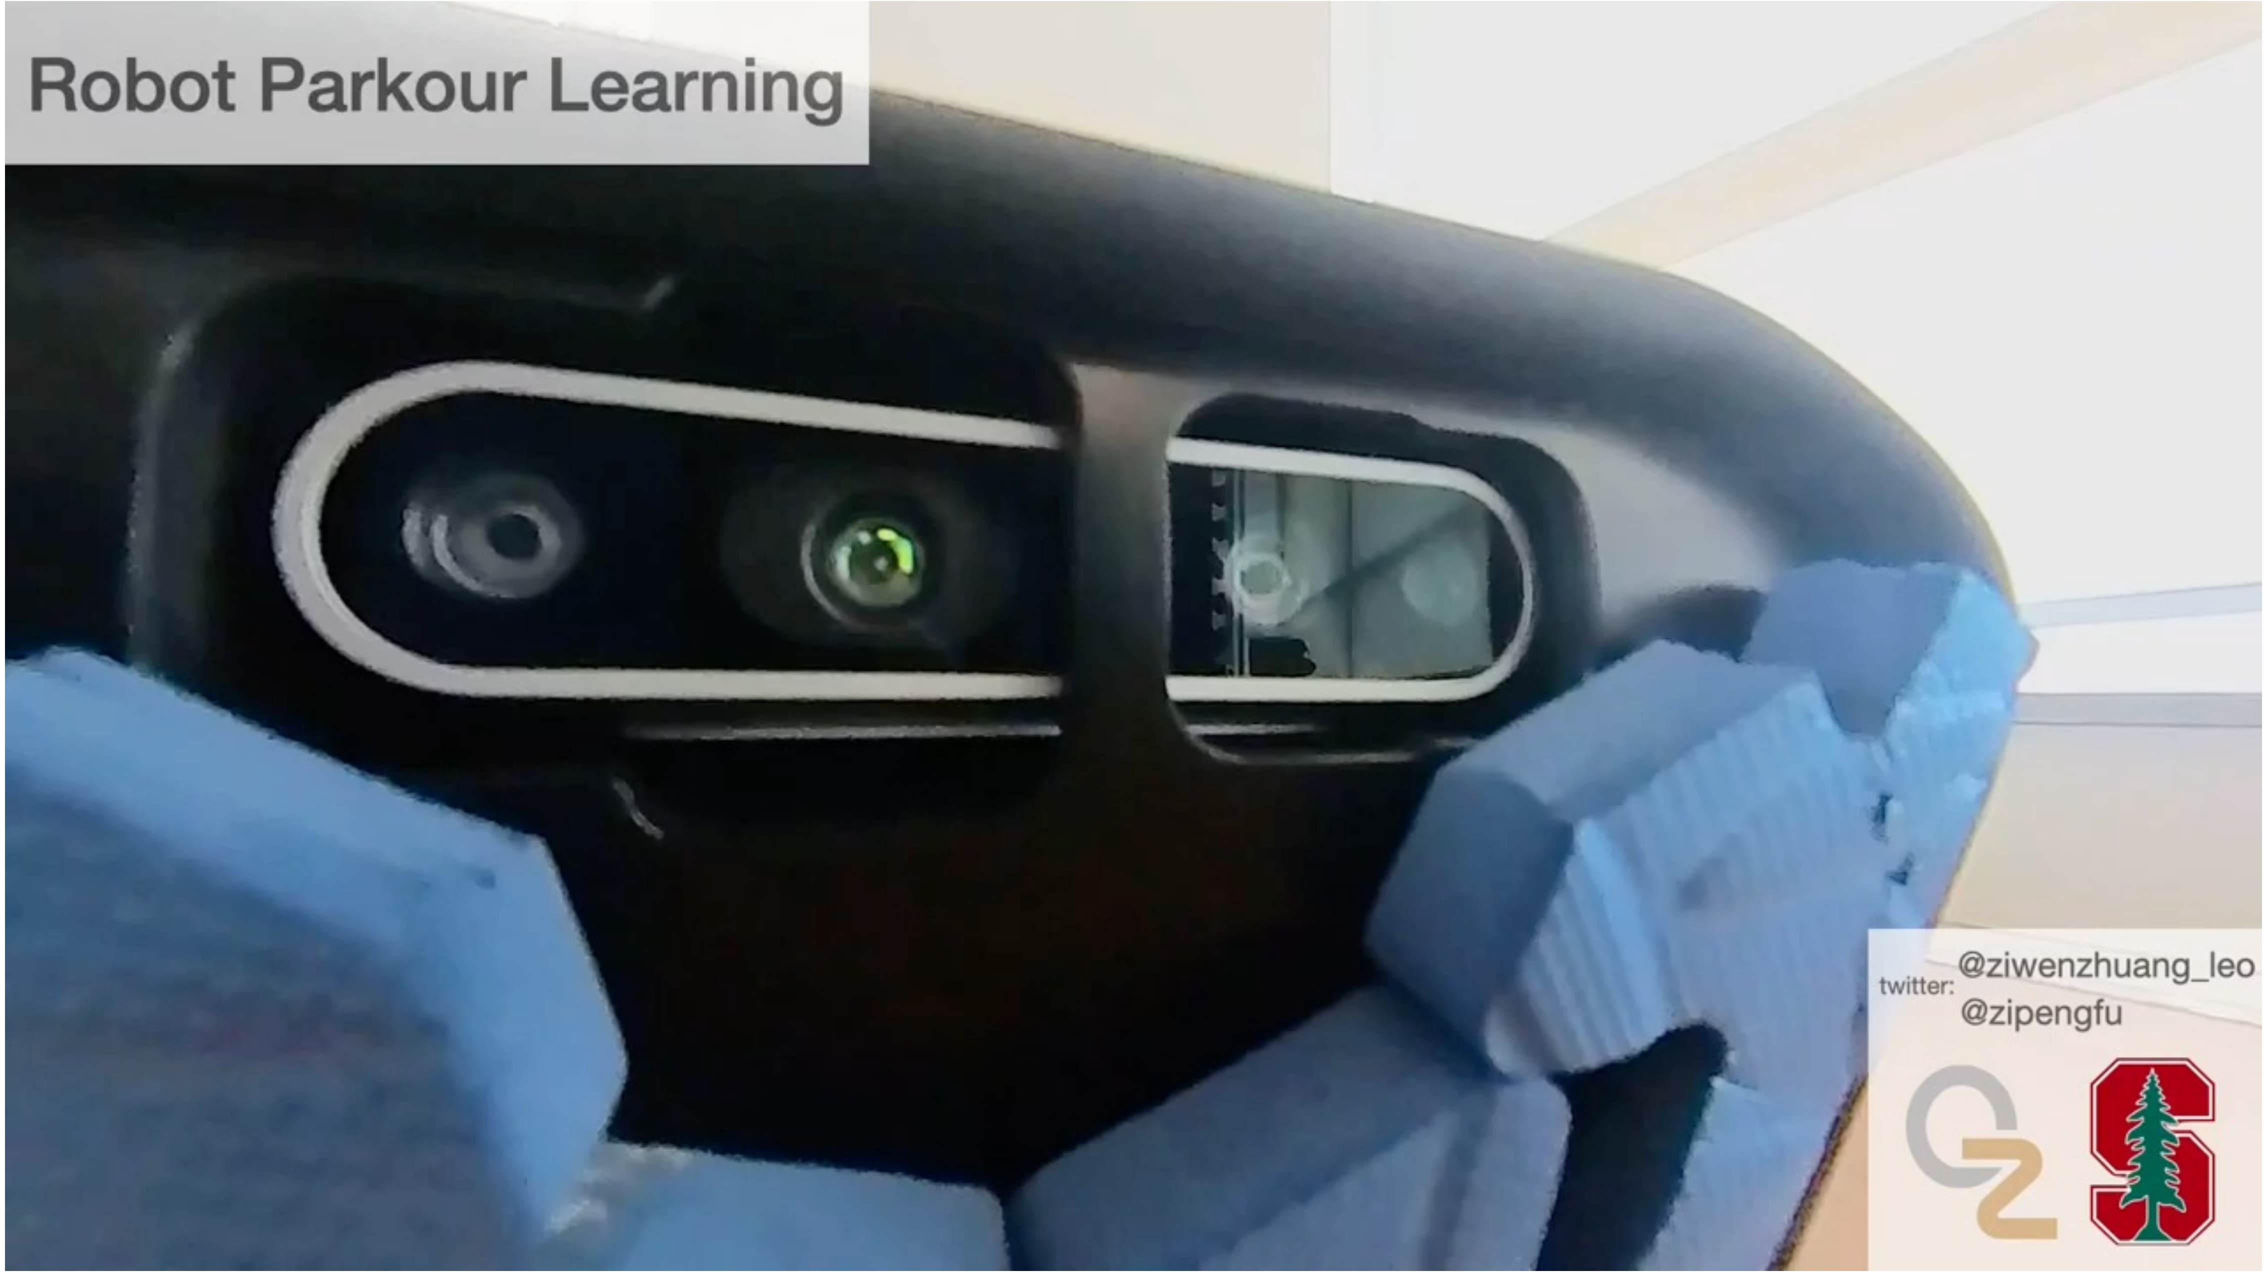

https://robot-parkour.github.io/

Playful DoggyBot: Learning Agile and Precise Quadrupedal Locomotion

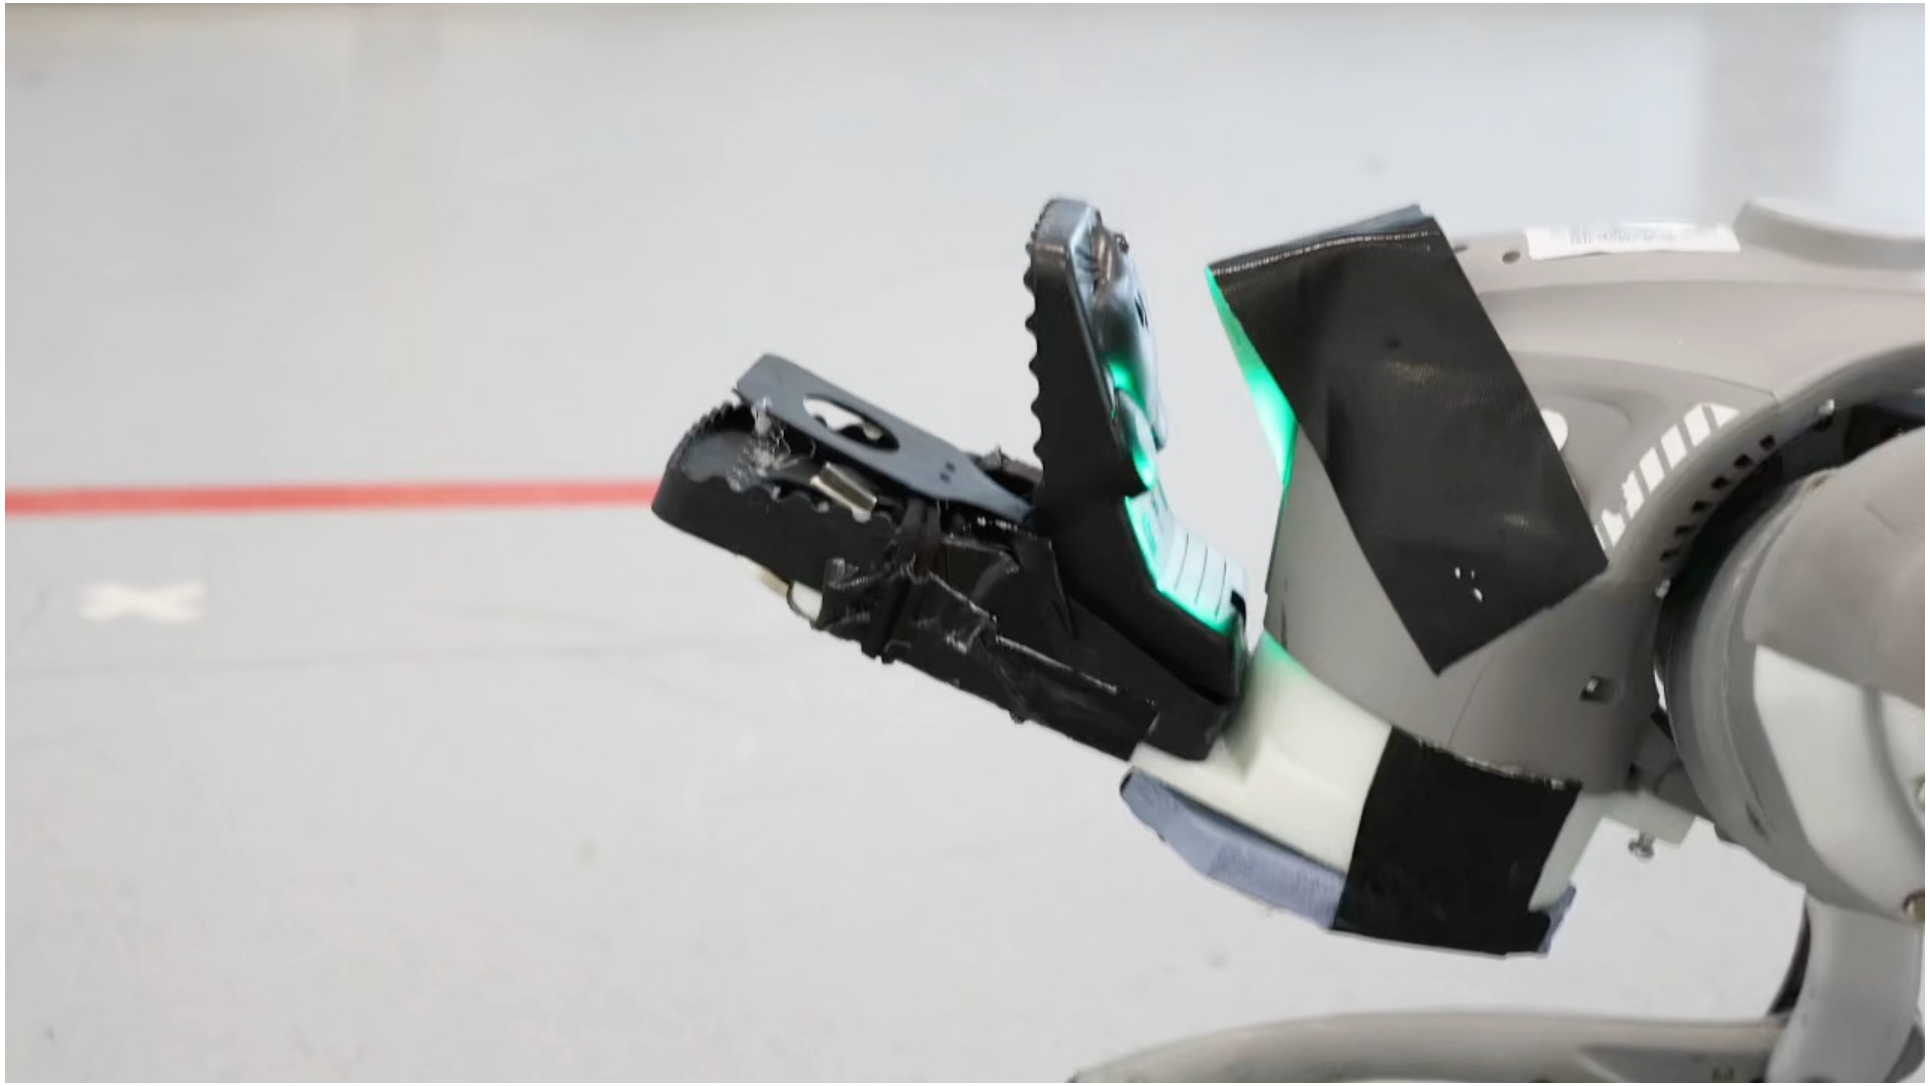

https://playful-doggybot.github.io/

Enhancing Embodied AI Performance by Integrating Traditional Algorithms with Large Language Models for Map-Based Tasks

Fujing Xie

Advisor: Prof. Dr. Sören Schwertfeger

ShanghaiTech University

Large Language AI Models for Autonomous Mobile Robotics

- Visual language manipulation

https://palm-e.github.io/

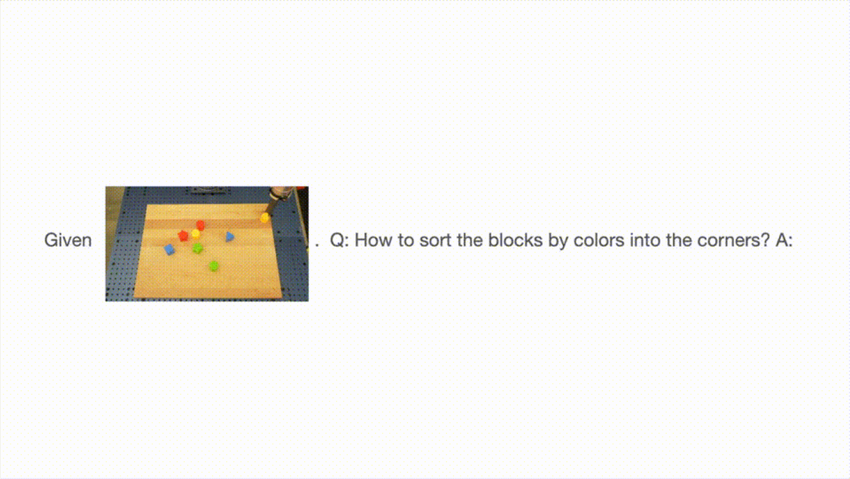

Task: How to sort the blocks by colors into the corners?

The input is the task and the images, the output of PaLM-E is command the robot could execute.

- Visual language navigation

https://www.microsoft.com/en-us/research/group/autonomous-systems-group-robotics/articles/chatgpt-for-robotics/

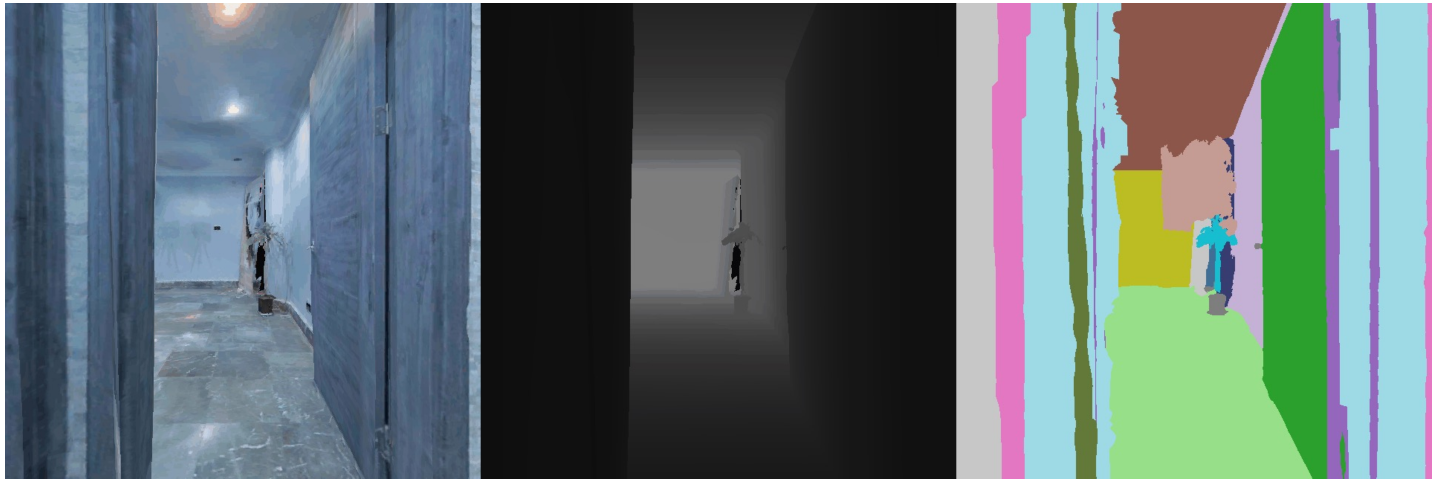

Task: navigate to the stairs

Objects in the scene: <door, 1.15 m, 12>, <door, 0.79 m, 15>, <door, 4.43 m, 0>, <plant, 3.36 m, 2>, <column, 3.40 m, 4>, <, 0.63 m, -13>, <door, 0.57 m, -24>, <door, 0.82 m, 4>

Command: move(1.5) # Move straight ahead 1.5 meters to explore the environment.

Objects in the scene: <stairs, 2.99 m, -21>, <objects, 2.68 m, -19>, <door, 2.93 m, 1>, <plant, 1.80 m, 5>, <column, 1.95 m, 7>, <objects, 1.67 m, 5>

Command: turn(-21) # Turn towards the stairs.

Objects in the scene: <stairs, 3.75 m, -3>, <door, 2.04 m, -18>, <objects, 3.30 m, -2>, <door, 2.47 m, 17>, <plant, 1.40 m, 22>, <door, 2.46 m, -11>, <column, 1.48 m, 25>, <objects, 1.29 m, 23>

Command: turn(-3) # Turn towards the stairs.

Objects in the scene: <stairs, 3.75 m, 3>, <door, 2.18 m, -11>, <objects, 3.27 m, 5>, <door, 2.36 m, 26>, <door, 2.56 m, -4>

Command: move(3.75) # Move towards the stairs.

Objects in the scene: <stairs, 1.02 m, -14>

Command: move(1) # Move closer to the stairs.

Why Maps Matter in Real Environment

Notice about Emergency Repair of Sewage Pipes

on the East Side of the Silk Road Canteen and Road Closures

Dear all,

Due to the emergency repair of the sewage pipe under the road on the east side of the Silk Road Canteen, the road between the east side of the Silk Road Canteen and Student Apartment Building 2 must be excavated and renovated (as marked in red in the picture below). During the repair, the road will be closed for 6 days (January 19 – January 24, 2024).

There will be noise and machinery works. Please stay away and go around the area to avoid accidental injuries.

We apologize for any inconvenience this may cause!

Fig. 1 The figure above depicts a real-life situation encountered by a 3rd-party delivery robot on our University campus, where it is blocked by an intersection closure. Below the e-mail sent by Office of General Services announcing this closure is shown.

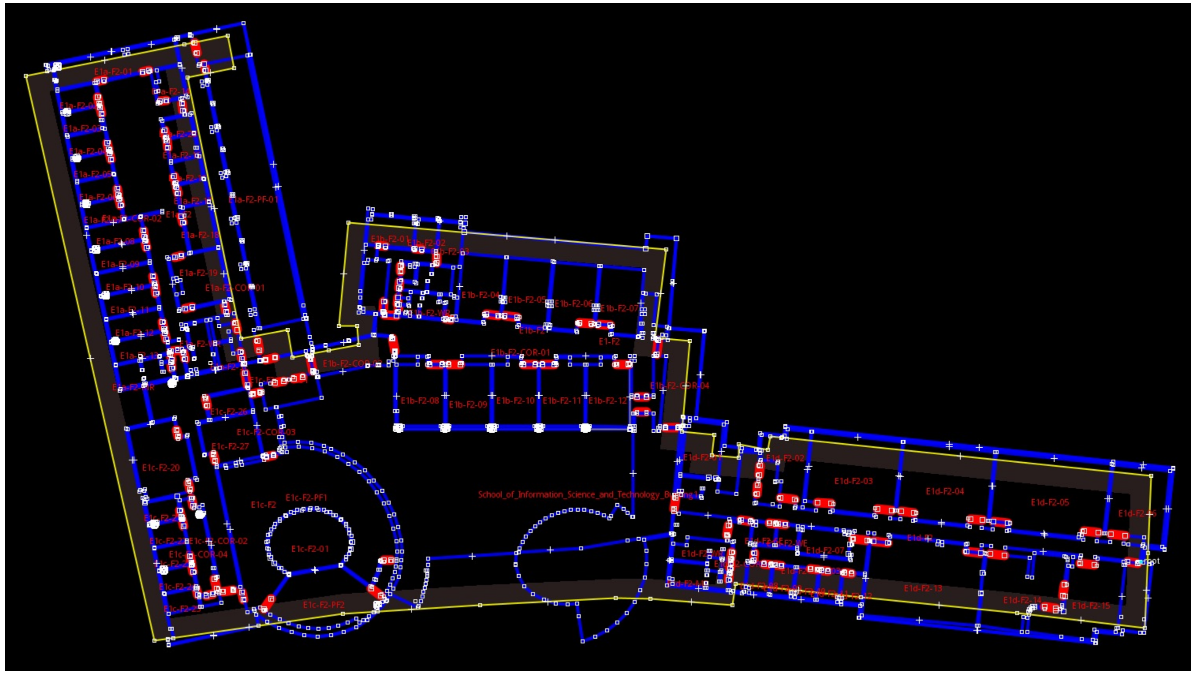

Fig. 2 A osmAG (Area Graph in OpenStreetMap textual format)

Map Representations

- **Traditional Map Representation**: Occupancy Grid Map, Point Cloud Map, Visual Keypoint Map

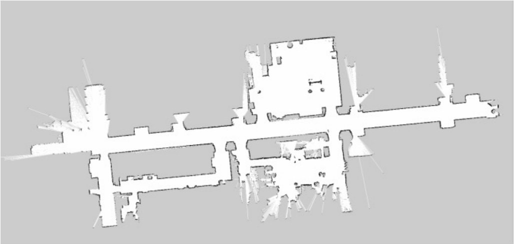
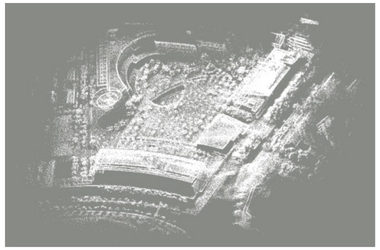


- **osmAG (Area Graph in OpenStreetMap)**

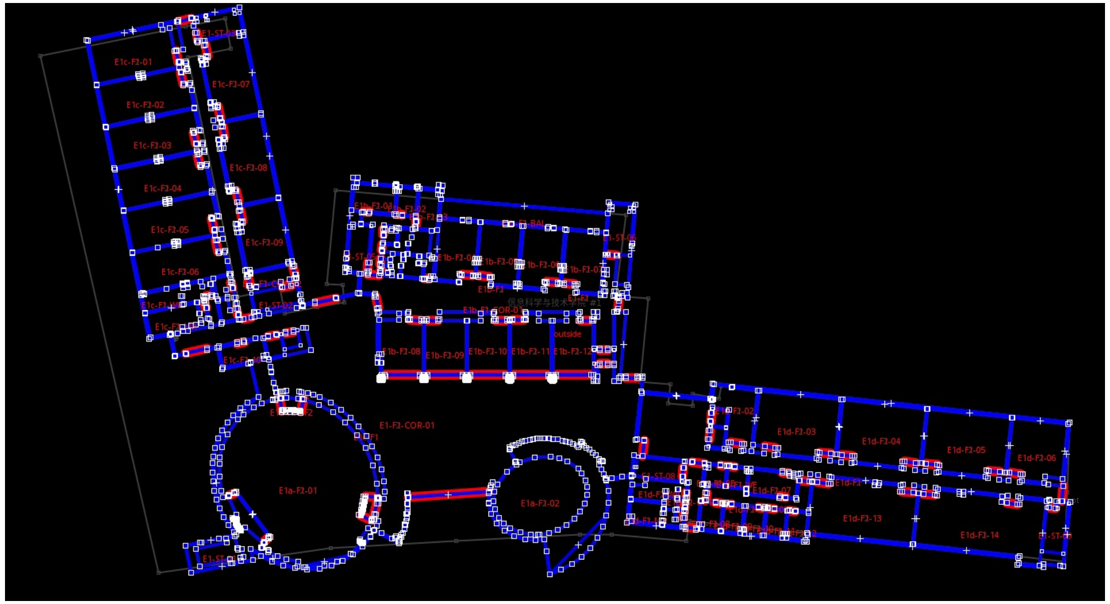


```xml
<way id='-189680' action='modify' visible='true'>
  <nd ref='-344769' />
  <nd ref='-345177' />
  <nd ref='-344900' />
  <nd ref='-344856' />
  <nd ref='-345171' />
  <nd ref='-345214' />
  <nd ref='-344812' />
  <nd ref='-344768' />
  <nd ref='-344769' />
  <tag k='height' v='3.2' />
  <tag k='indoor' v='room' />
  <tag k='level' v='2' />
  <tag k='name' v='E1d-F2-08' />
  <tag k='osmAG:areaType' v='room' />
  <tag k='osmAG:occupied_by' v='Soeren Schwertfeger' />
  <tag k='osmAG:parent' v='E1d-F2' />
  <tag k='osmAG:room_number' v='1D-203.A' />
  <tag k='osmAG:type' v='area' />
</way>
```

```xml
<way id='-189722' action='modify' visible='true'>
  <nd ref='-345171' />
  <nd ref='-345214' />
  <tag k='level' v='2' />
  <tag k='name' v='E1d-F2-COR-03_to_E1d-F2-08' />
  <tag k='osmAG:degree' v='0' />
  <tag k='osmAG:from' v='E1d-F2-COR-03' />
  <tag k='osmAG:passage_type' v='handle' />
  <tag k='osmAG:to' v='E1d-F2-08' />
  <tag k='osmAG:type' v='passage' />
</way>
```


- **Comprehensible by LLMs, traditional robotic algorithms and humans**

osmAG – Map Representation

- **Localization**: Xie F, Schwertfeger S. Robust lifelong indoor lidar localization using the area graph[J]. IEEE Robotics and Automation Letters, 2023, 9(1): 531-538.

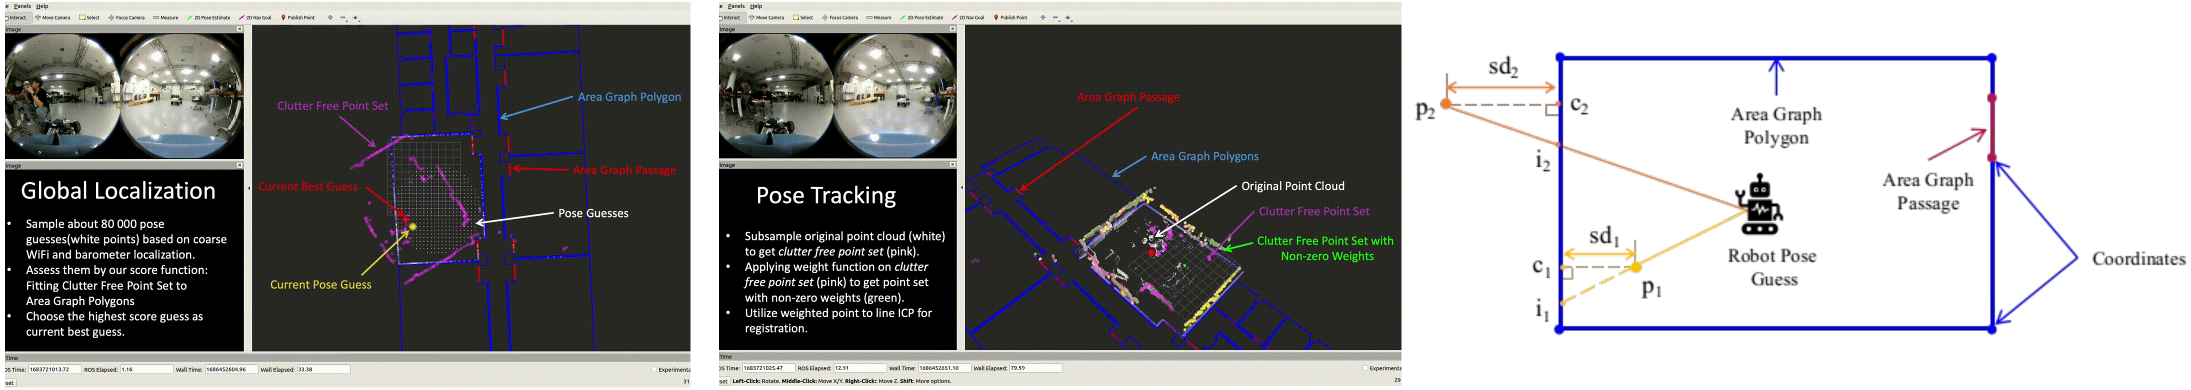

Fig. 1

- **Path Planning:** Xie F, Zhang J, Schwertfeger S. Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot[J]. arXiv preprint arXiv:2409.08493, 2024.

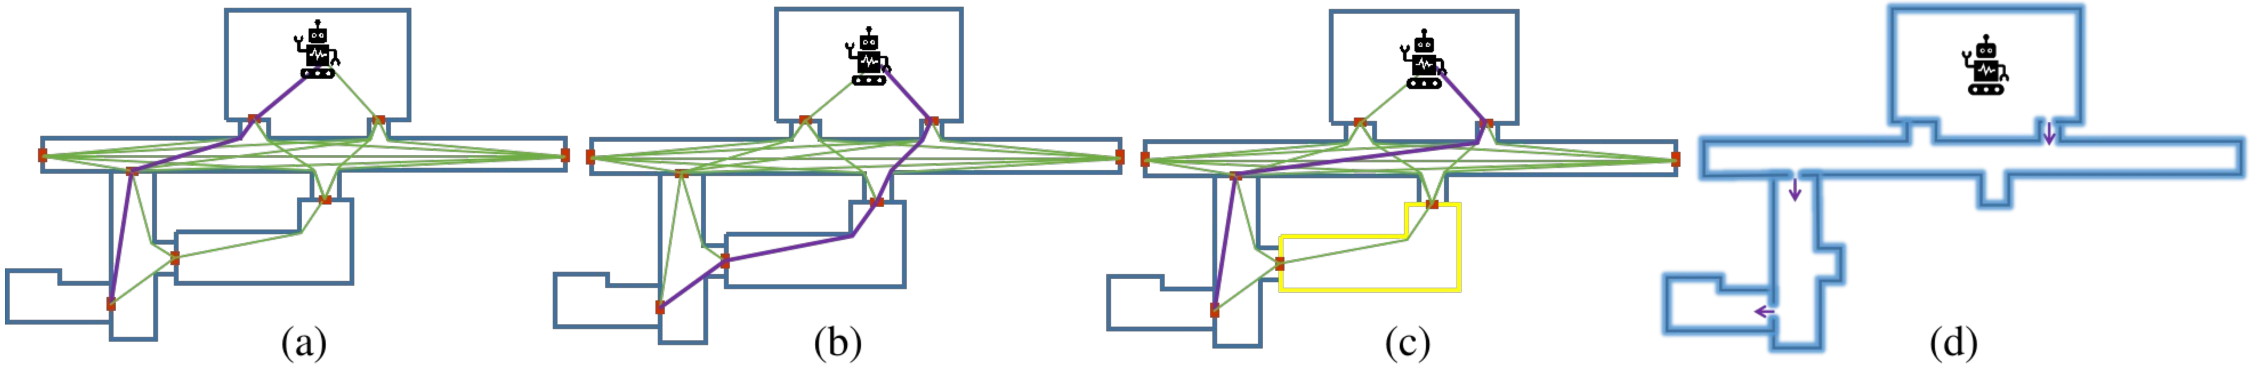

Fig. 2

Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot

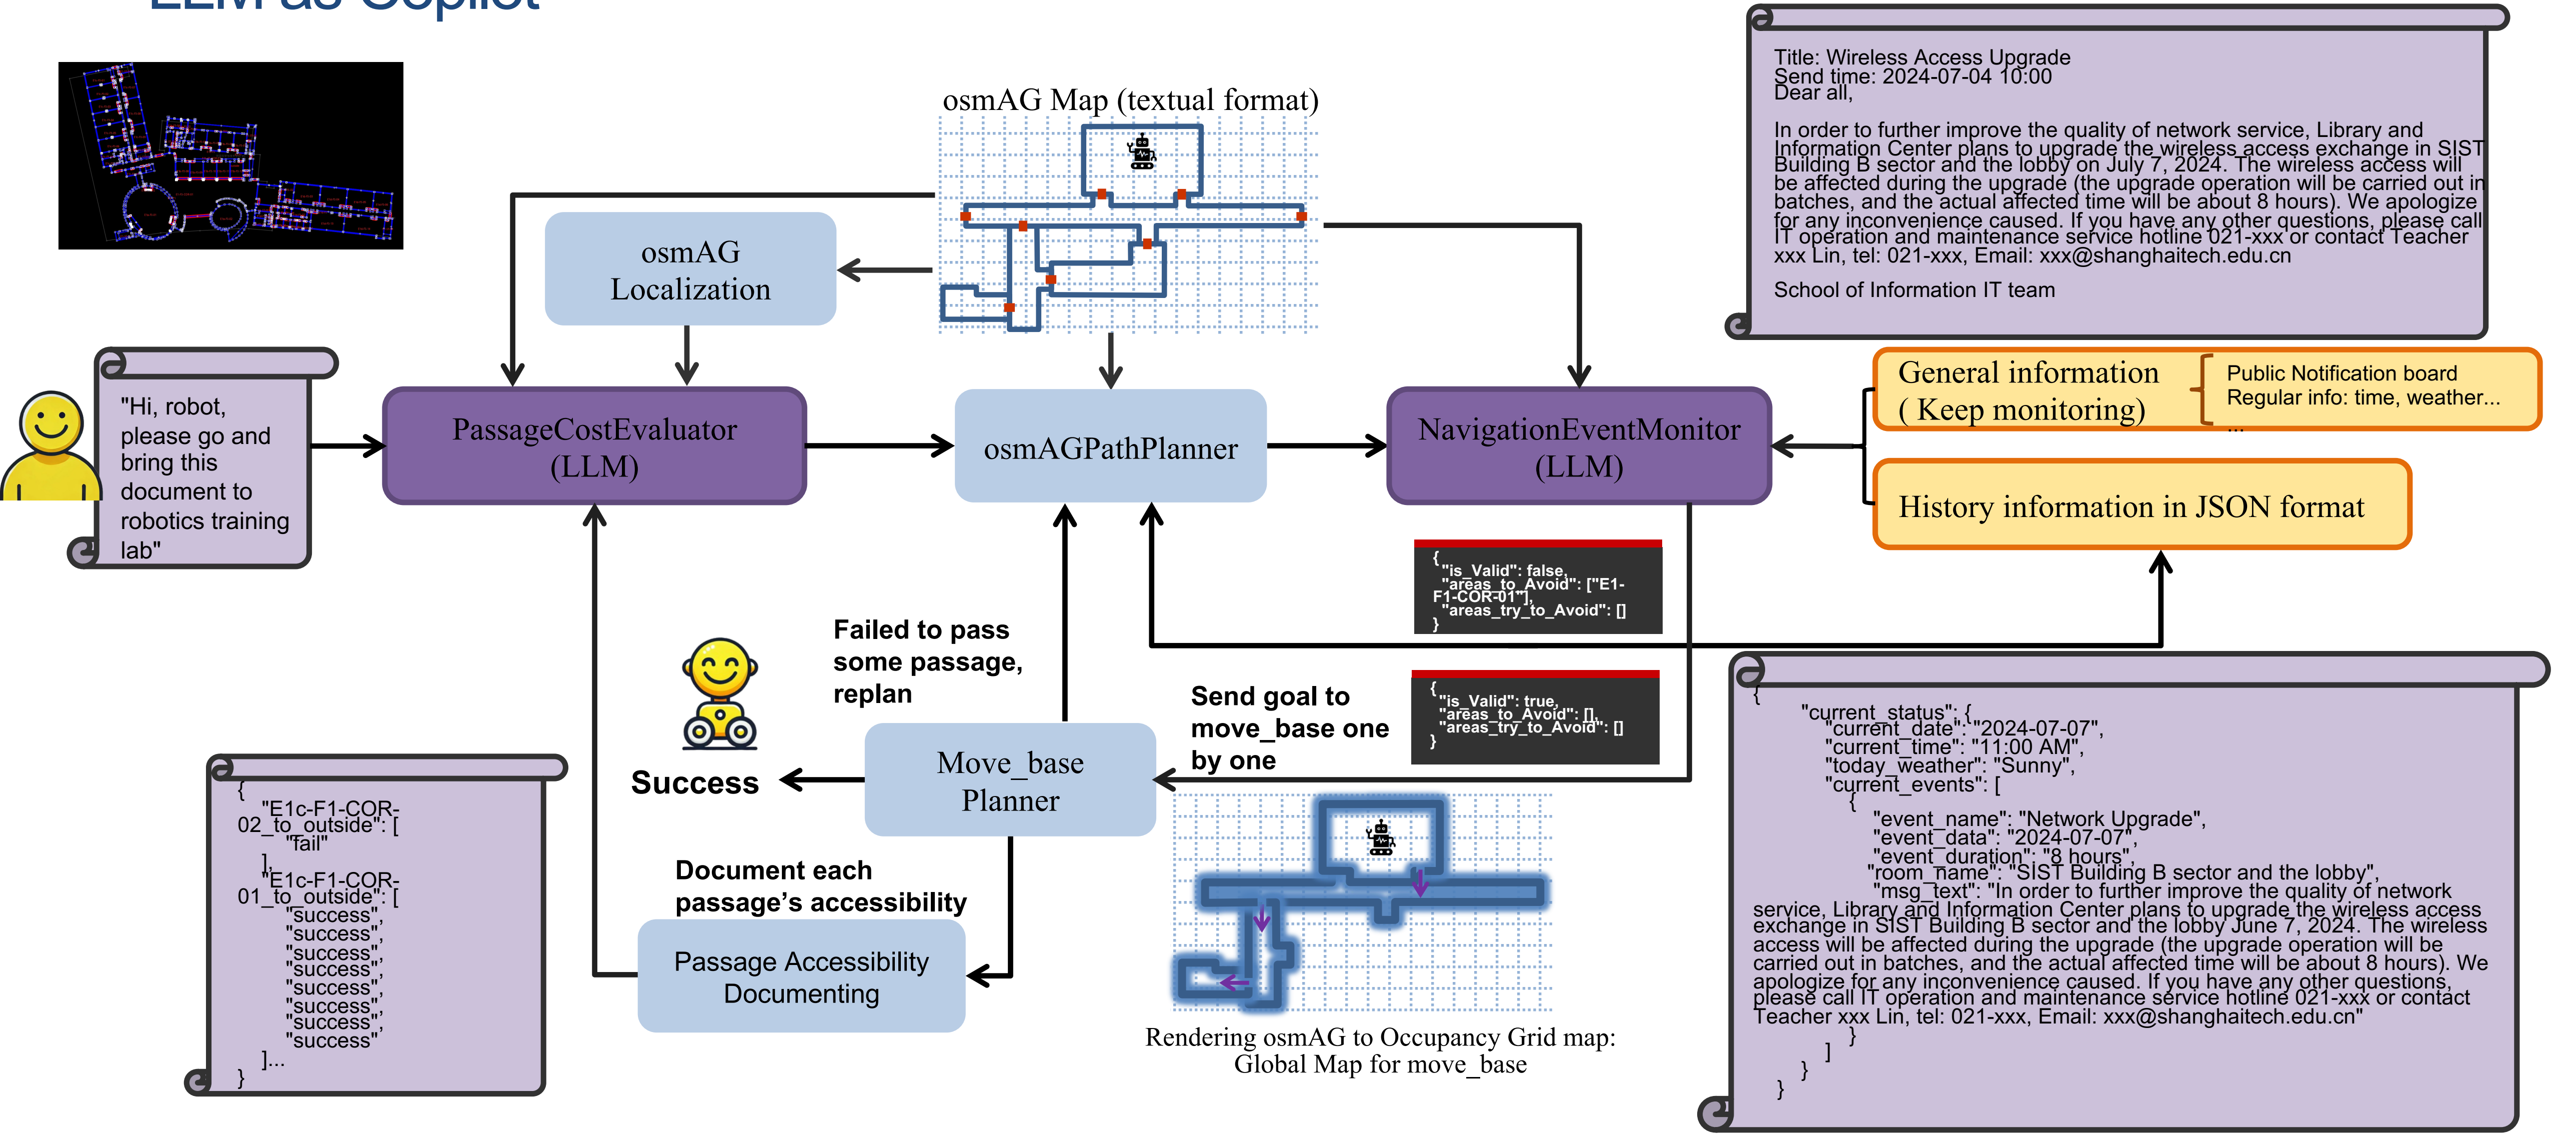

Xie F, Zhang J, Schwertfeger S. Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot[J]. arXiv preprint arXiv:2409.08493, 2024.

Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot

- Combine the strength of traditional planning with AI:
  - Use A* for planning:
    - If input data (map, start and goal pose) are correct:
    - Guaranteed optimal solution
    - Inherently safe (does not suffer from prompt injection)
    - Fast
  - With the help of LLM for:
    - LLM can understand the human input (human robot interaction)
    - LLM can parse external information (texts written for humans, e.g. maintenance announcements)
    - LLM can reason about things, predict stuff, is an AI

Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot

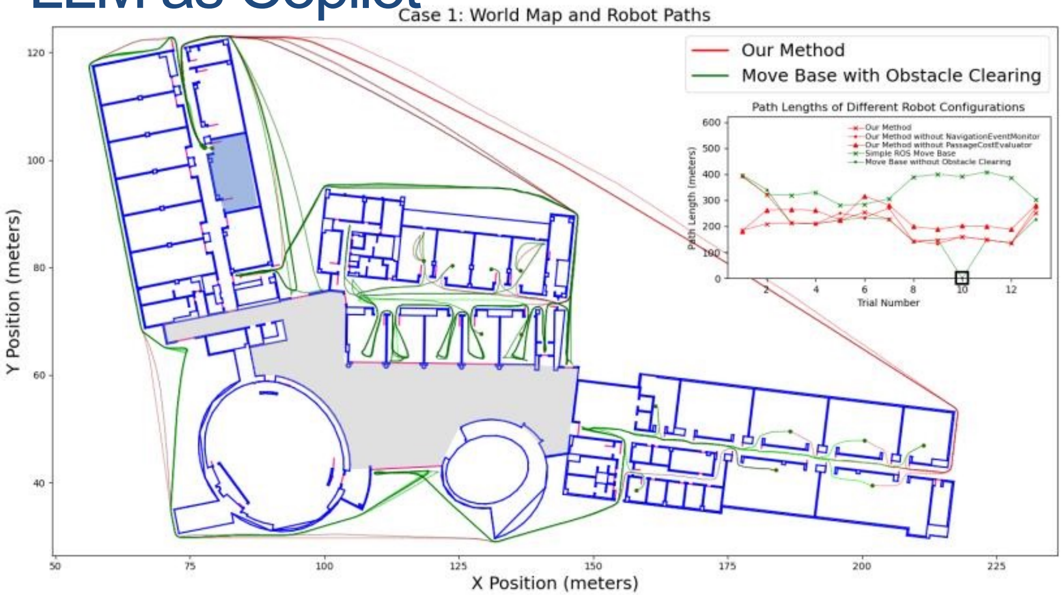

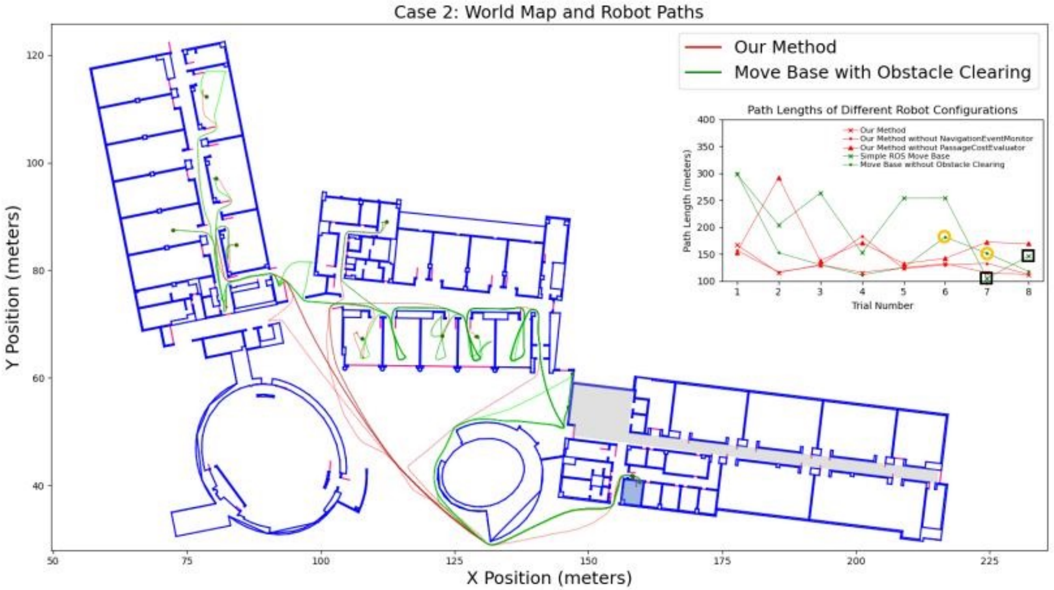

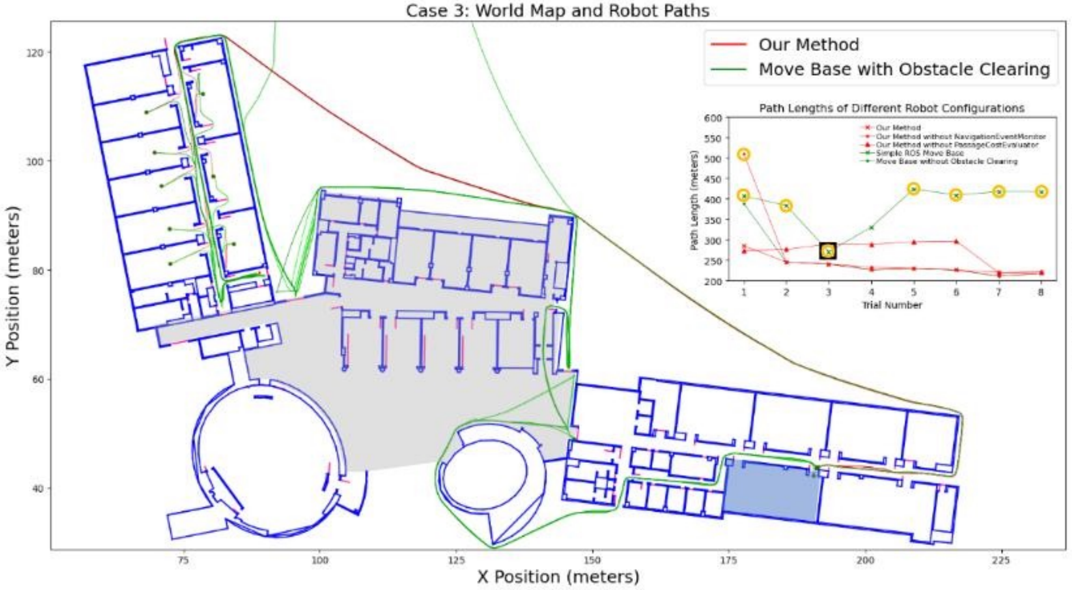

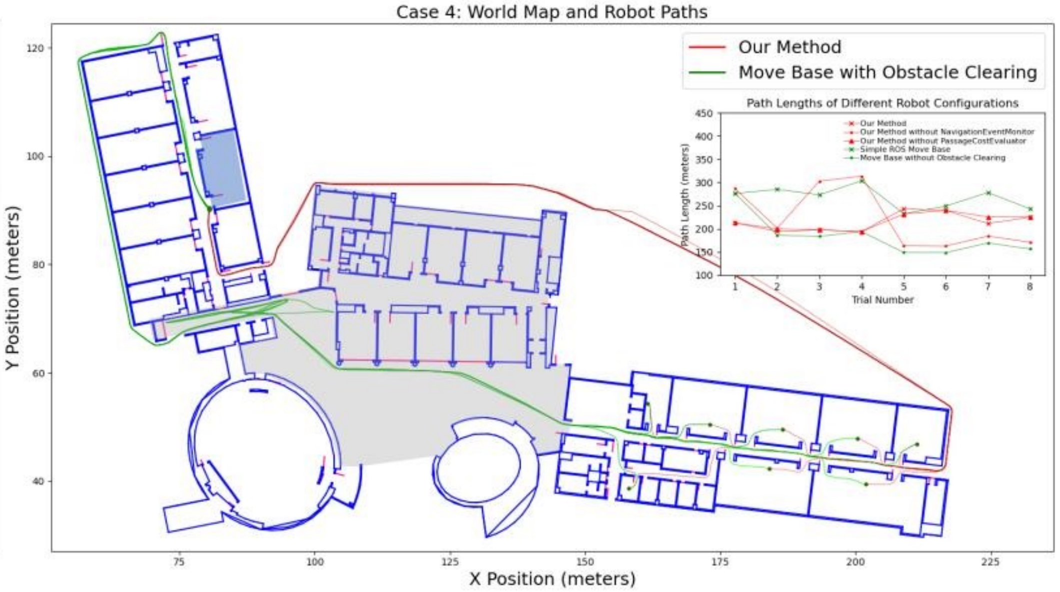

Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot

Fujing Xie, Jiajie Zhang, Sören Schwertfeger

ShanghaiTech University

Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot

TABLE I: Performance Comparison of ChatGPT-4o and DeepSeek-V3 in Path Validation (%)

Model	Accuracy	Recall	False Positive Rate
ChatGPT-4o	0.96	1.0	0.93
DeepSeek-V3	0.93	0.9	0.95

TABLE II: Comparison of Average Path Lengths (m) Across Different Navigation Configurations for Various Cases

Navigation configurations	Case 1	Case 2	Case 3	Case 4
Our method	**192.5**	**126.5**	**236.3**	214.7
Our method w/o *NavigationEventMonitor*	221.2	135.4	264.3	223.3
Our method w/o *PassageCostEvaluator*	234.6	171.4	270.0	215.9
Move_base	347.3	202.2	398.5	267.5
Move_base w/o obstacle layer clearing	220.1	158.7	248.1	**183.1**

Furthermore -- Object Goal Navigation

- **Object-Goal Navigation**

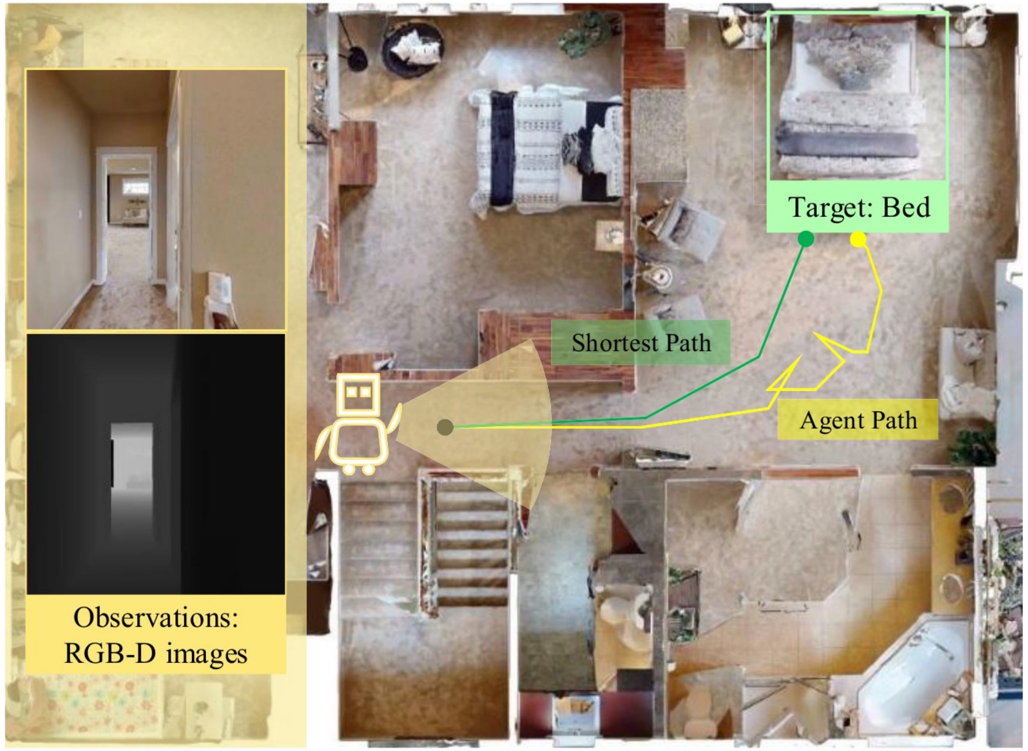

[1] Sun J, Wu J, Ji Z, et al. A survey of object goal navigation[J]. IEEE Transactions on Automation Science and Engineering, 2024.

(a)

- **Example 1**

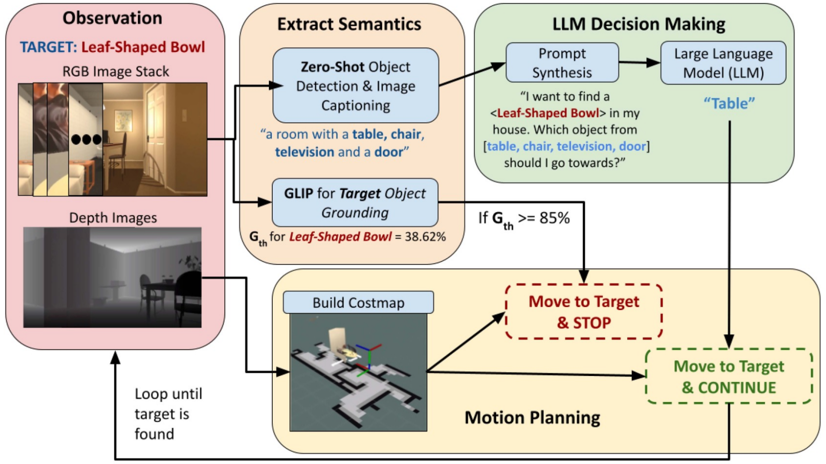

[2] Dorbala V S, Mullen J F, Manocha D. Can an embodied agent find your "cat-shaped mug"? llm-based zero-shot object navigation[J]. IEEE Robotics and Automation Letters, 2023, 9(5): 4083-4090.

(b)

- **Example 2**

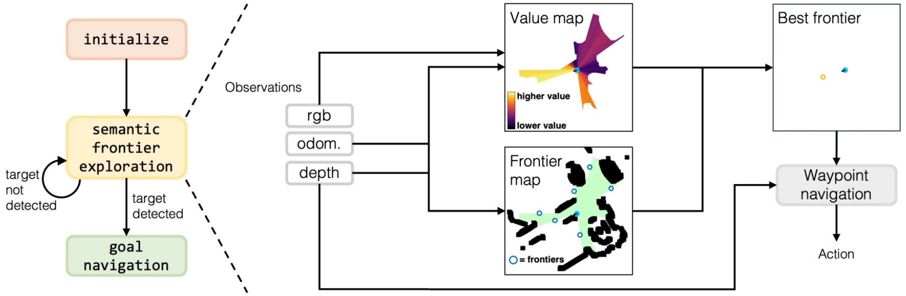

[3] Yokoyama N, Ha S, Batra D, et al. Vlfm: Vision-language frontier maps for zero-shot semantic navigation[C]//2024 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2024: 42-48.

(c)

Furthermore -- Object Goal Navigation with osmAG

- **Mapping Phase:**
  - Adding room description, and nodes represent objects.
  - Off-line, using handhold apple scanner, using LabelMaker[1] and VLMs to label objects.

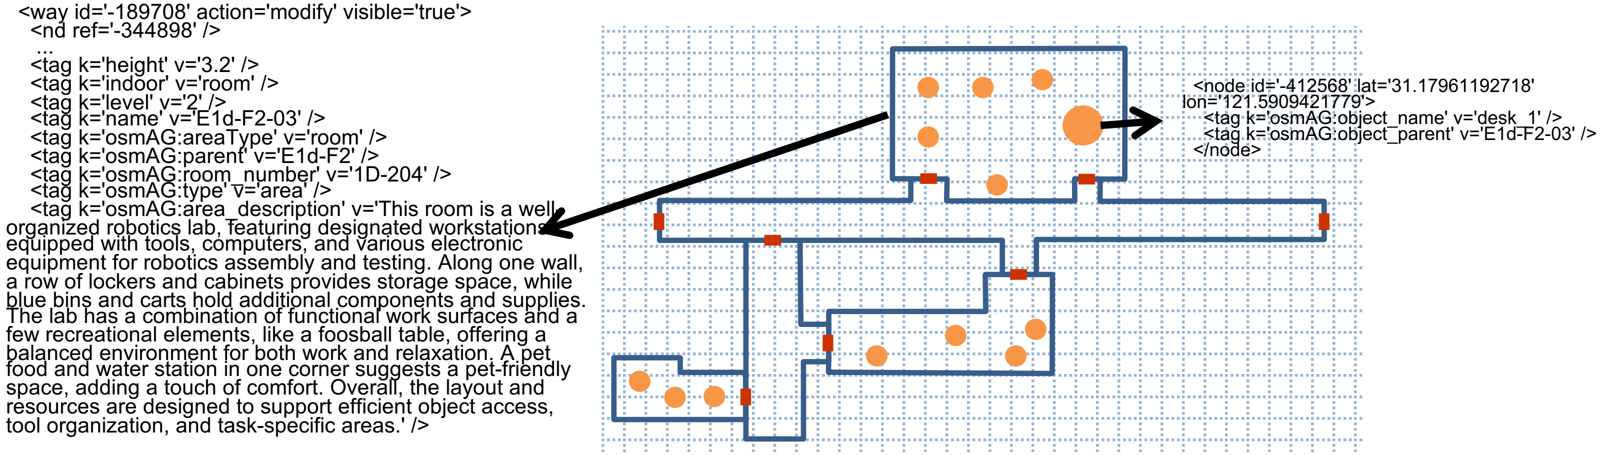

- **Navigation Phase:**
  - Online, using a robot equipped with camera and Jetson Orin

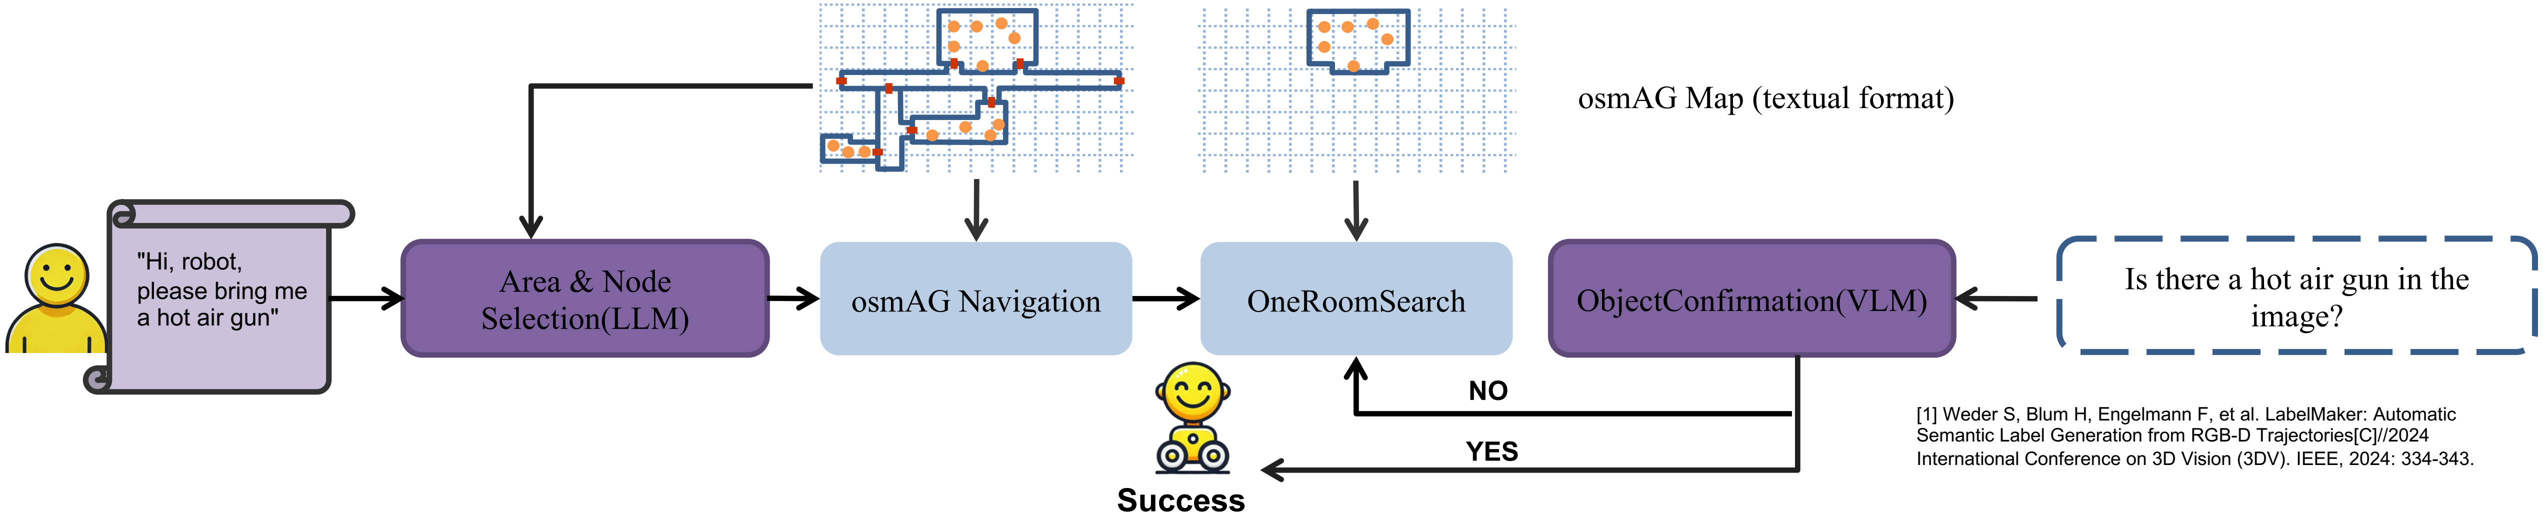

[1] Weder S, Blum H, Engelmann F, et al. LabelMaker: Automatic Semantic Label Generation from RGB-D Trajectories[C]//2024 International Conference on 3D Vision (3DV). IEEE, 2024: 334-343.

Advantages of osmAG

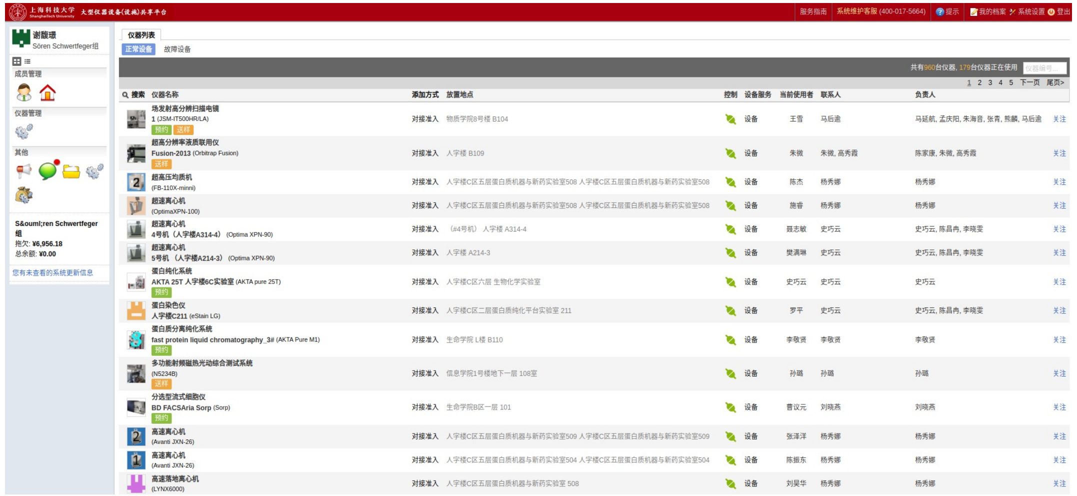

(a)

```xml
<way id='-189680' action='modify' visible='true'>
  <nd ref='-344769' />
  <nd ref='-345177' />
  <nd ref='-344900' />
  <nd ref='-344856' />
  <nd ref='-345171' />
  <nd ref='-345214' />
  <nd ref='-344812' />
  <nd ref='-344768' />
  <nd ref='-344769' />
  <tag k='height' v='3.2' />
  <tag k='indoor' v='room' />
  <tag k='level' v='2' />
  <tag k='name' v='E1d-F2-08' />
  <tag k='osmAG:areaType' v='room' />
  <tag k='osmAG:occupied_by' v='Soeren Schwertfeger' />
  <tag k='osmAG:parent' v='E1d-F2' />
  <tag k='osmAG:room_number' v='1D-203.A' />
  <tag k='osmAG:type' v='area' />
</way>
```

(b)

Mapping

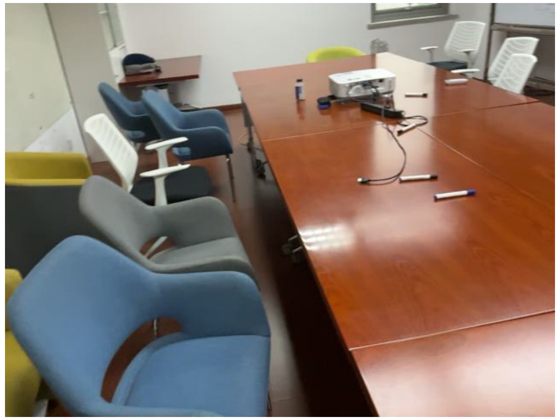

Testing

(c)

```xml
<tag k='height' v='3.2' />
<tag k='indoor' v='room' />
<tag k='level' v='2' />
<tag k='name' v='E1d-F2-13' />
<tag k='osmAG:areaType' v='room' />
<tag k='osmAG:occupied_by' v='Mars Lab' />
<tag k='osmAG:parent' v='E1d-F2' />
<tag k='osmAG:room_number' v='1D-203-1' />
<tag k='osmAG:type' v='area' />
<tag k='osmAG:usable area' v='120 square meters' />
<tag k='semantic_osmAG:area_description' v='&quot;Based on the substantial image descriptions provided, the room appears to be a
multifunctional workspace, likely an electronics or robotics laboratory or workshop. Here&apos;s a composite description of the
room:\n\nThe room is organized and functional, designed to accommodate various technical activities. It features multiple workstations
equipped with long desks aligning the walls. These desks are cluttered with electronic devices such as computer monitors, desktop towers,
laptops, and components including oscilloscopes, soldering irons, and power supplies.\n\nThe workspace shows evidence of active technical
work with numerous tools, cables, and small components scattered across tables. Pegboards mounted on the walls display tools like
screwdrivers and pliers. Numerous storage solutions are present, including white cabinets labeled with organizational tags and colorful
plastic bins organizing electronic parts and tools.\n\nThere are also elements of a communal or collaborative environment visible, such as
filing cabinets, office chairs, and desks with typical office supplies and computer setups. The workspace appears to be well-lit, with
natural light coming in through windows linking the room to the outside environment. Some windowsills are adorned with indoor plants and
decor items, suggesting attempts to personalize the space or enhance the ambiance.\n\nDistinctive elements within the room include blue
trash bins with black liners, which appear to be scattered in strategic positions for easy access. Safety equipment like first aid kits,
fire extinguishers, and noticeable signage (including control panels or switches) highlight considerations for regulatory compliance and
user safety.\n\nA significant portion of the room seems dedicated to robotics with several mobile robotic platforms and robotic arms
present. These pieces of equipment indicate focuses on development or testing in robotics and automation. Some setups show detailed work
involving mechanical components or robotics, suggesting ongoing experimental or research activities.\n\nIn leisure or social areas, a
foosball table is visible, juxtaposing the technical environment with recreational space, hinting at a workspace culture that values breaks
and team interaction.\n\nFlooring in the room is primarily light-colored, featuring smooth surfaces like concrete or tiles which ensure
ease of movement, especially for wheeled items like mobile robots and chairs. Various tripods and technical equipment mounted on mobile
stands indicate the dynamic nature of experimentation or testing processes.\n\nThrough these accumulated elements, the overall environment
suggests an industrious, multipurpose technical space, capable of supporting a wide range of engineering activities, particularly in
electronics, robotics, or similar fields. The space is well-organized yet actively used, displaying a balance between functionality,
safety, and possible user comfort.&quot;' />
```

(d)

Furthermore -- Object Goal Navigation with osmAG

- **Mapping Phase:**

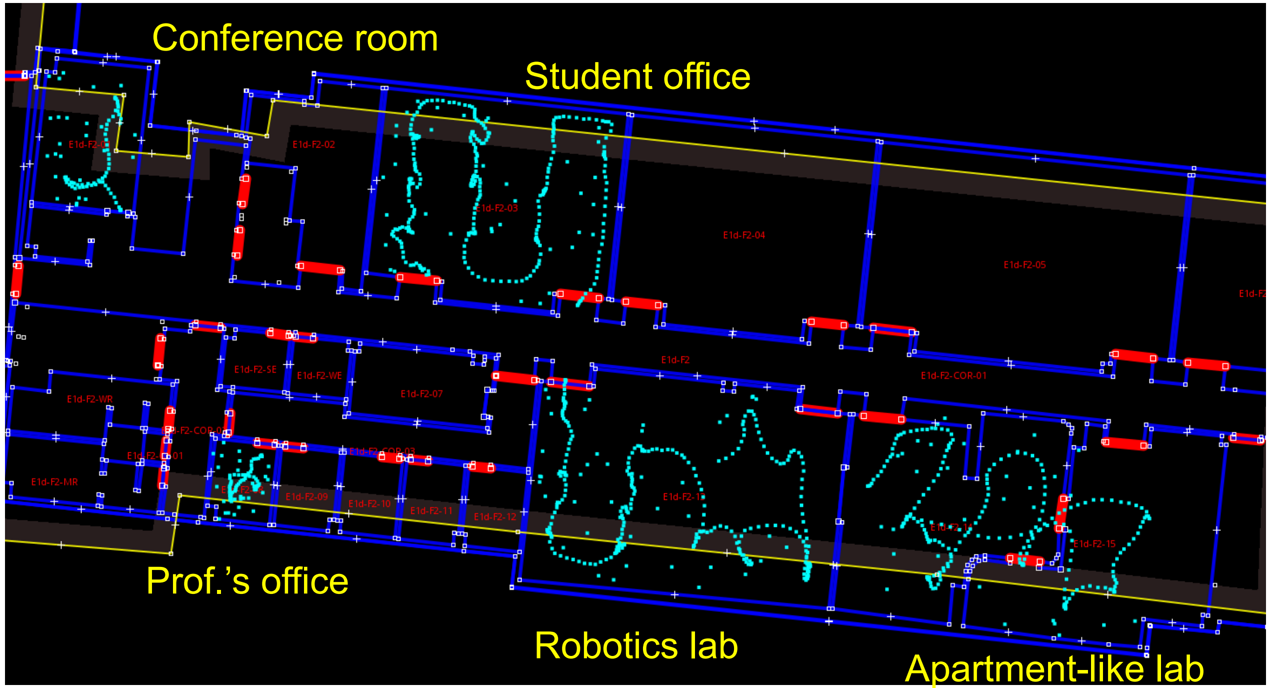

Fig. 1 Mapping result of 5 rooms in SIST building1

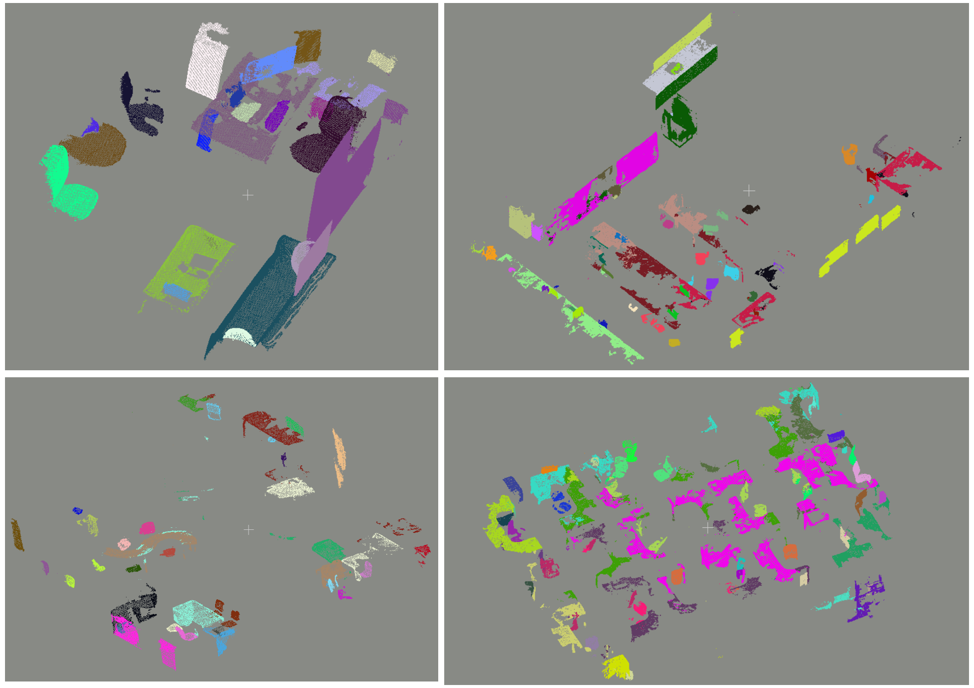

Fig. 2 Instance segmentation result from LabelMaker

Furthermore -- Object Goal Navigation with osmAG

- **Navigation Phase:**

(a) (b) (c)

- Static Objects:
  - Trash can
  - Sink
- Relocated Objects:
  - dry erase marker
  - robot dog
- Objects not exist during mapping:
  - Onion
  - Soeren's Excellent Faculty Award

[c] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, “Hierarchical Open-Vocabulary 3D Scene Graphs for Language-Grounded Robot Navigation,” Robotics: Science and Systems, 2024.

File Panels Help Robotics ShanghaiTech University - SIST - May 13, 2025

Interact | Move Camera | Select | Focus Camera | Measure | 2D Pose Estimate | 2D Nav Goal | Publish Point

Image_raw

/dinox_detection

No Image

object_bbox

No Image

```
Human instruction: I need a dry erase marker for the
whiteboard. Can you find one?
LLM response: {'E1d-F2-01': {'-8765487': (121.59071601107,
31.17962332427), '-8765488': (121.59068371109,
31.17966054779), '-8765489': (121.59070128246,
31.17959962954)}, 'E1d-F2-08': {'-8765575': (121.5907953599,
31.17946797566), '-8765576': (121.5907859773, 31.17947713376),
'-8765577': (121.59078196734, 31.17945507626)}, 'E1d-F2-14':
{'-8767262': (121.59114740408, 31.17947579278), '-8767263':
(121.59111653048, 31.17947414198), '-8767264':
(121.59110945308, 31.17948100105)}}
```

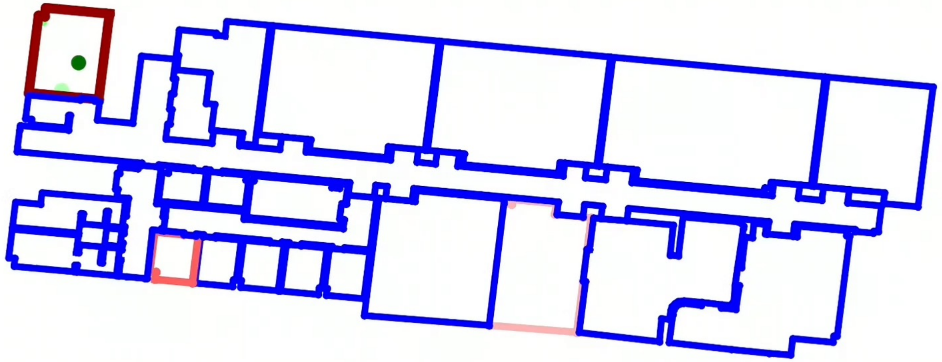

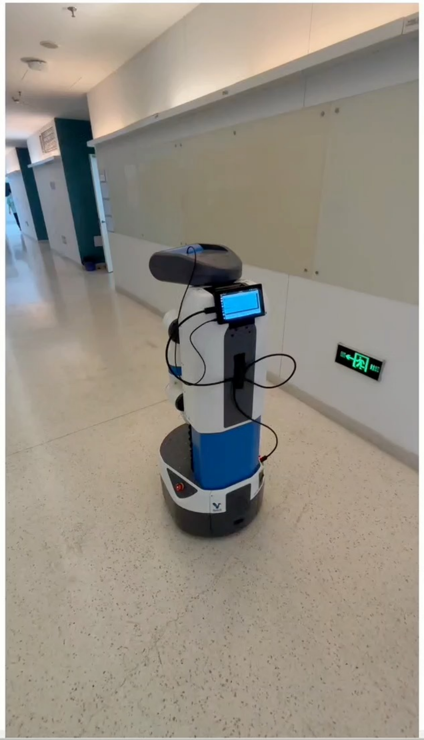

Time

Pause | Synchronization: Off | ROS Time: 1745042271.54 | ROS Elapsed: 13.09 | Wall Time: 1745042271.56 | Wall Elapsed: 13.09

Reset

31 fps

AI ETHICS

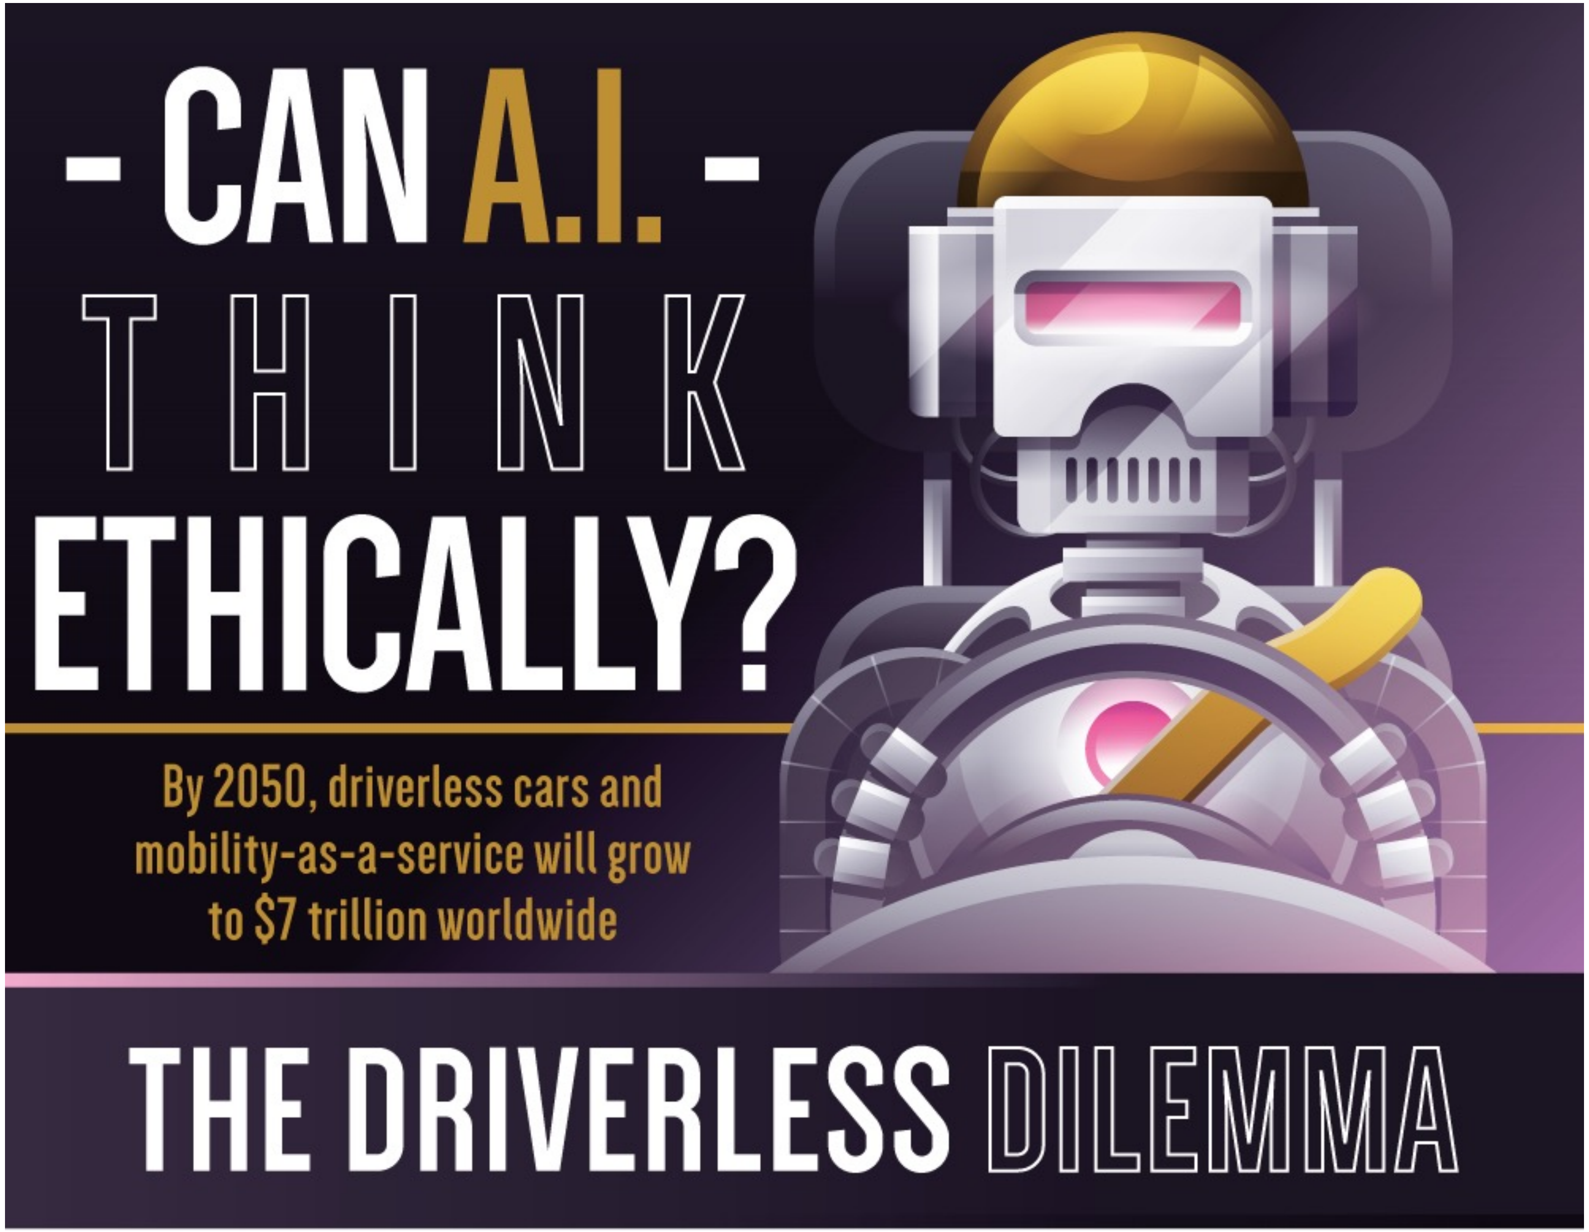

https://cybersecuritydegrees.com/ethical-ai/

FROM 2035 TO 2045

Consumers will regain up to **250 MILLION HOURS OF FREE TIME** from behind the wheel

$234 BILLION IN PUBLIC COSTS will be saved by reducing accidents from human error

Driverless cars can **ELIMINATE 90% OF TRAFFIC FATALITIES** — Saving 1 million lives every year

HOW WILL DRIVERLESS CARS DETERMINE WHOSE LIFE SHOULD BE SPARED?

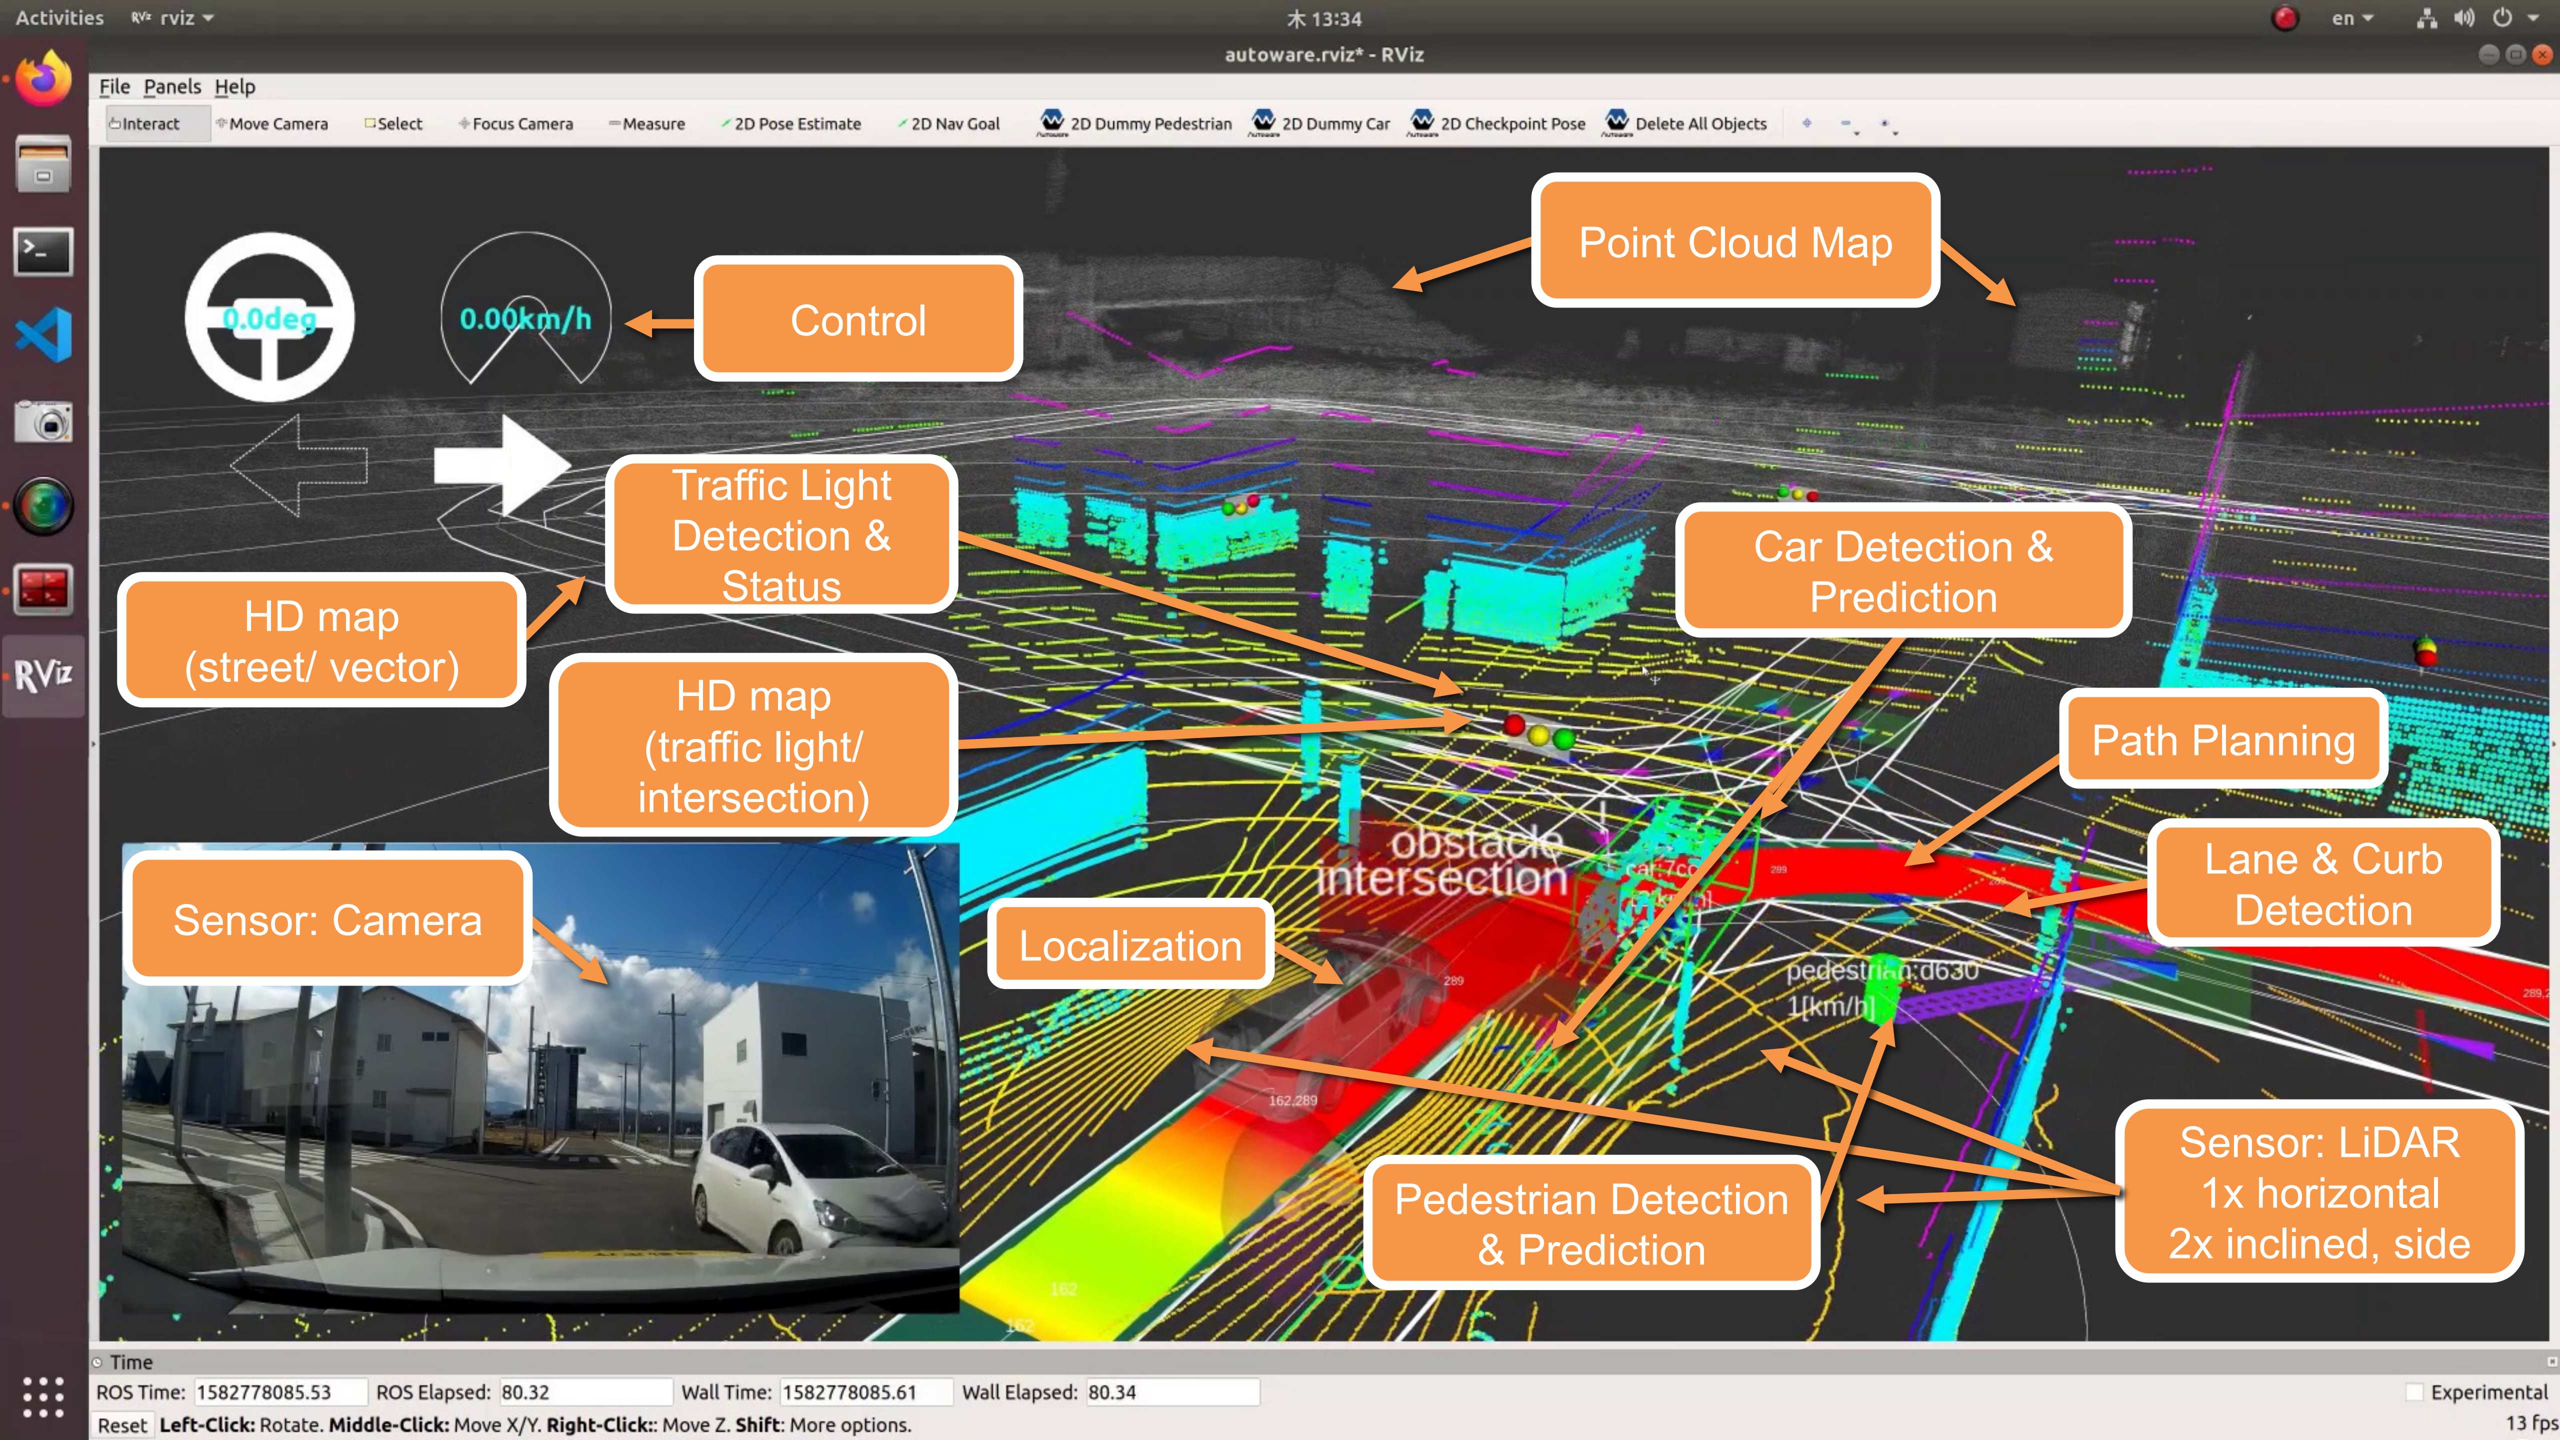

Activities
rviz
13:34
autoware.rviz* - RViz
File Panels Help
Interact
Move Camera
Select
Focus Camera
Measure
2D Pose Estimate
2D Nav Goal
2D Dummy Pedestrian
2D Dummy Car
2D Checkpoint Pose
Delete All Objects
Point Cloud Map
0.0deg
0.00km/h
Control
Traffic Light Detection & Status
Car Detection & Prediction
HD map (street/ vector)
HD map (traffic light/ intersection)
Path Planning
obstacle intersection
Lane & Curb Detection
Sensor: Camera
Localization
pedestrian:d630
1[km/h]
Sensor: LiDAR 1x horizontal 2x inclined, side
Pedestrian Detection & Prediction
Time
ROS Time: 1582778085.53
ROS Elapsed: 80.32
Wall Time: 1582778085.61
Wall Elapsed: 80.34
Experimental
Reset Left-Click: Rotate. Middle-Click: Move X/Y. Right-Click:: Move Z. Shift: More options.
13 fps

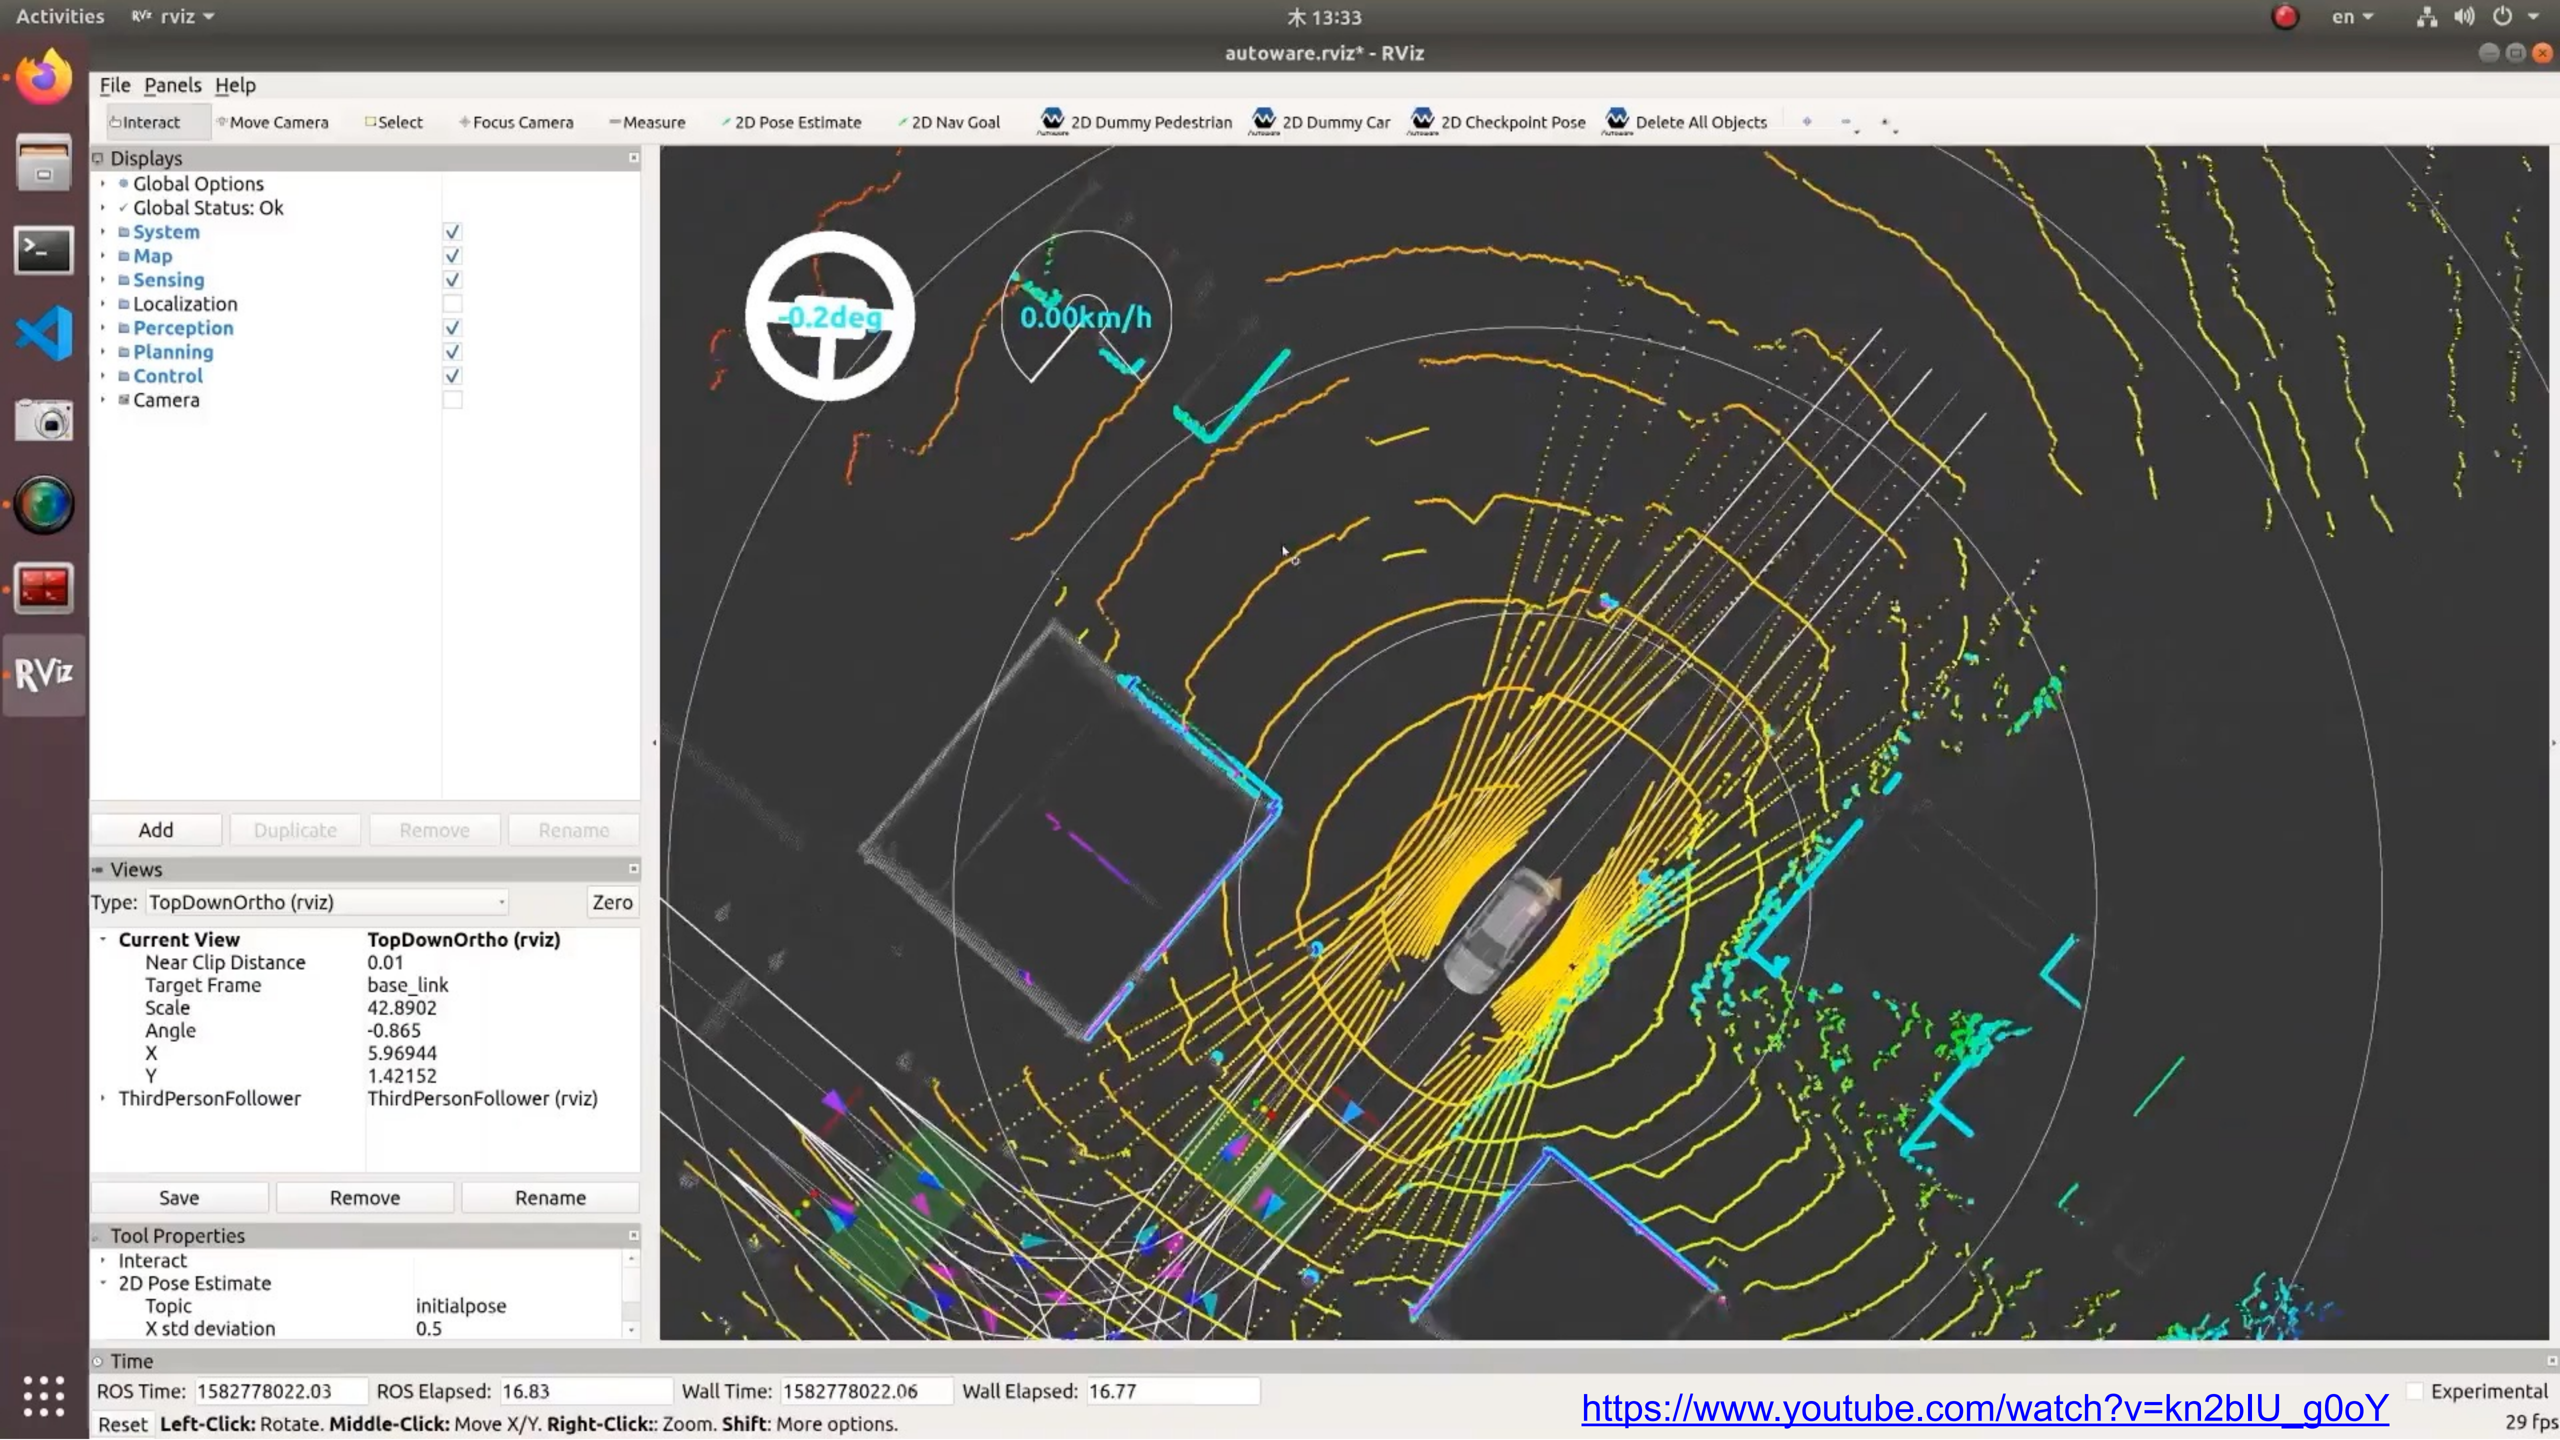

https://www.youtube.com/watch?v=kn2bIU_g0oY

https://www.youtube.com/watch?v=HSeIOg4SyOg

电车托尼 bilibili
高速幽灵刹车
TESLA PHANTOM BRAKING

https://www.youtube.com/watch?v=X3hrKnv0dPQ

https://ed.ted.com/lessons/would-you-sacrifice-one-person-to-save-five-eleanor-nelsen

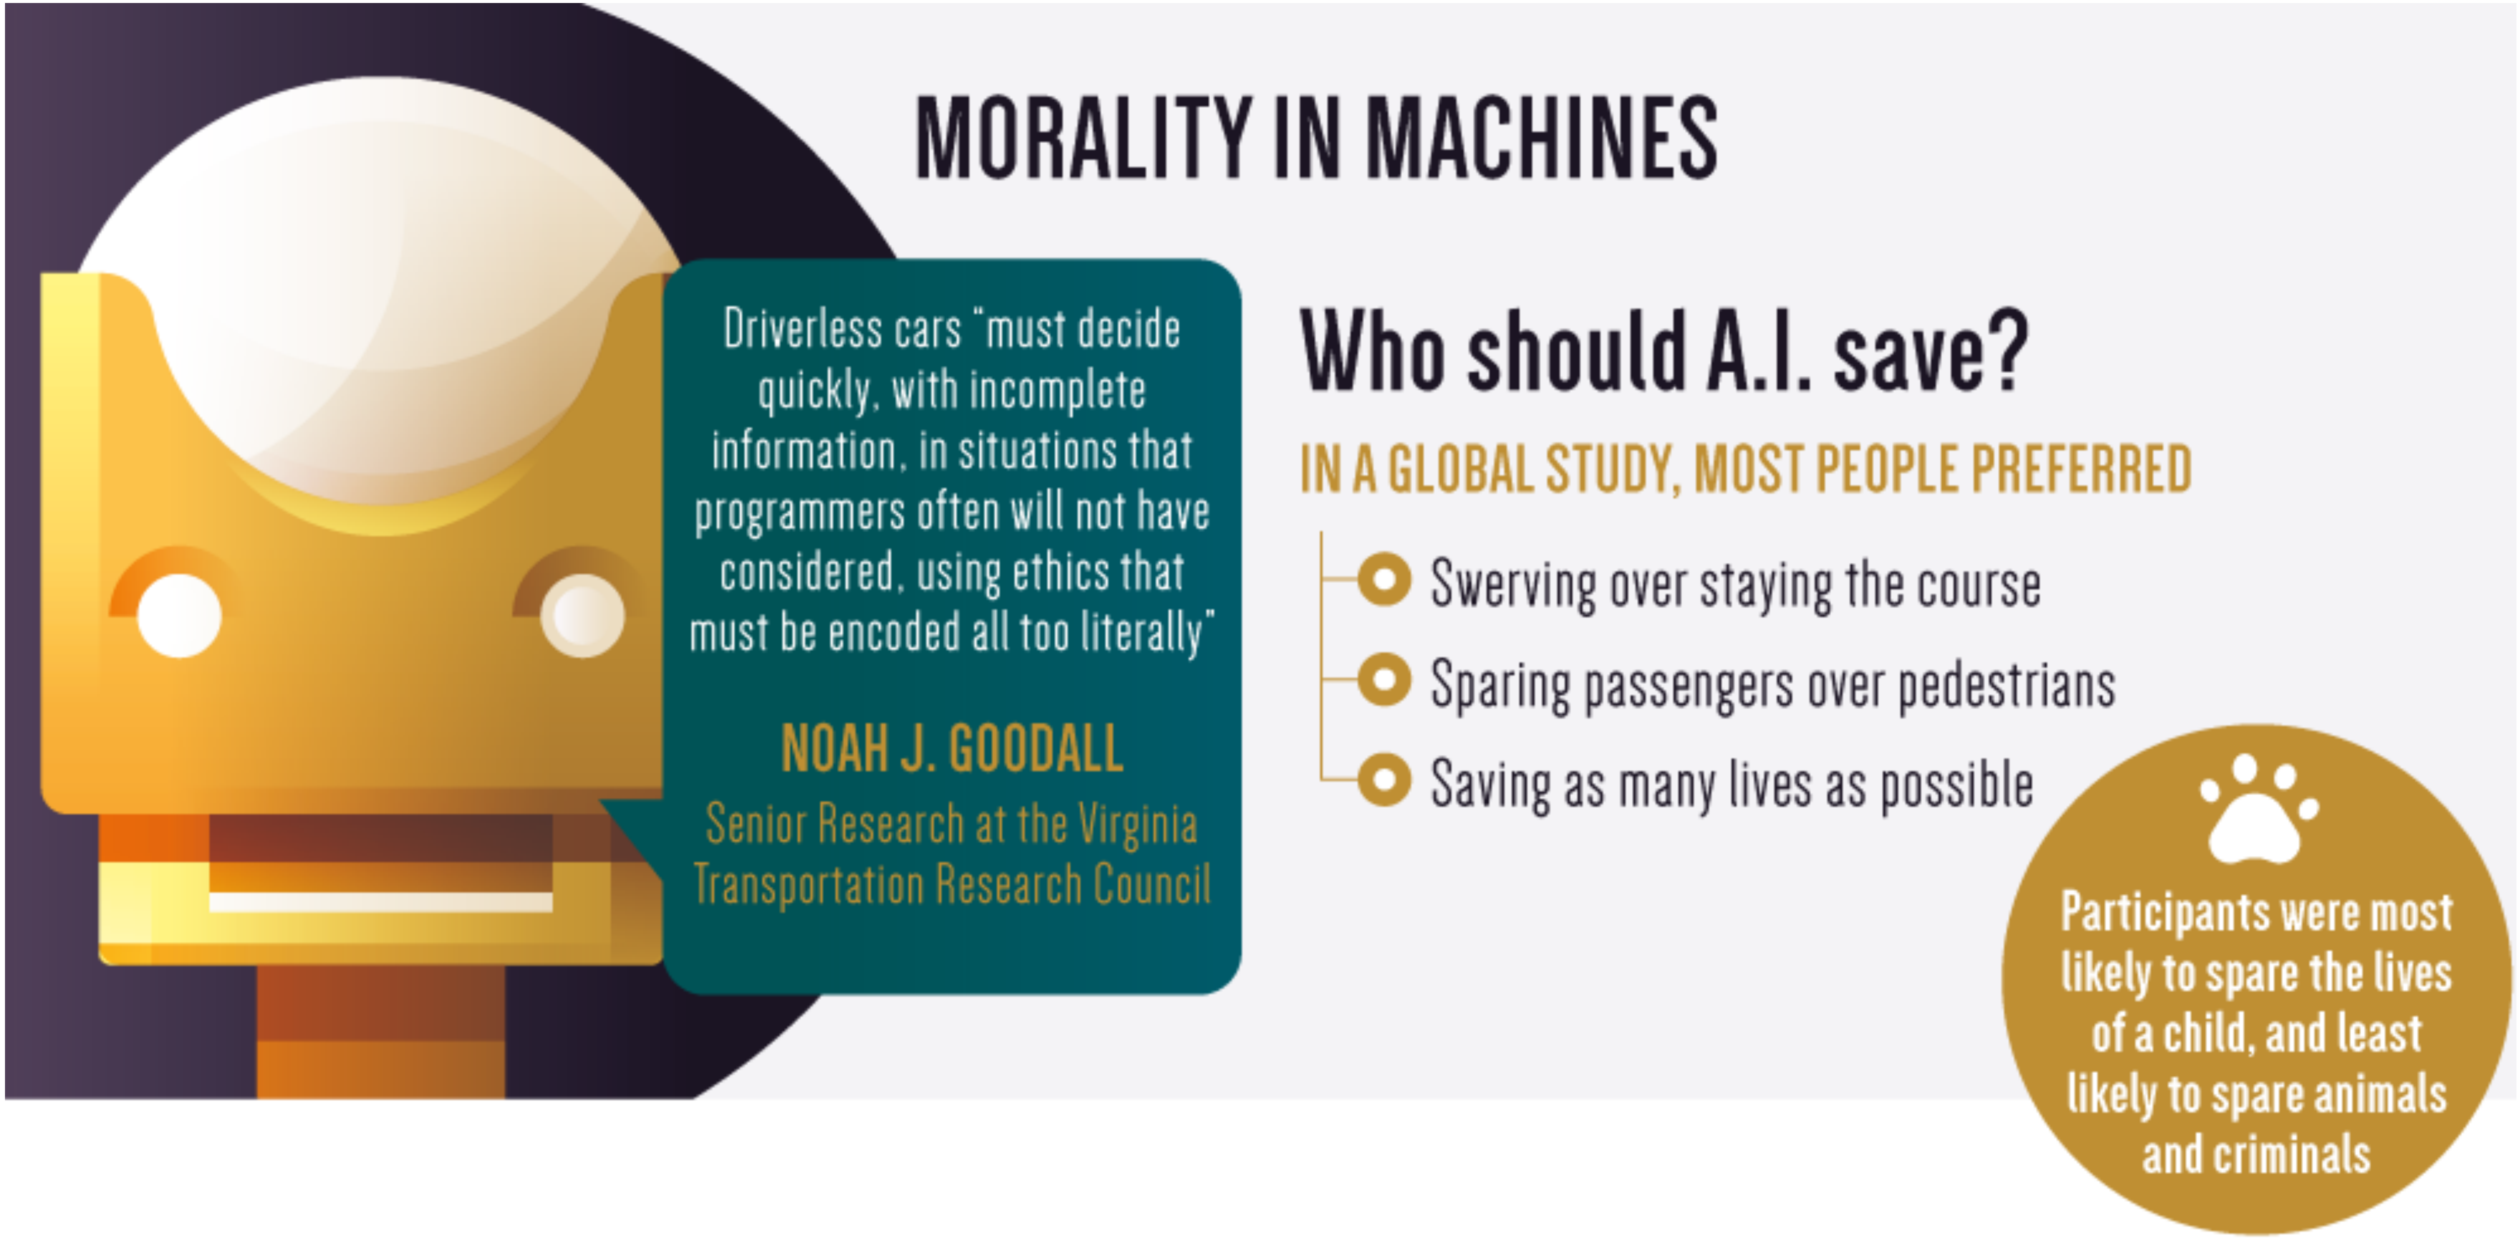
MORALITY IN MACHINES
Driverless cars "must decide quickly, with incomplete information, in situations that programmers often will not have considered, using ethics that must be encoded all too literally"
NOAH J. GOODALL
Senior Research at the Virginia Transportation Research Council
Who should A.I. save?
IN A GLOBAL STUDY, MOST PEOPLE PREFERRED
Swerving over staying the course
Sparing passengers over pedestrians
Saving as many lives as possible
Participants were most likely to spare the lives of a child, and least likely to spare animals and criminals

MINIMIZED HARM VS. PASSENGER PROTECTION

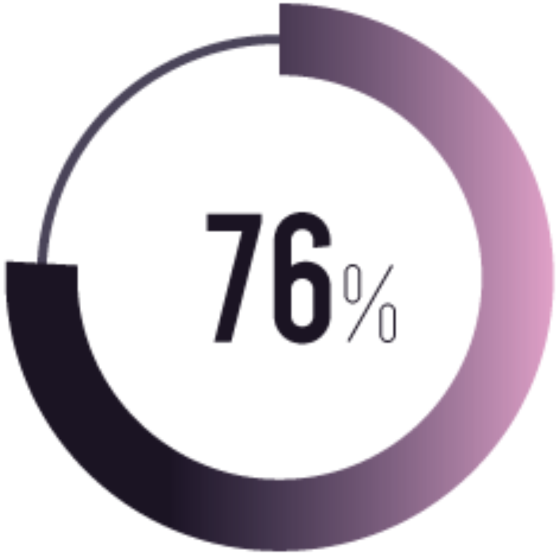

of people felt driverless cars should **SAVE AS MANY LIVES AS POSSIBLE**

<

BUT, very few were willing to buy a vehicle programmed to minimize harm

They prefer cars programmed to **PROTECT PASSENGERS AT ALL COSTS**

DRIVERLESS CARS WILL SAVE LIVES, BUT PROGRAMMING THEM TO DO SO COULD SLOW THEIR ADOPTION AND COST MANY MORE LIVES