



上海科技大学
ShanghaiTech University

CS283: Robotics Spring 2026: Learning & Ethics

Sören Schwertfeger / 师泽仁

ShanghaiTech University

ADMIN

HW

- HW 3 is due this Friday...
- HW 4:
 - Use gazebo simulation to implement a simple frontier exploration algorithm with a SLAM algorithm for mapping (e.g. the one from HW3) and navigation 2 for ... navigation.
- Due Fri, June 19, 23:59

Project conclusion

- Due June 22:
 - Project Report:
 - Like intermediate report
 - With nice results and proper quantitative evaluation
 - Make look like a scientific paper
 - Use bibtex!
 - Put into git (folder: doc/final)
 - Include everything that is needed to generate the report in the git!
 - So don't forget images/ the bib file
 - Project Demo:
 - Make an appointment with Prof. Schwertfeger and your TA to show the final demo of your project
 - Before: June 22, 22:00 !

Project Webpage

- Write a text (word document) about your project for the general public – not too technical – not too many details
 - Some details can be written
 - Do not just copy the abstract/ intro from your final report – write a nice text for the general public!
- Provide a few images with captions (as images also extra files)
- Put into your group git (folder: doc/webpage)
- Prof. Schwertfeger will upload the data to the website – e.g. look at : <https://robotics.shanghaitech.edu.cn/teaching> all previous robotics and MoMa courses
- Also make a NICE video about your project. 4-8 minutes. Leave the video at good quality – size maybe 100 – 300 MB (MP4) – Prof. Schwertfeger will compress it to make a web version
 - Avoid showing other people; do not talk in the video; do not add music;
 - Add a title page: same info as on report
 - Add to your git folder

Final

- June 23
 - 16:00 – 18:00 in 1A-108
- Content:
 - All lectures
 - Take a look at facts, algorithms, concepts
 - Take a look at the homeworks again
 - Sample exam: https://robotics.shanghaitech.edu.cn/sites/default/files//files/final_Example.pdf
- You are allowed to bring **3** A4 sheets (so 6 pages) of info to the exams. You can write/ print anything on those sheets. On top of **every page** (so 6 times) there needs to be your **name (pinyin), student ID and ShanghaiTech email** address. We will check every cheat sheet before the exam and **confiscate** every sheet without name or with a name that is not yours.
- No electronics/ calculator/ smartwatch allowed

Midterm -> Final

- Questions in final can be:
 - Similar to the midterm
 - New questions (e.g. other algorithm)
- Midterm grading will be finished shortly after ICRA...
- => learn for your Midterm mistakes!

Project Meetings!

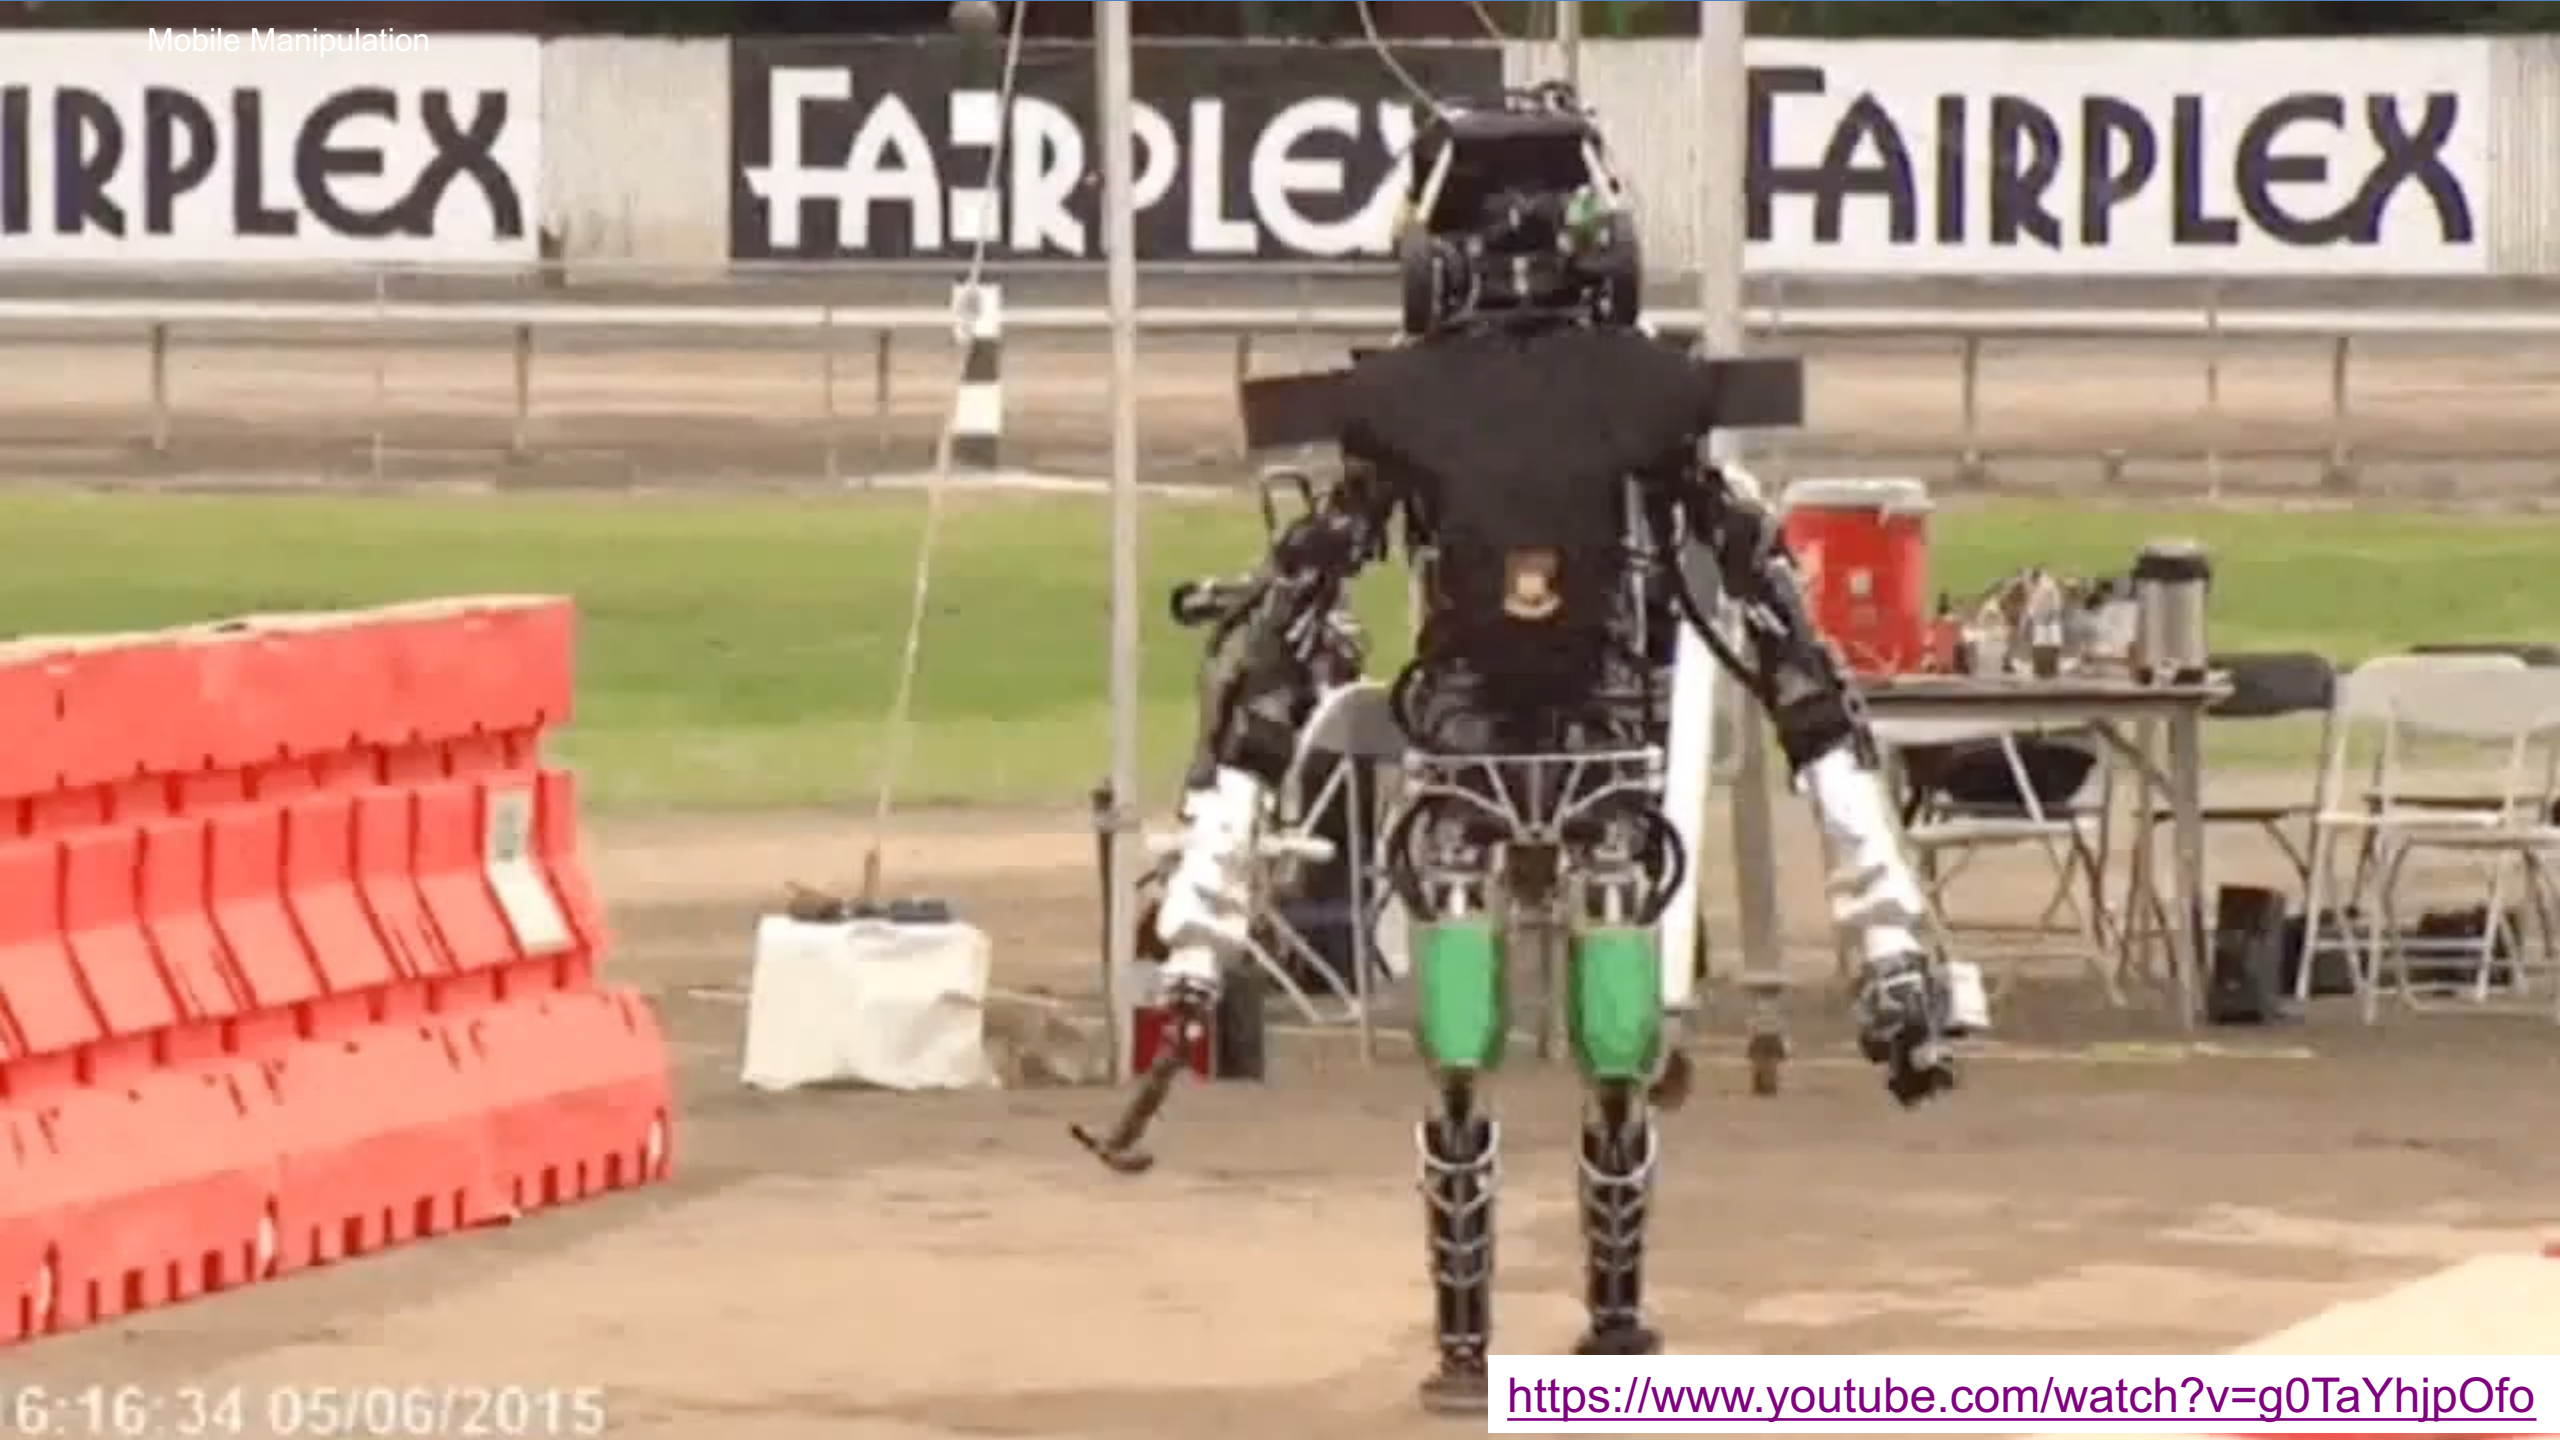
- Today is last class => spend lots of time on your project!
- We have scheduled more “after intermediate report meetings” today – see me in my office
- We will make at least one another project meeting between June 11 and 16 schedule with Bichi!
- Also, other meetings if needed!

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



Modern Robotics with the help of Learning/ AI

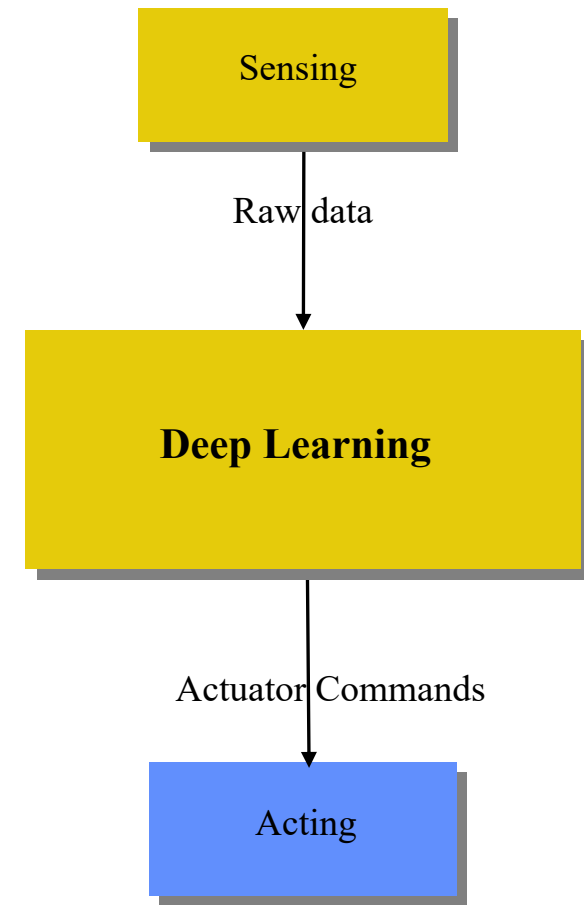
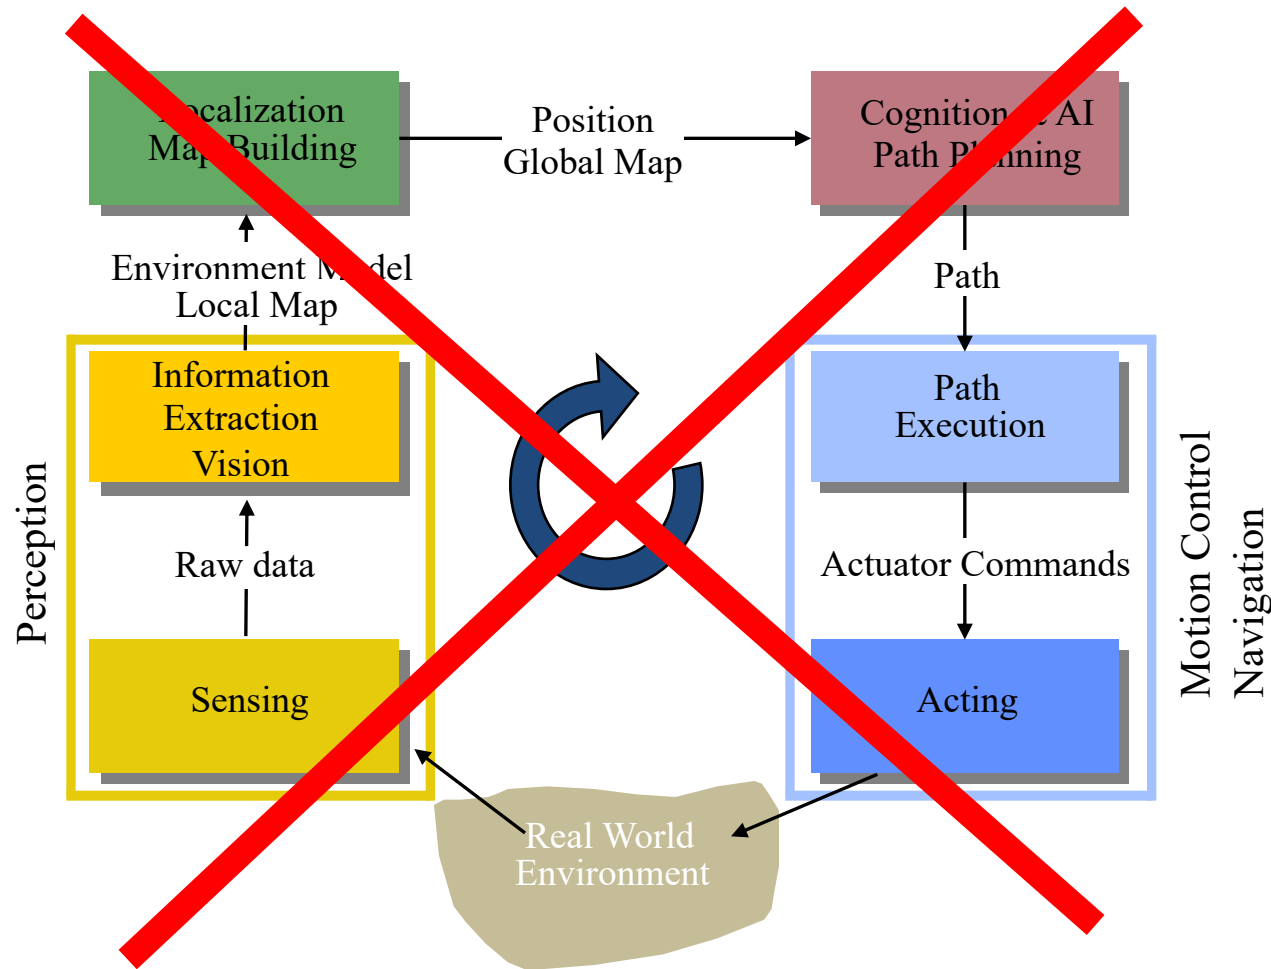
- Learning for Robotics: Train a Neural Network to do the things that are too hard to program by hand, e.g.:
- DL based computer vision and recognition (e.g. also 3D point clouds)
- Reinforcement Learning for control
- LLM for human robot interaction and intelligence
- ...

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

End-to-End Deep Learning



End-to-End Deep Reinforcement Learning

- From sensors to actuation: one layered or recurrent neural network! =>
 - NOT classical general control scheme (Perception, SLAM, Cognition & Planning, Navigation)
- Needs reward signal: sparse, noisy, delayed!
- Take time into account: input frames are related!
- Gained interest 2013 again with:
 - Deep Mind (google) playing ATARI 2600 games
 - Video: Breakout
 - Learned 7 games
 - Surpasses human expert in 3

BRETT: Berkeley Robot for Learning Tedious Tasks: Deep Reinforcement Learning

- "There are no labeled directions, no examples of how to solve the problem in advance. There are no examples of the correct solution like one would have in speech and vision recognition programs"
- Learn simple tasks in 10 minutes;
learn vision and control together in 3 hours
- Pieter Abbeel of UC Berkeley;
- 2015

<http://news.berkeley.edu/2015/05/21/deep-learning-robot-masters-skills-via-trial-and-error/>

BRETT

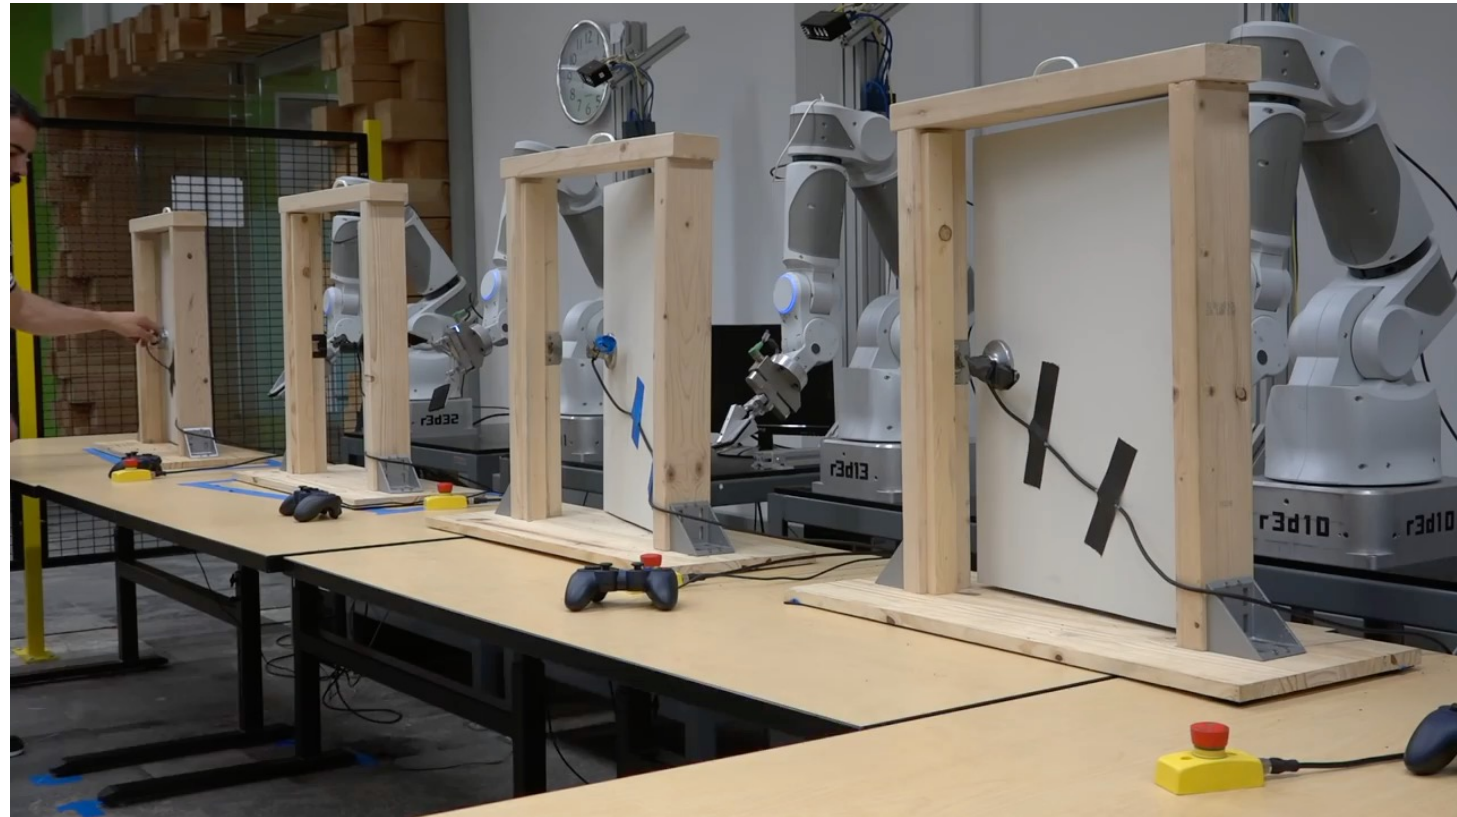
(Berkeley Robot for the Elimination of Tedious Tasks)

BRETT has acquired the ability to learn to perform tasks on its own through trial and error.

Google Door Opening Project

- Learn to open doors using Reinforcement learning
 - Learning reward: opening the door
 - Much harder than purely digital learning: very slow iterations!
 - Simulation only helps a bit: real world much more complex
- Google and UC Berkeley Sergey Levine
- Google very secretive ...

<https://www.wired.com/2017/01/googles-gol-playing-machine-opens-door-robots-learn/>



Nvidia end-to-end deep learning self driving car

- Raw pixel of single camera => steering command of car
 - 30 FPS; single-image control (no history)
 - Nvidia Drive PX car computer (ARM cores and Nvidia GPU)
 - Human steering angle for training
- End-to-end: NO explicit:
 - Lane detection
 - Road detection
 - Obstacle detection

<https://devblogs.nvidia.com/parallelforall/deep-learning-self-driving-cars/>



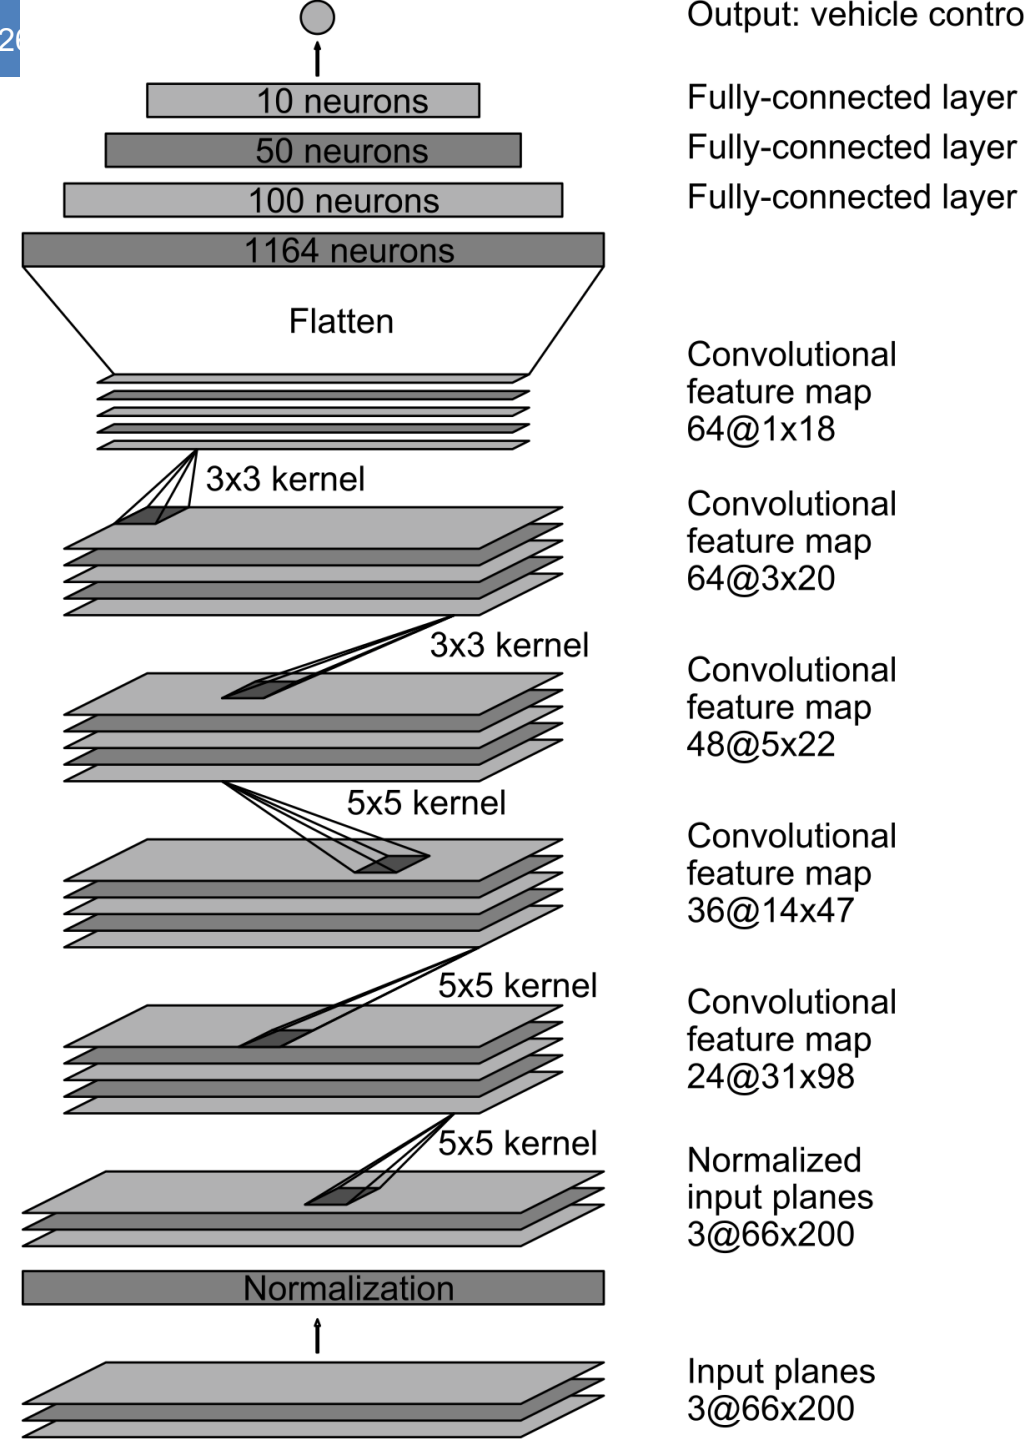
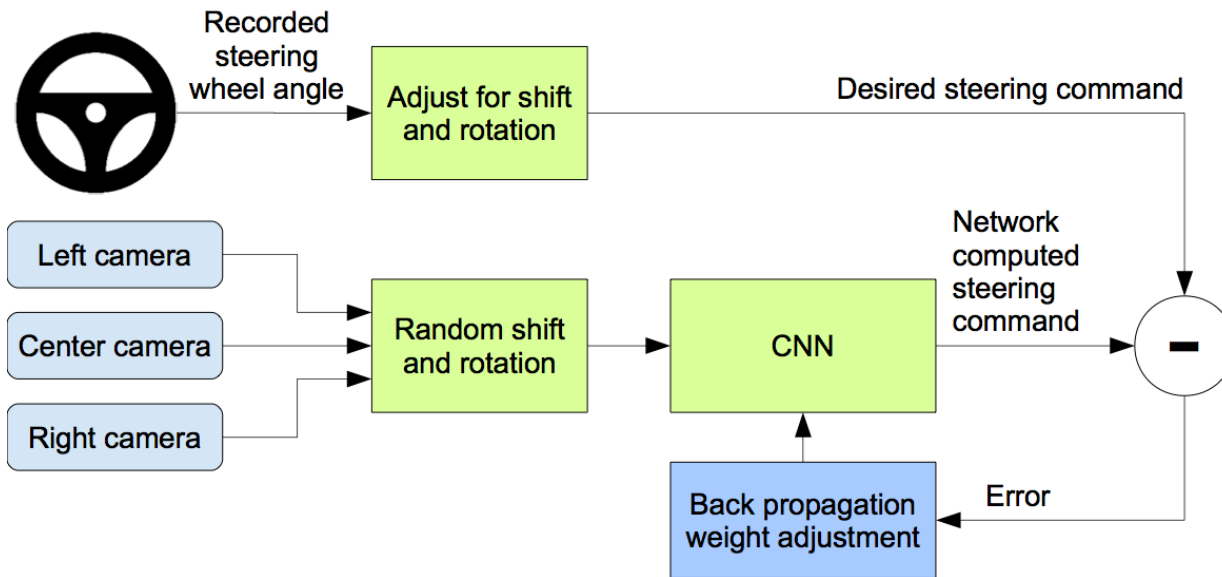
DAVE 2 Driving a Lincoln

- A convolutional neural network
- Trained by human drivers
- Learns perception, path planning, and control "pixel in, action out"
- Front-facing camera is the only sensor

• Training:

- Additionally use left and right camera: negative examples!
- Highway, residential roads, unpaved roads, car parks
- Different weather conditions

• Result: Autonomous 98% of the time => 2% driver intervention



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

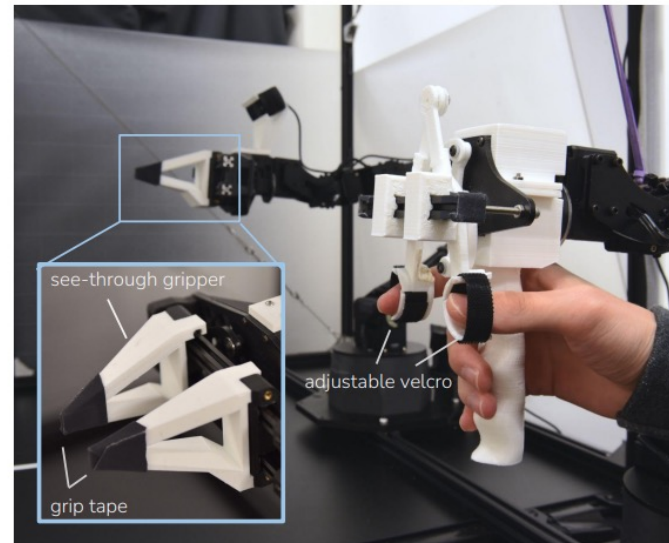
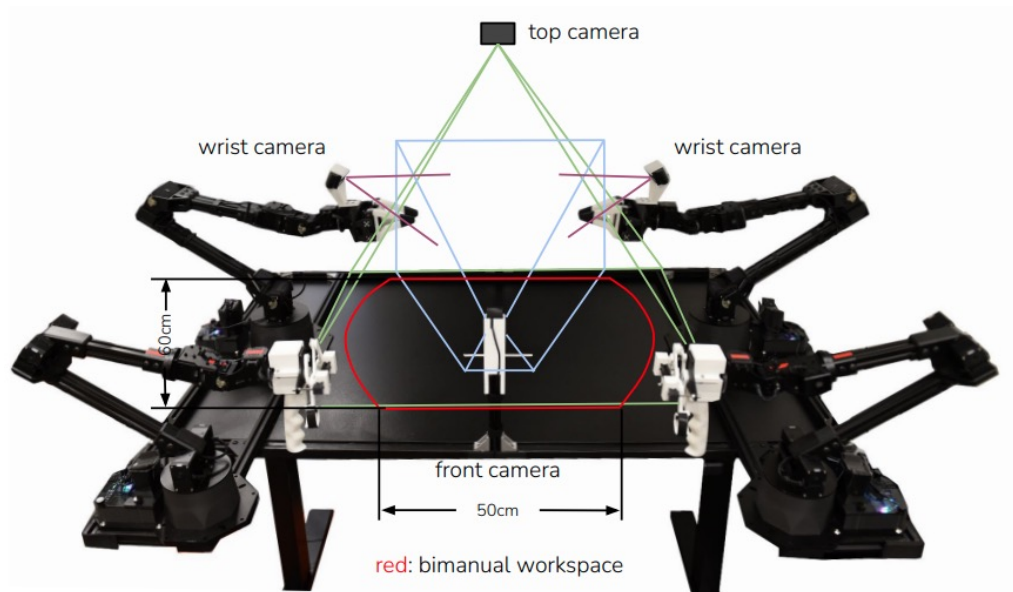
Do As I Can, Not As I Say: Grounding Language in Robotic Affordances

Customers want a
robot that handles

all household tasks

and is **commanded
by natural
language**

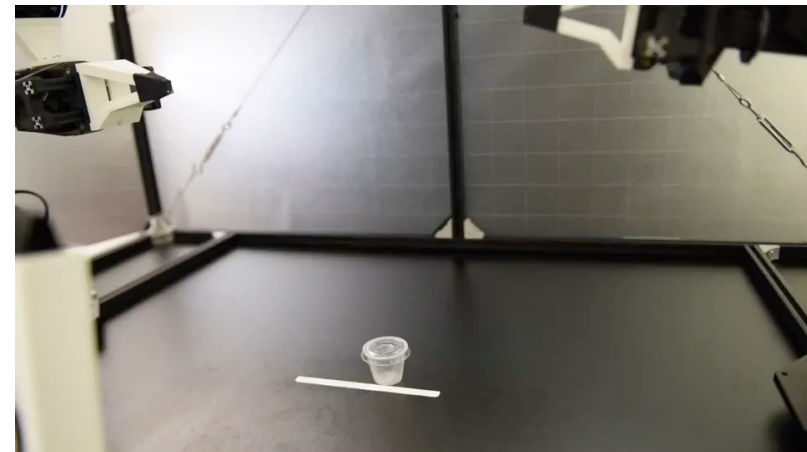
Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware (ACT / ALOHA)



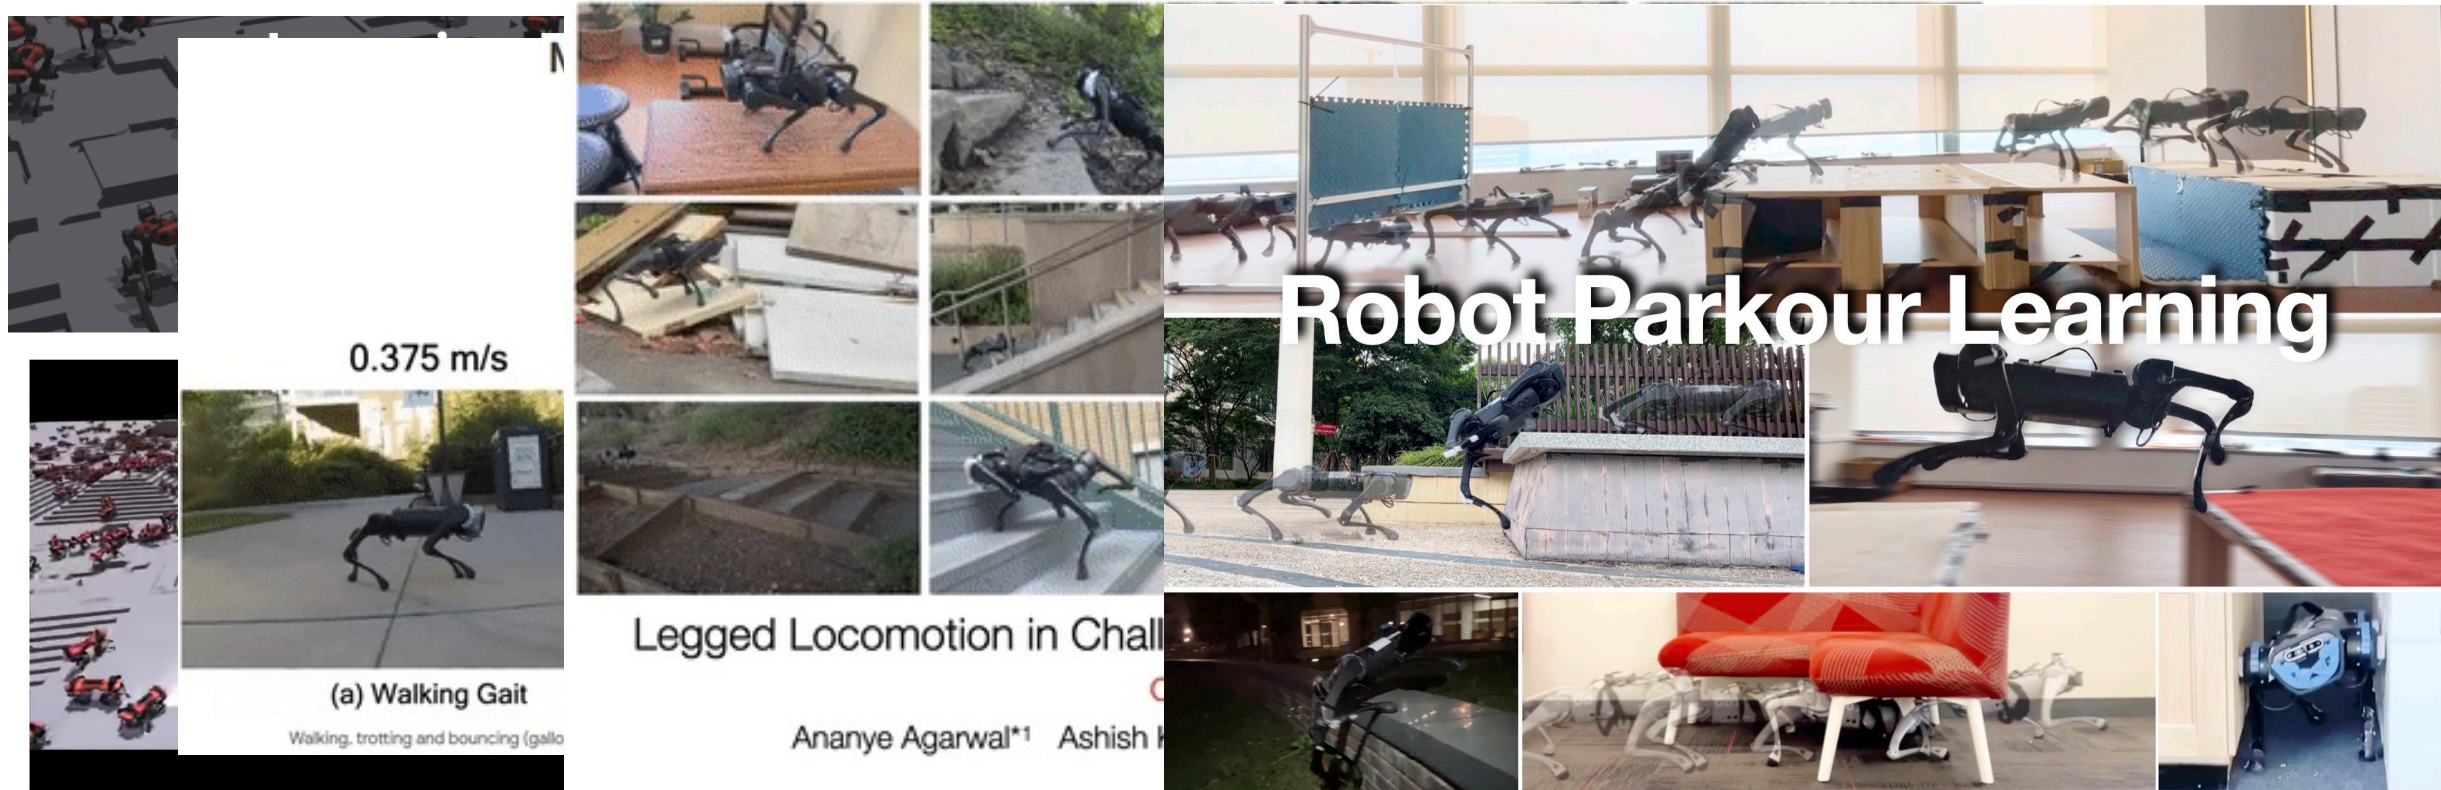
ViperX 6dof Arm (follower)

#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

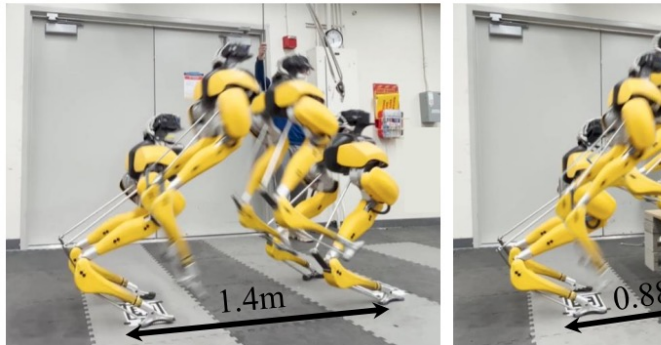
Athletic Intelligence (bipedal robot)

Learning Humanoid Locomotion with Transformers

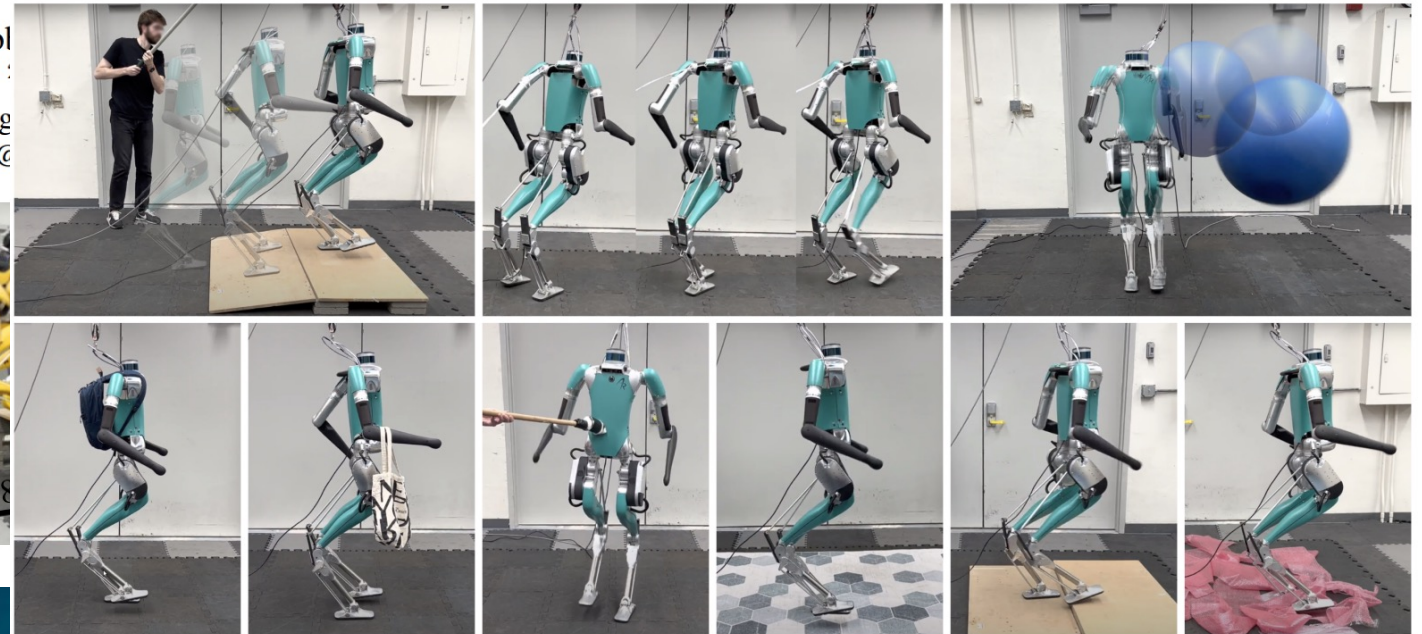
Robust and Versatile Locomotion through Reinforcement Learning

Ilija Radosavovic* Tete Xiao* Bike Zhang* Trevor Darrell† Jitendra Malik† Koushil Sreenath†
University of California, Berkeley

Zhongyu Li¹, Xue Bin Peng², Pieter Abbeel¹
¹University of California, Berkeley,
Email: zhongyu_li@berkeley.edu, xbpeng@eecs.berkeley.edu, pietera@eecs.berkeley.edu



(a)



accordingly.

Athletic Intelligence

Connections between perception and control allow Atlas to adapt—quite literally—on the fly.

world in meaningful ways.

environment.

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

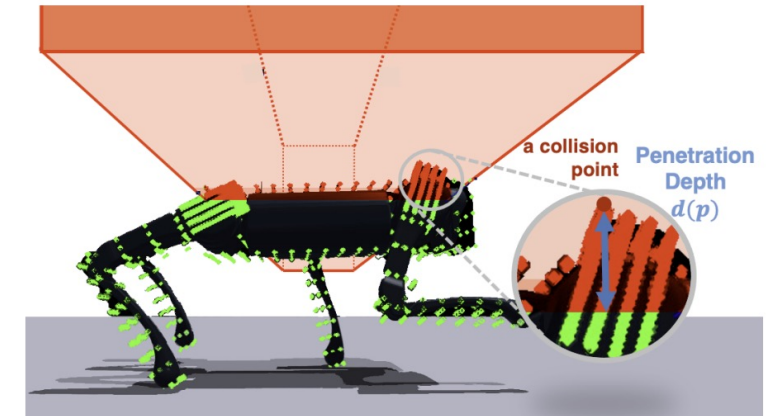
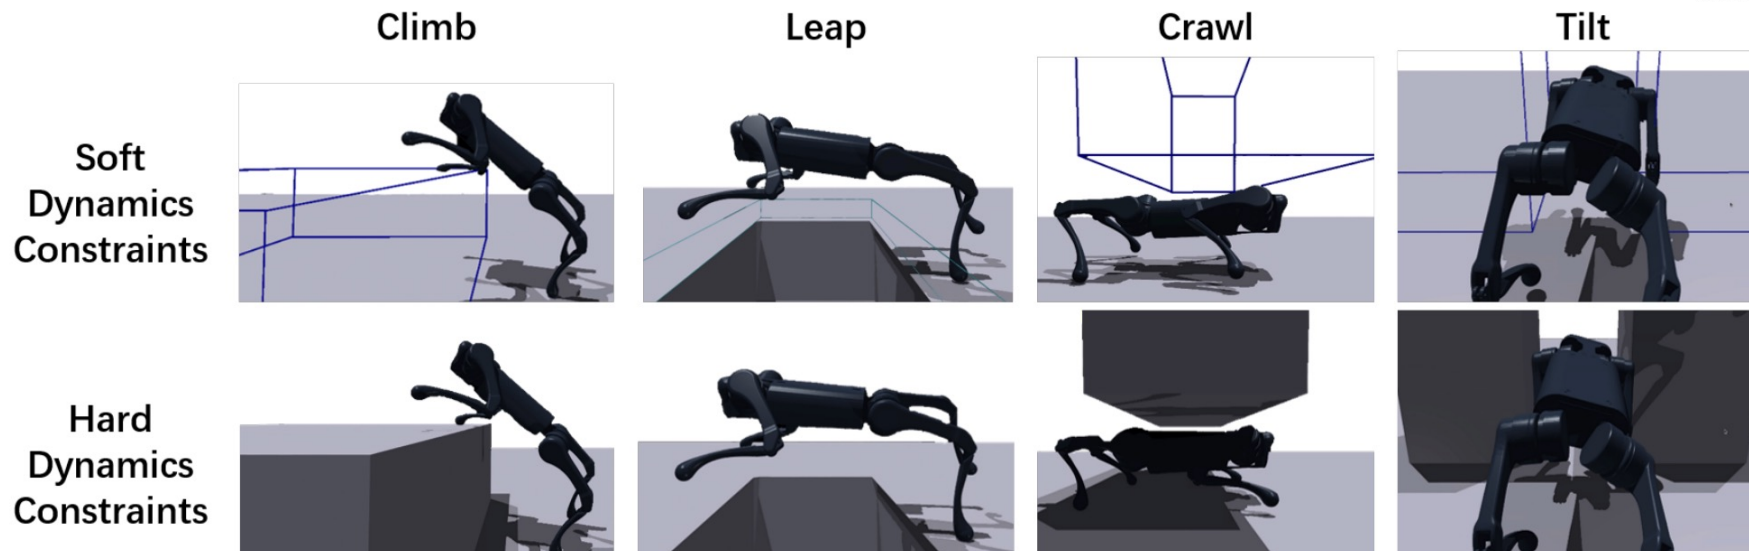
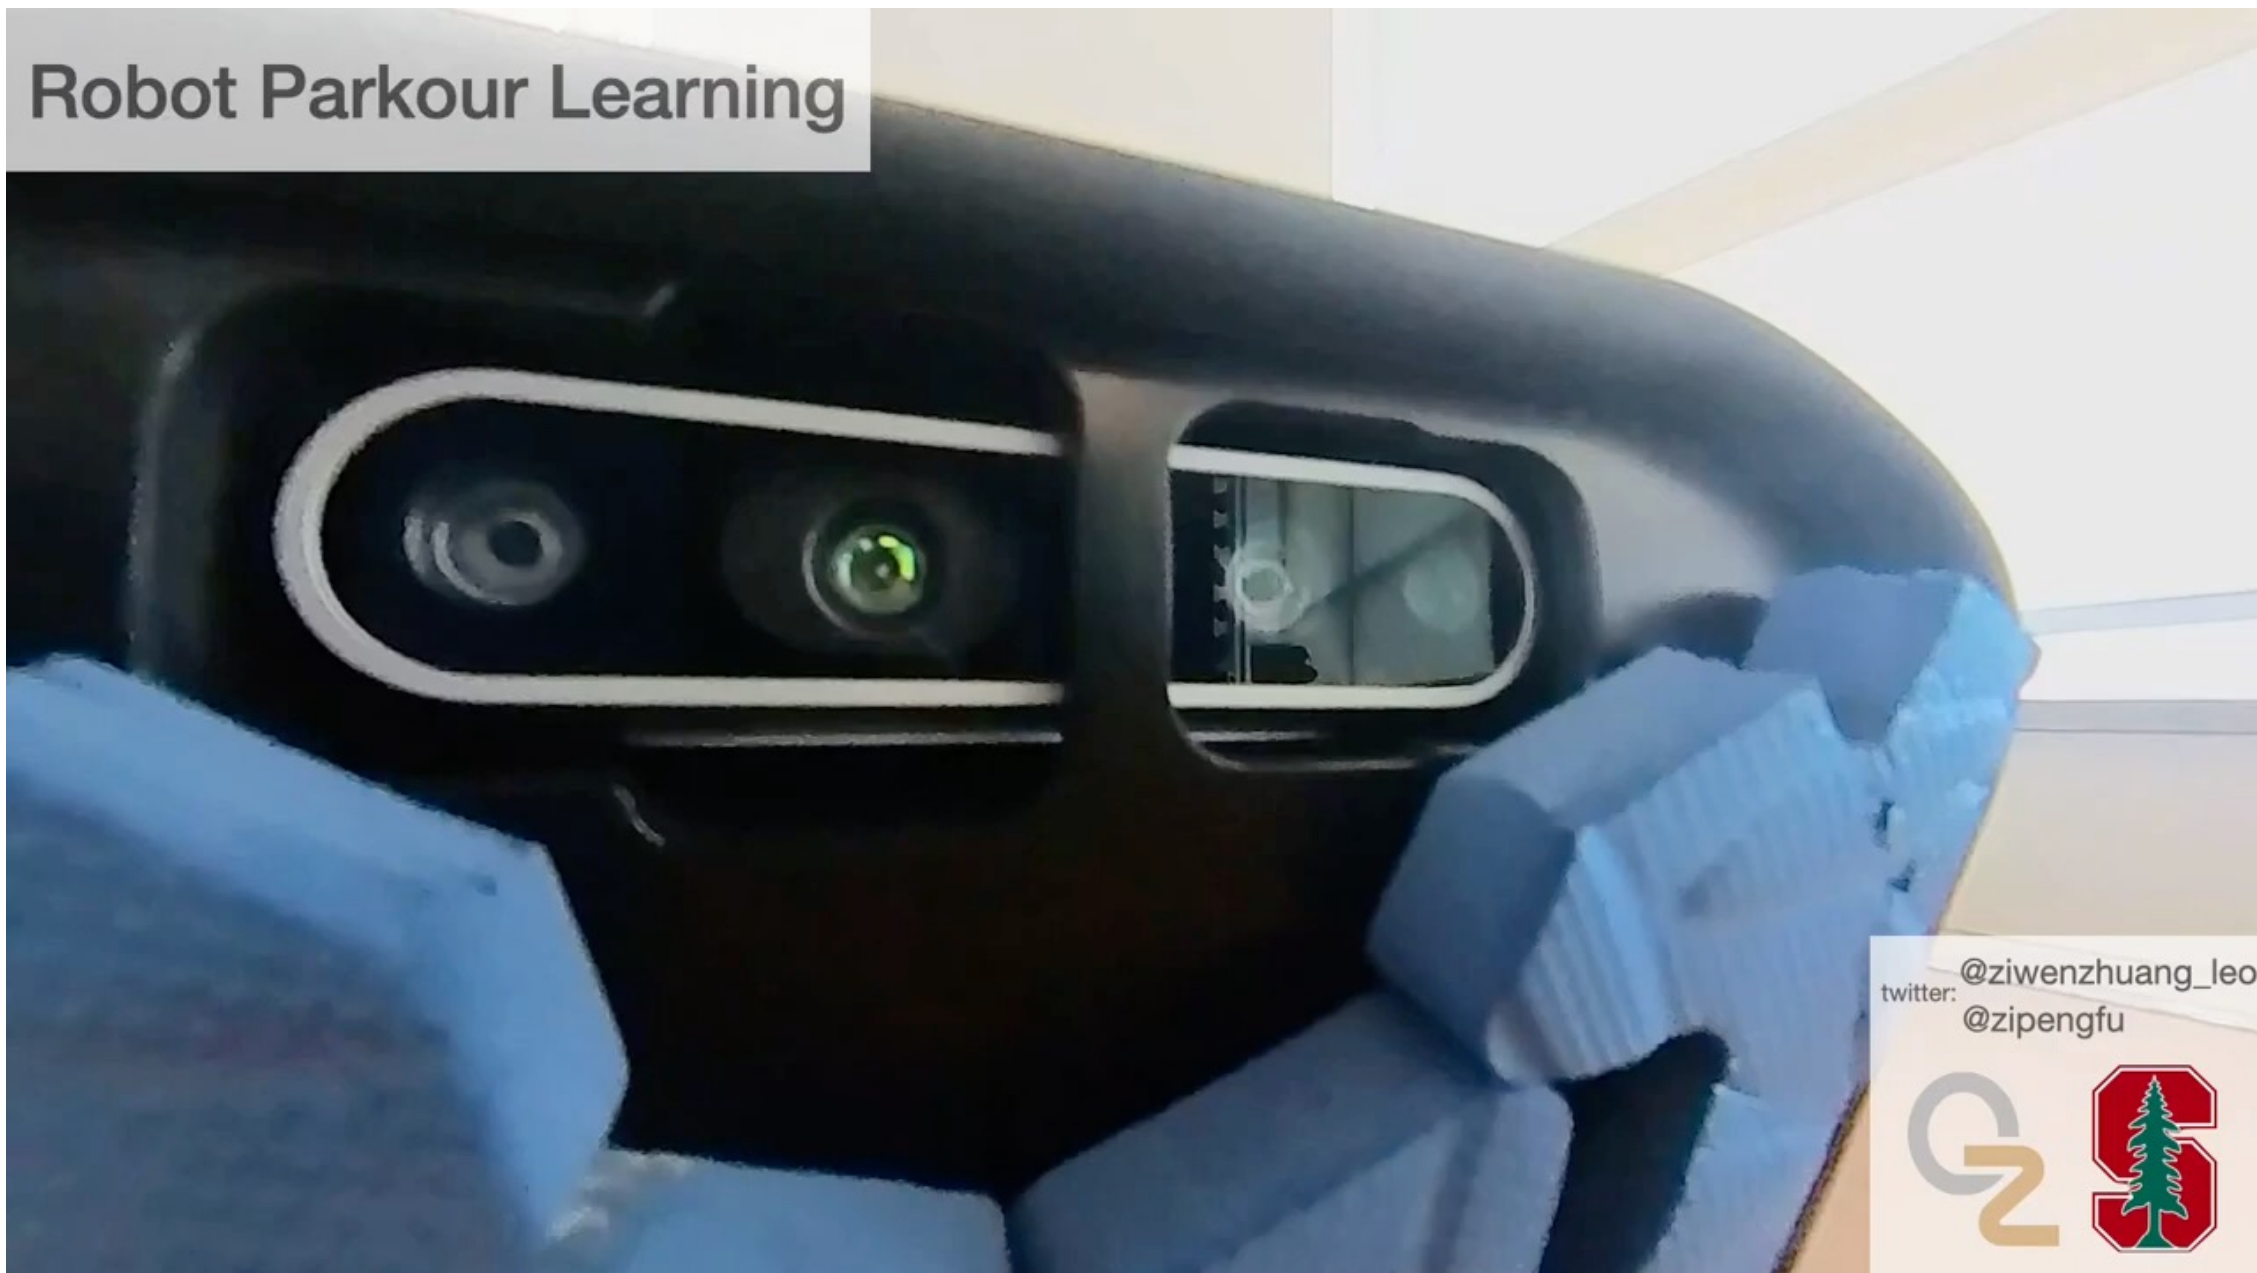


Figure 4: We show collisions points on the robot. Collision points that penetrate obstacles are in red.



Robot Parkour Learning



<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

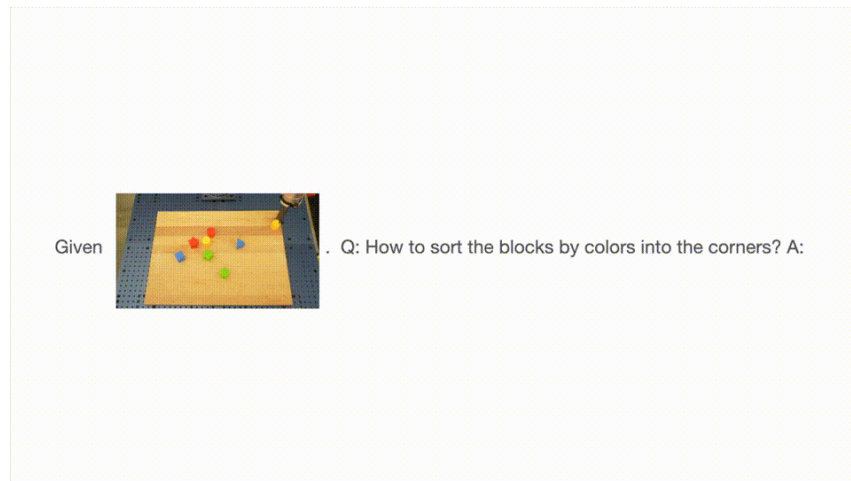
ShanghaiTech University

Large Language AI Models for Autonomous Mobile Robotics

- Visual language manipulation



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

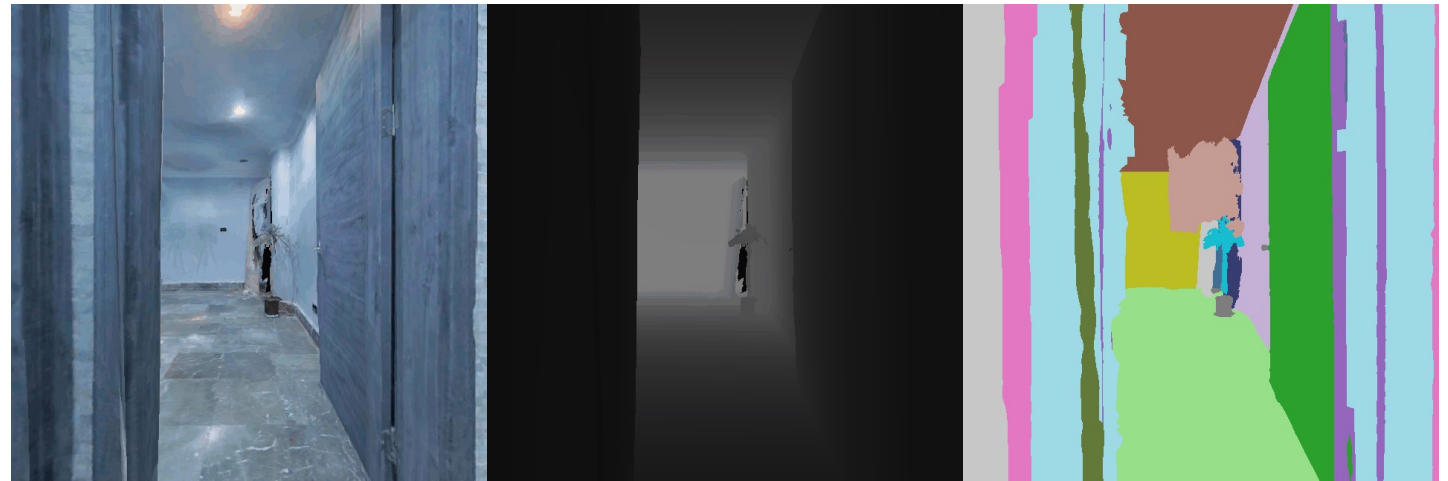


Task: How to sort the blocks by colors into the corners?

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>, <column, 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.

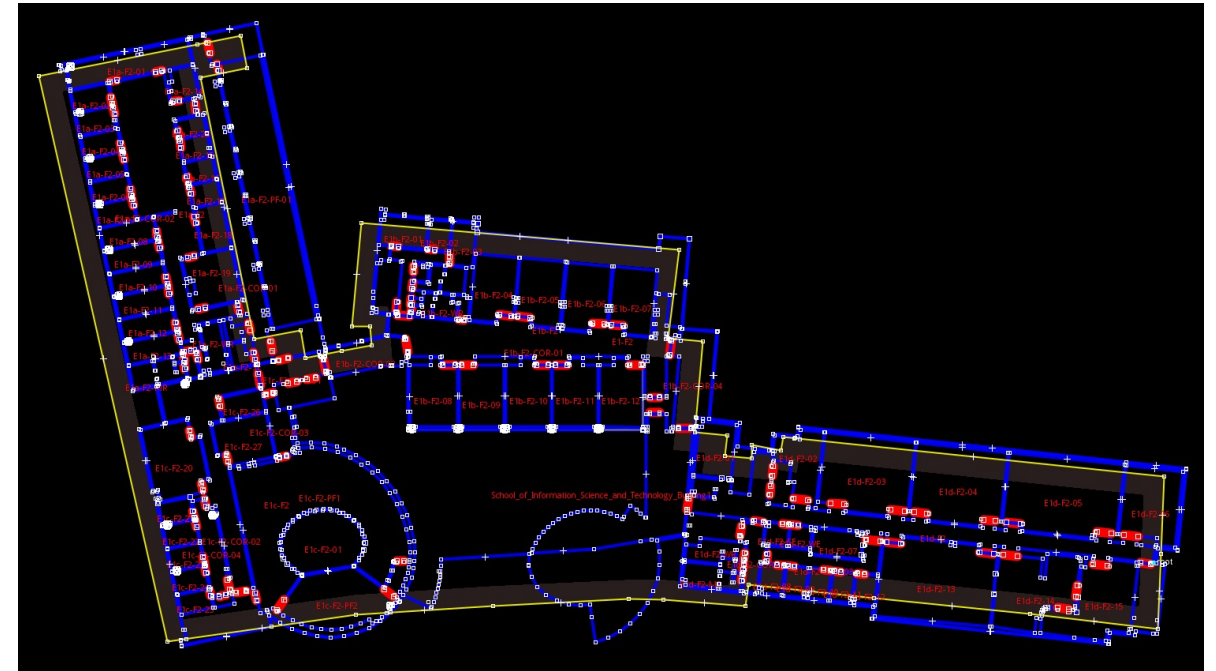
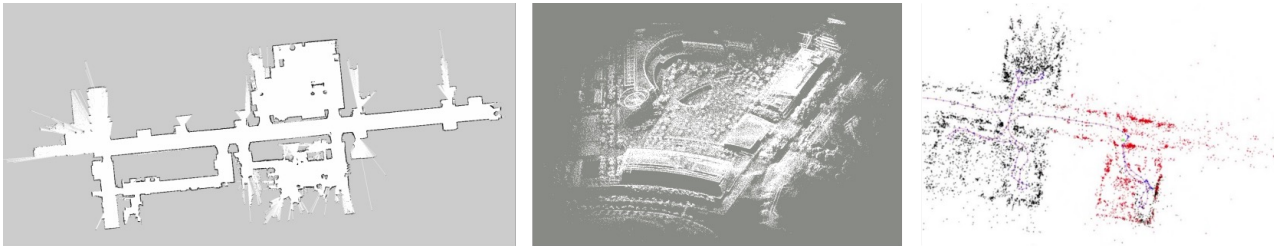


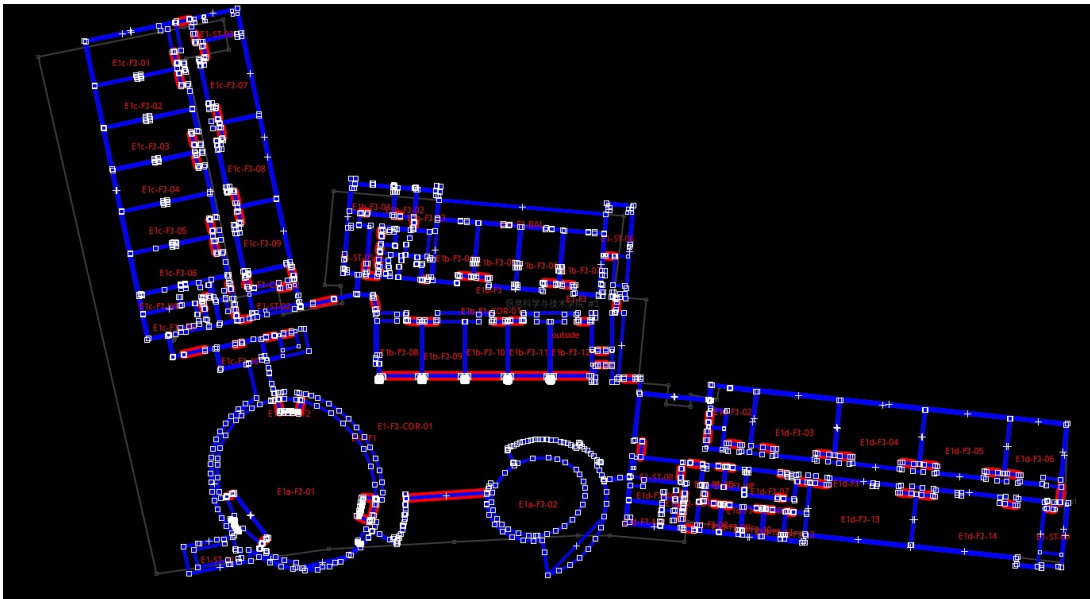
Fig. 2 A osmAG (Area Graph in OpenStreetMap textual format)

Map Representations

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



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



```
<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>
```

```
<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.

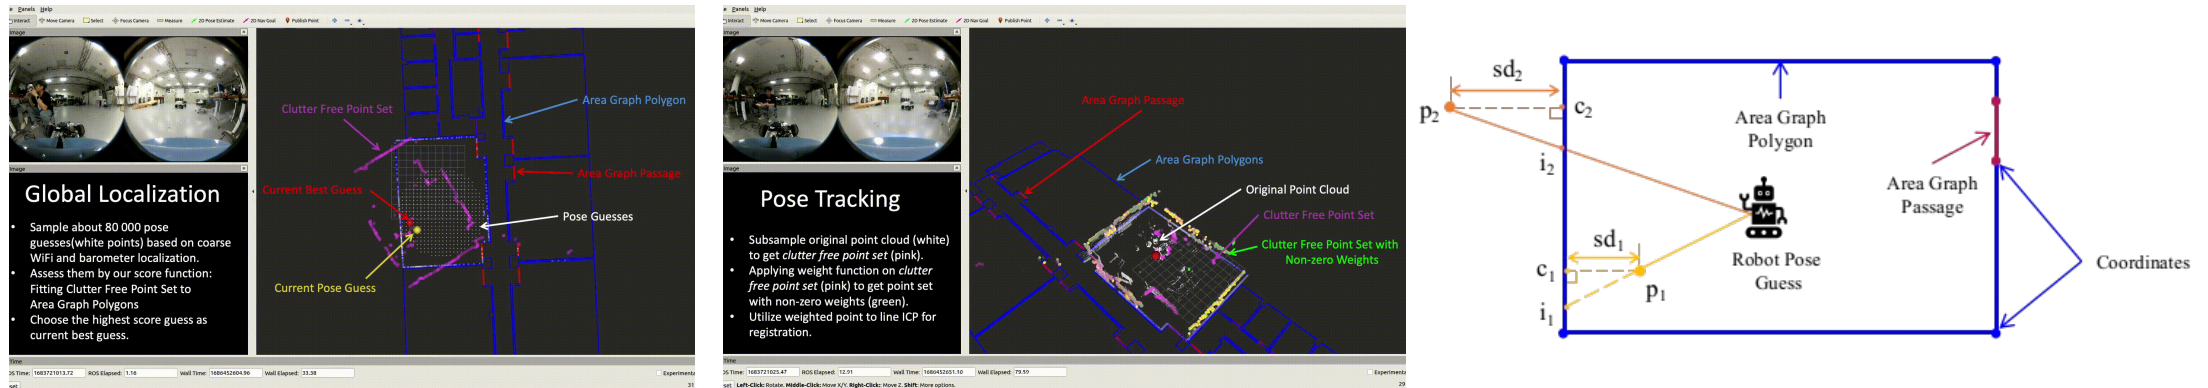


Fig. 1

- **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.

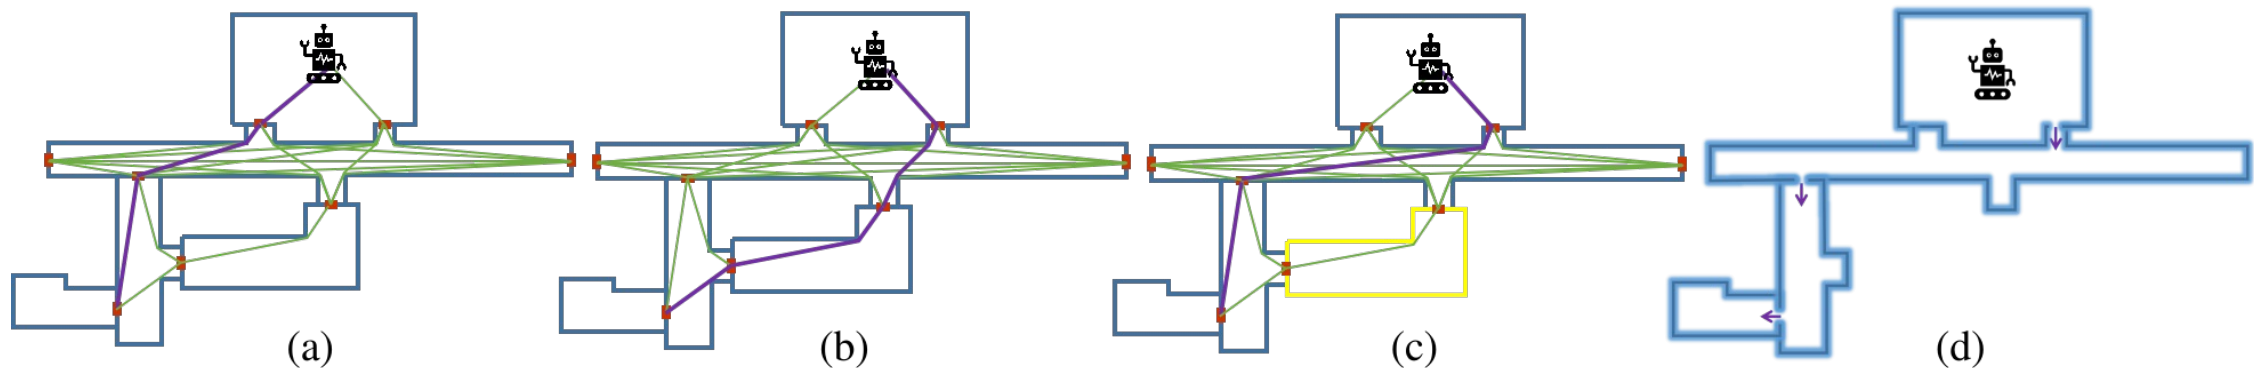
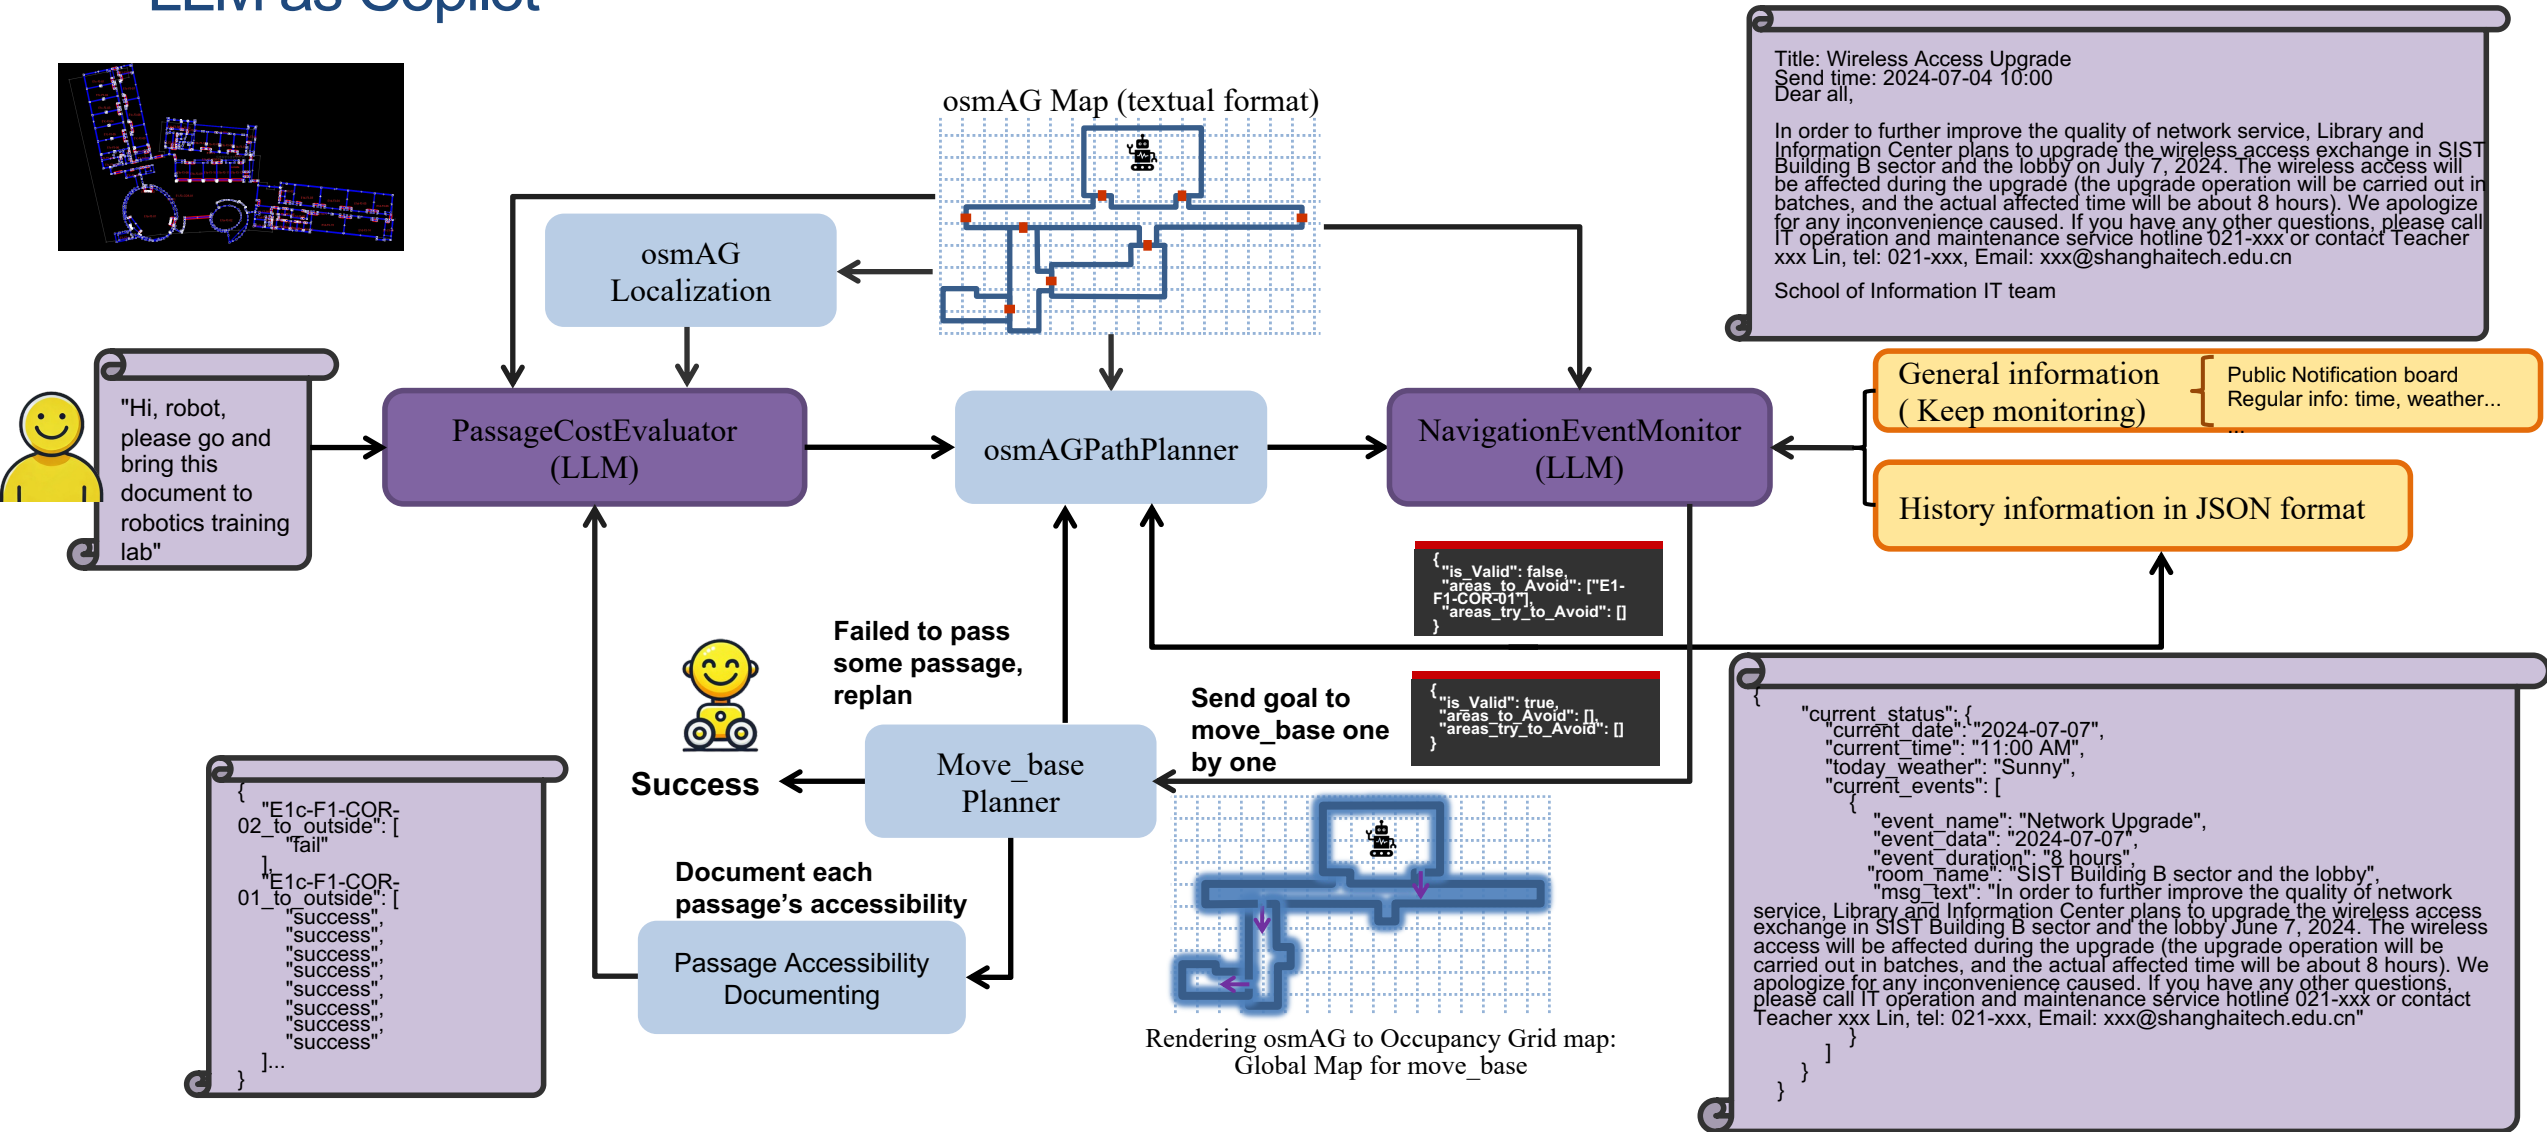


Fig. 2

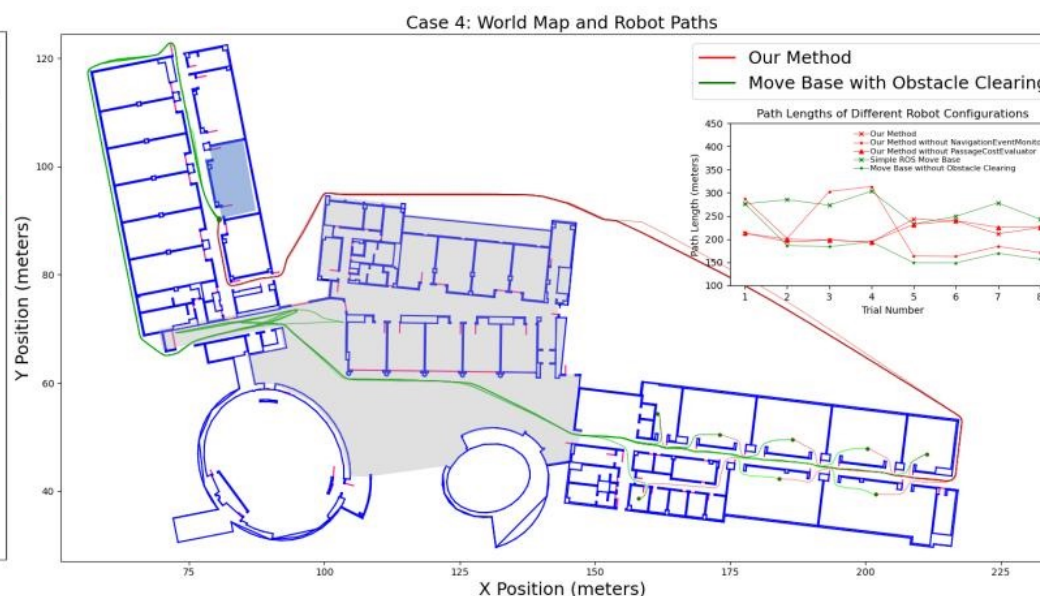
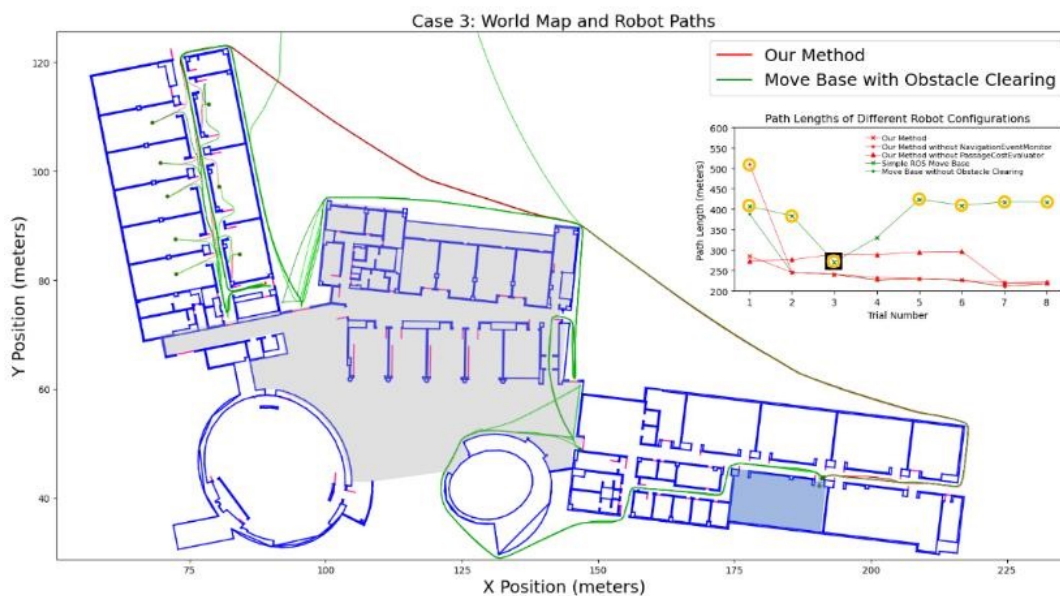
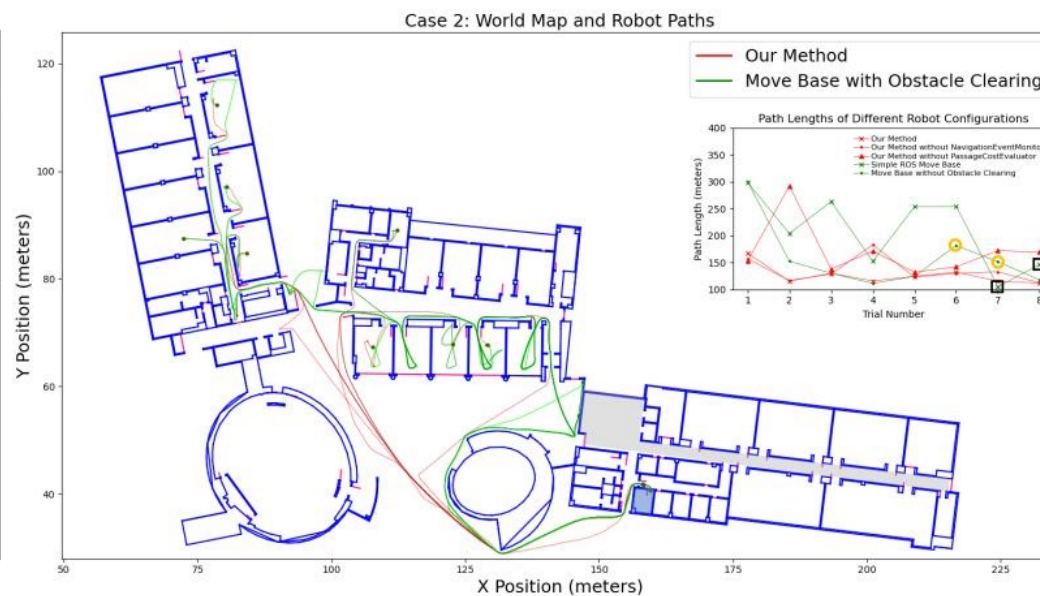
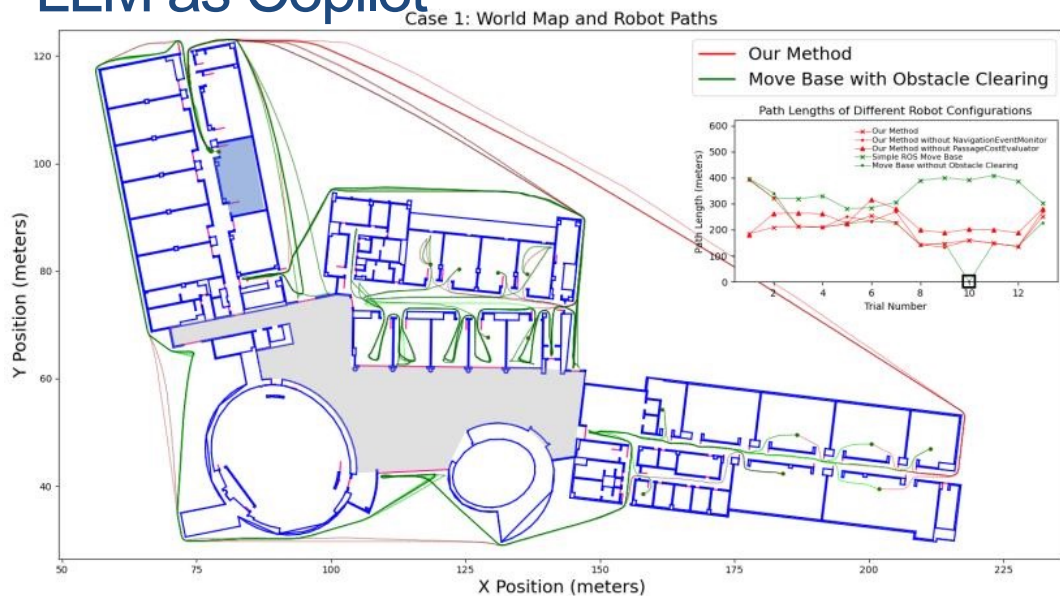
Intelligent LiDAR Navigation: Leveraging External Information and Semantic Maps with LLM as Copilot



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 (%)

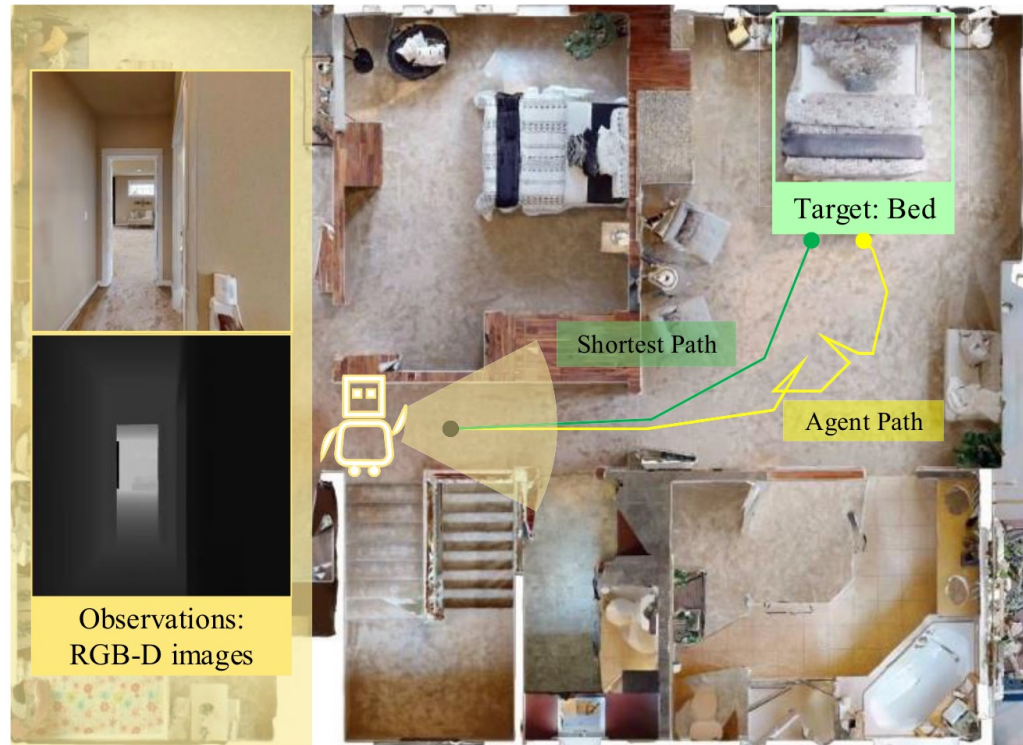
Model	Accuracy	Recall	False Positive Rate
ChatGPT-4o	0.96	1.0	0.93
DeepSeek-V3	0.93	0.9	0.95

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 <i>NavigationEventManager</i>	221.2	135.4	264.3	223.3
Our method w/o <i>PassageCostEvaluator</i>	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

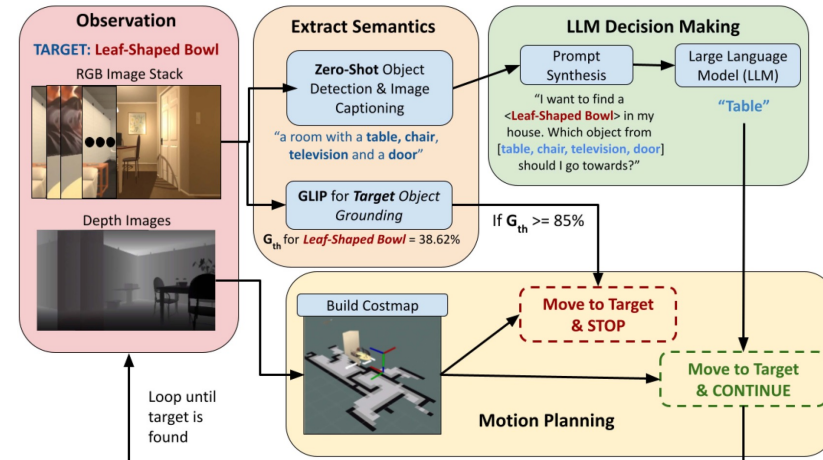
Furthermore -- Object Goal Navigation

- Object-Goal Navigation



(a)

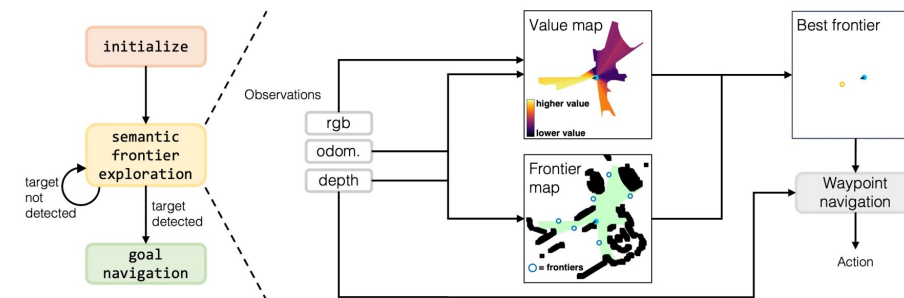
- Example 1



[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



(c)

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

[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.

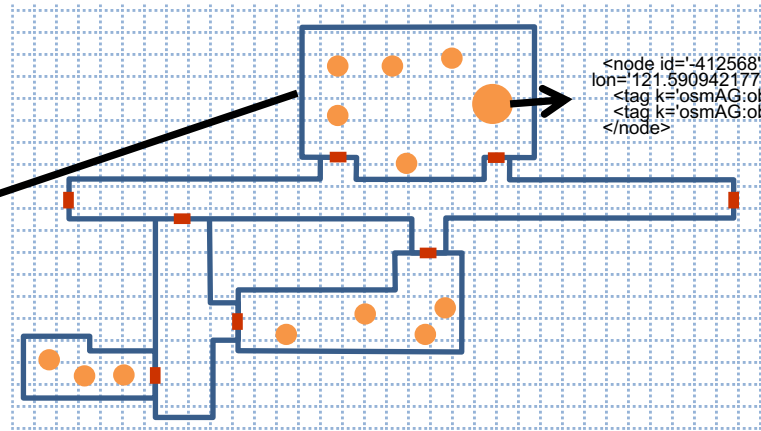
Furthermore -- Object Goal Navigation with osmAG

• Mapping Phase:

- Adding room description, and nodes represent objects.
- Off-line, using handheld apple scanner, using LabelMaker[1] and VLMs to label objects.

```
<way id='-189708' action='modify' visible='true'>
  <nd ref='-344898' />
```

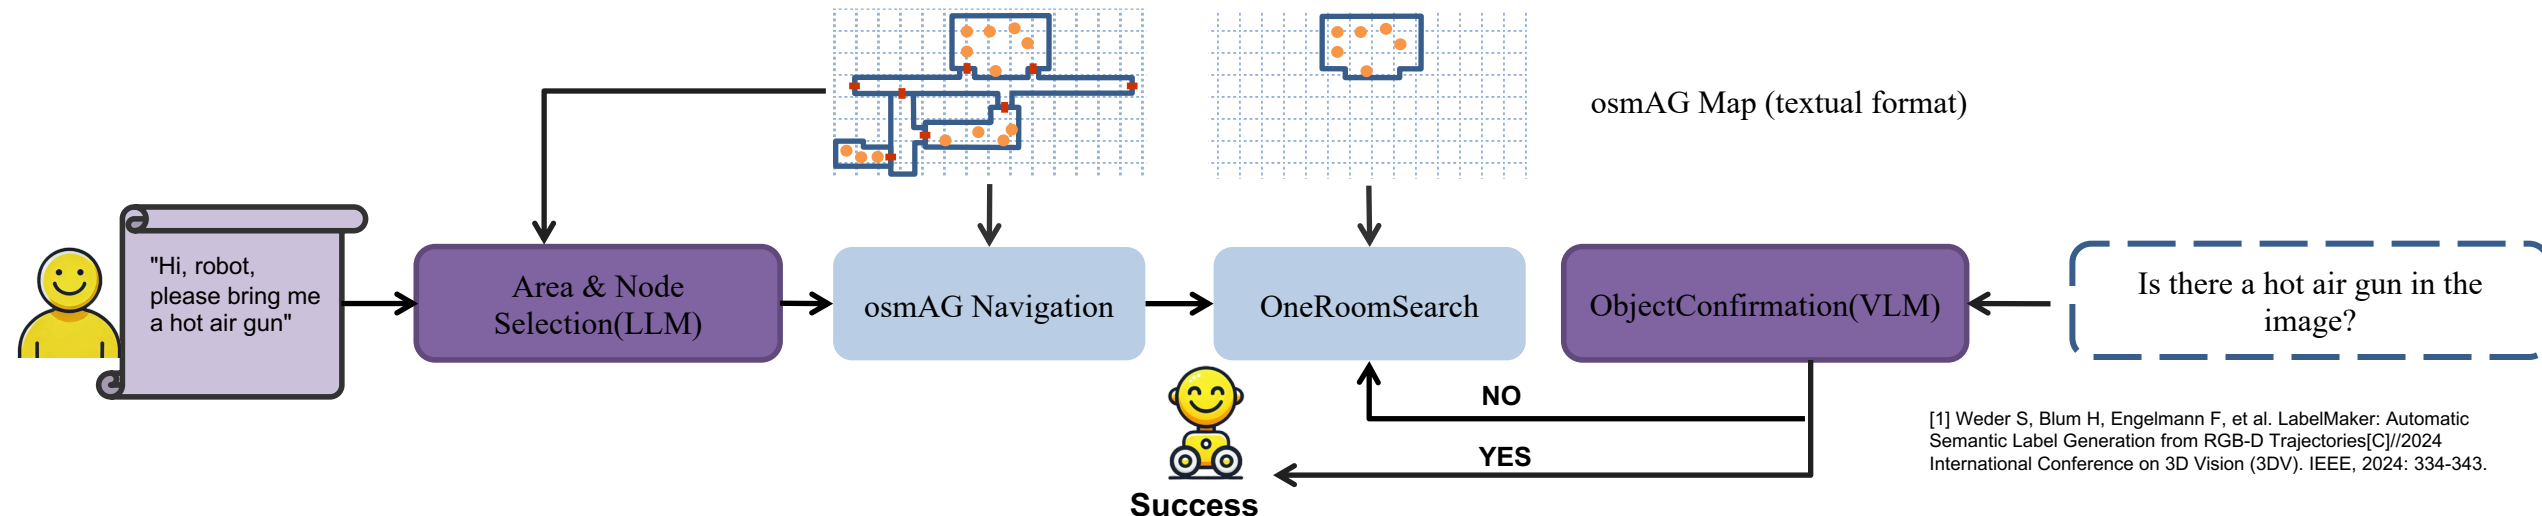
```
  <tag k='height' v='3.2' />
  <tag k='indoor' v='room' />
  <tag k='level' v='2' />
  <tag k='name' v='E1d-F2-03' />
  <tag k='osmAG:areaType' v='room' />
  <tag k='osmAG:parent' v='E1d-F2' />
  <tag k='osmAG:room_number' v='1D-204' />
  <tag k='osmAG:type' v='area' />
  <tag k='osmAG:area_description' v='This room is a well-organized robotics lab, featuring designated workstations equipped with tools, computers, and various electronic equipment for robotics assembly and testing. Along one wall, a row of lockers and cabinets provides storage space, while blue bins and carts hold additional components and supplies. The lab has a combination of functional work surfaces and a few recreational elements, like a foosball table, offering a balanced environment for both work and relaxation. A pet food and water station in one corner suggests a pet-friendly space, adding a touch of comfort. Overall, the layout and resources are designed to support efficient object access, tool organization, and task-specific areas.' />
```



```
<node id='-412568' lat='31.17961192718'
  lon='121.5909421779'>
  <tag k='osmAG:object_name' v='desk_1' />
  <tag k='osmAG:object_parent' v='E1d-F2-03' />
</node>
```

• 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

设备名称	添加方式	地点	控制	设备型号	当前使用者	联系人	负责人
图像分析扫描设备	对接录入	物资字间特楼 B104			王雷	马后康	马后康, 孟宪阳, 朱海阳, 张青, 陈晨, 马后康
Fusion-2012 (Optima Fusion)	对接录入	人字楼 B109			朱斌	朱斌, 高秀霞	陈宏源, 朱斌, 高秀霞
超速离心机 (OptimaXP90)	对接录入	人字楼C区五层置自热机器与新药实验室508			陈杰	杨秀娜	杨秀娜
超速离心机 (OptimaXP90)	对接录入	人字楼C区五层置自热机器与新药实验室508			陈杰	杨秀娜	杨秀娜
超速离心机 (OptimaXP90)	对接录入	人字楼A214-3			陈杰	史巧云	史巧云, 陈晶内, 李瑞莹
超速离心机 (OptimaXP90)	对接录入	人字楼 A214-3			史巧云	史巧云, 陈晶内, 李瑞莹	史巧云, 陈晶内, 李瑞莹
AKTA 25+ 人字楼C区生物化学实验室	对接录入	人字楼C区二层的置自热平台实验室 211			罗甲	史巧云	史巧云, 陈晶内, 李瑞莹
蛋白质纯化系统 fast protein liquid chromatography_3d (AKTA Pure M1)	对接录入	生命科学楼 1楼 B110			李俊贤	李俊贤	李俊贤
多功能材料制备与测试系统 (N2348)	对接录入	信息学院1号楼楼下一层 108室			孙楠	孙楠	孙楠
分选型流式细胞仪 BD FACSAria Sapp (Sapp)	对接录入	生命科学楼一层 101			曹元元	刘佩燕	刘佩燕
超速离心机 (Averé 20x 20)	对接录入	人字楼C区五层置自热机器与新药实验室508			张泽洋	杨秀娜	杨秀娜
超速离心机 (Averé 20x 20)	对接录入	人字楼C区五层置自热机器与新药实验室504			陈丽东	杨秀娜	杨秀娜
超速离心机 (LYN6000)	对接录入	人字楼C区五层置自热机器与新药实验室 508			刘建华	杨秀娜	杨秀娜

(a)

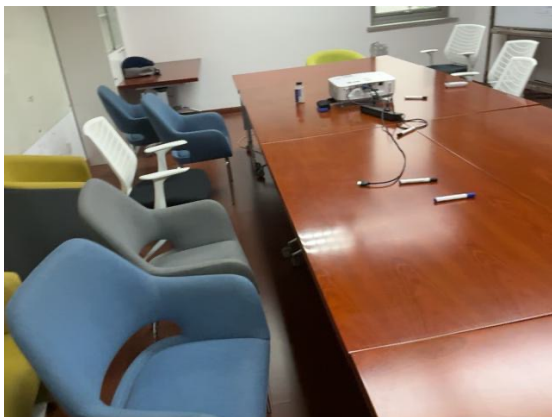
```

<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)

```

<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='Based on the substantial image descriptions provided, the room appears to be a multifunctional workspace, likely an electronics or robotics laboratory or workshop. Here follows 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\nAlthough 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.' />
    
```

(d)

Furthermore -- Object Goal Navigation with osmAG

- Mapping Phase:

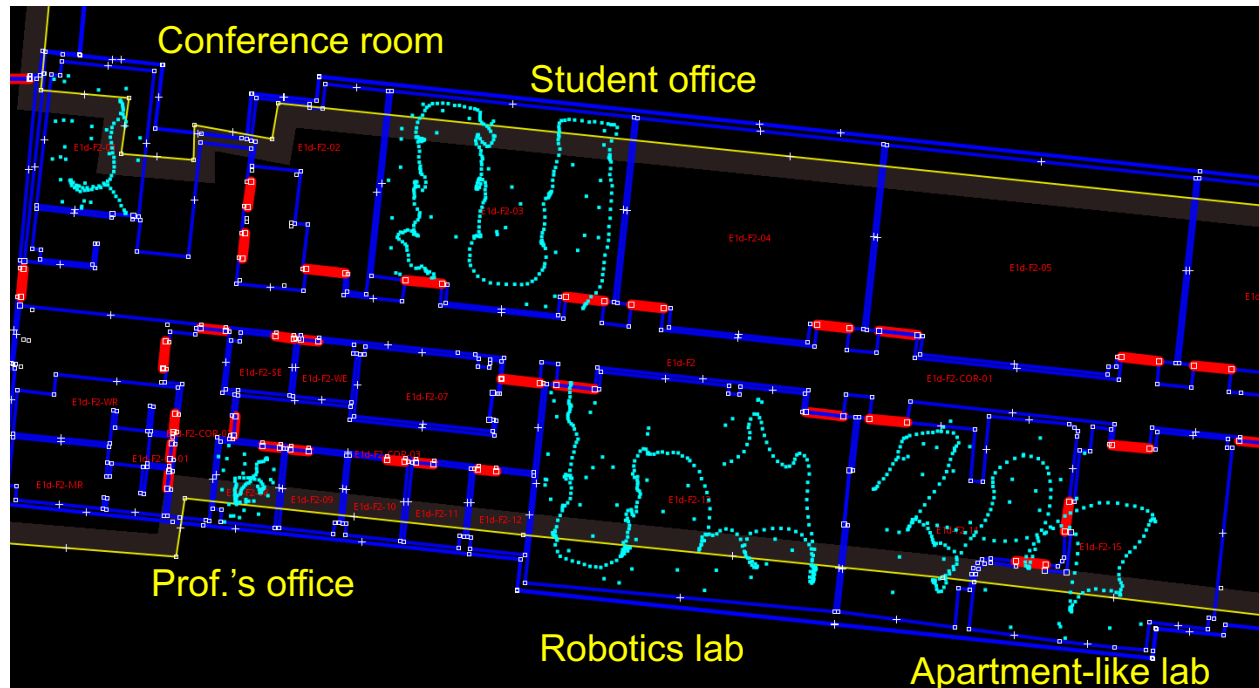


Fig. 1 Mapping result of 5 rooms in SIST building1

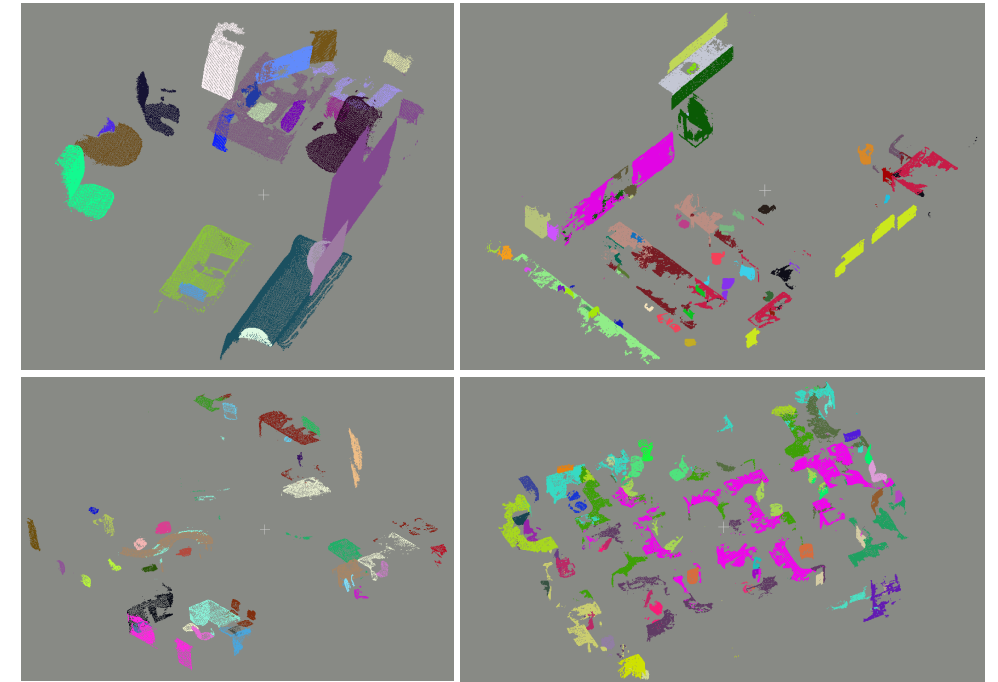


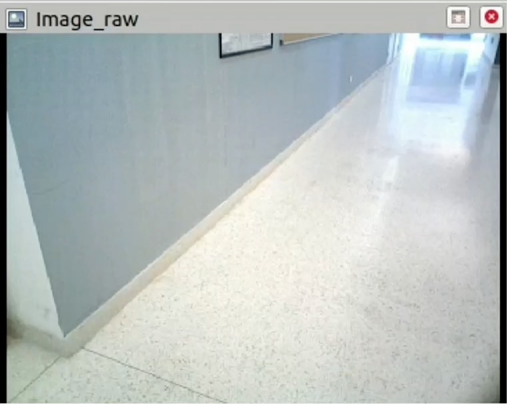
Fig. 2 Instance segmentation result from LabelMaker

Furthermore -- Object Goal Navigation with osmAG

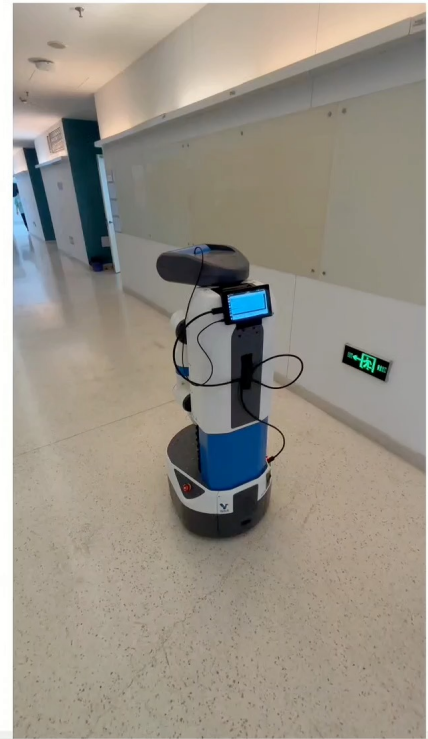
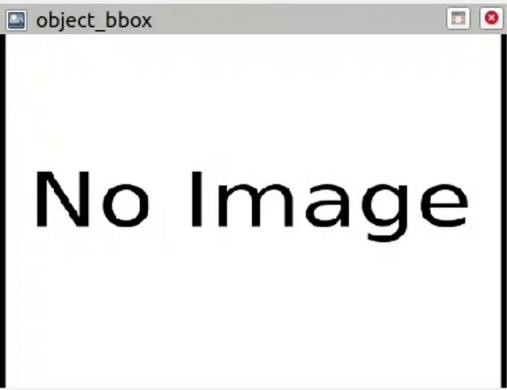
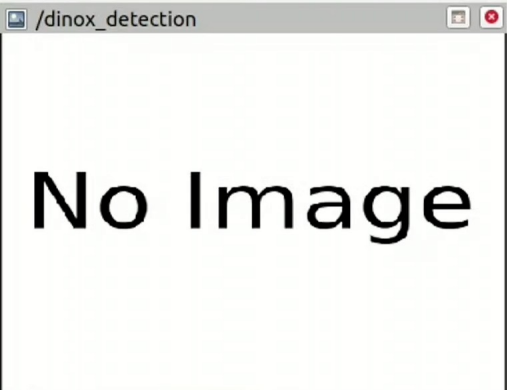
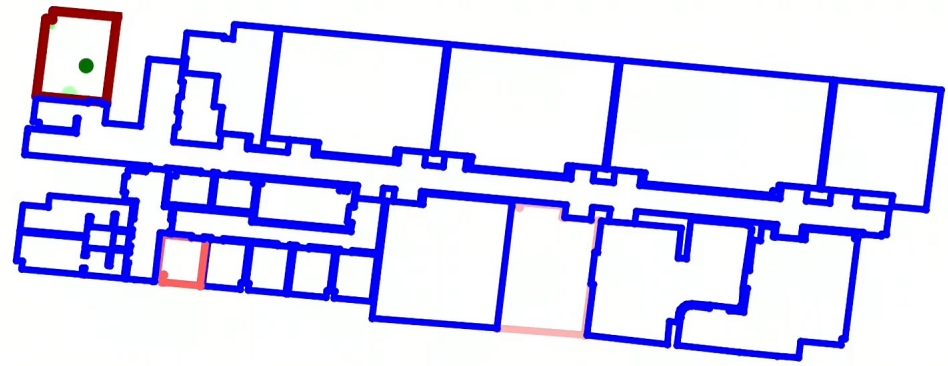
- **Navigation Phase:**




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



```
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)}}
```



AI ETHICS



**- CAN A.I. -
THINK
ETHICALLY?**

By 2050, driverless cars and
mobility-as-a-service will grow
to \$7 trillion worldwide

THE DRIVERLESS DILEMMA

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

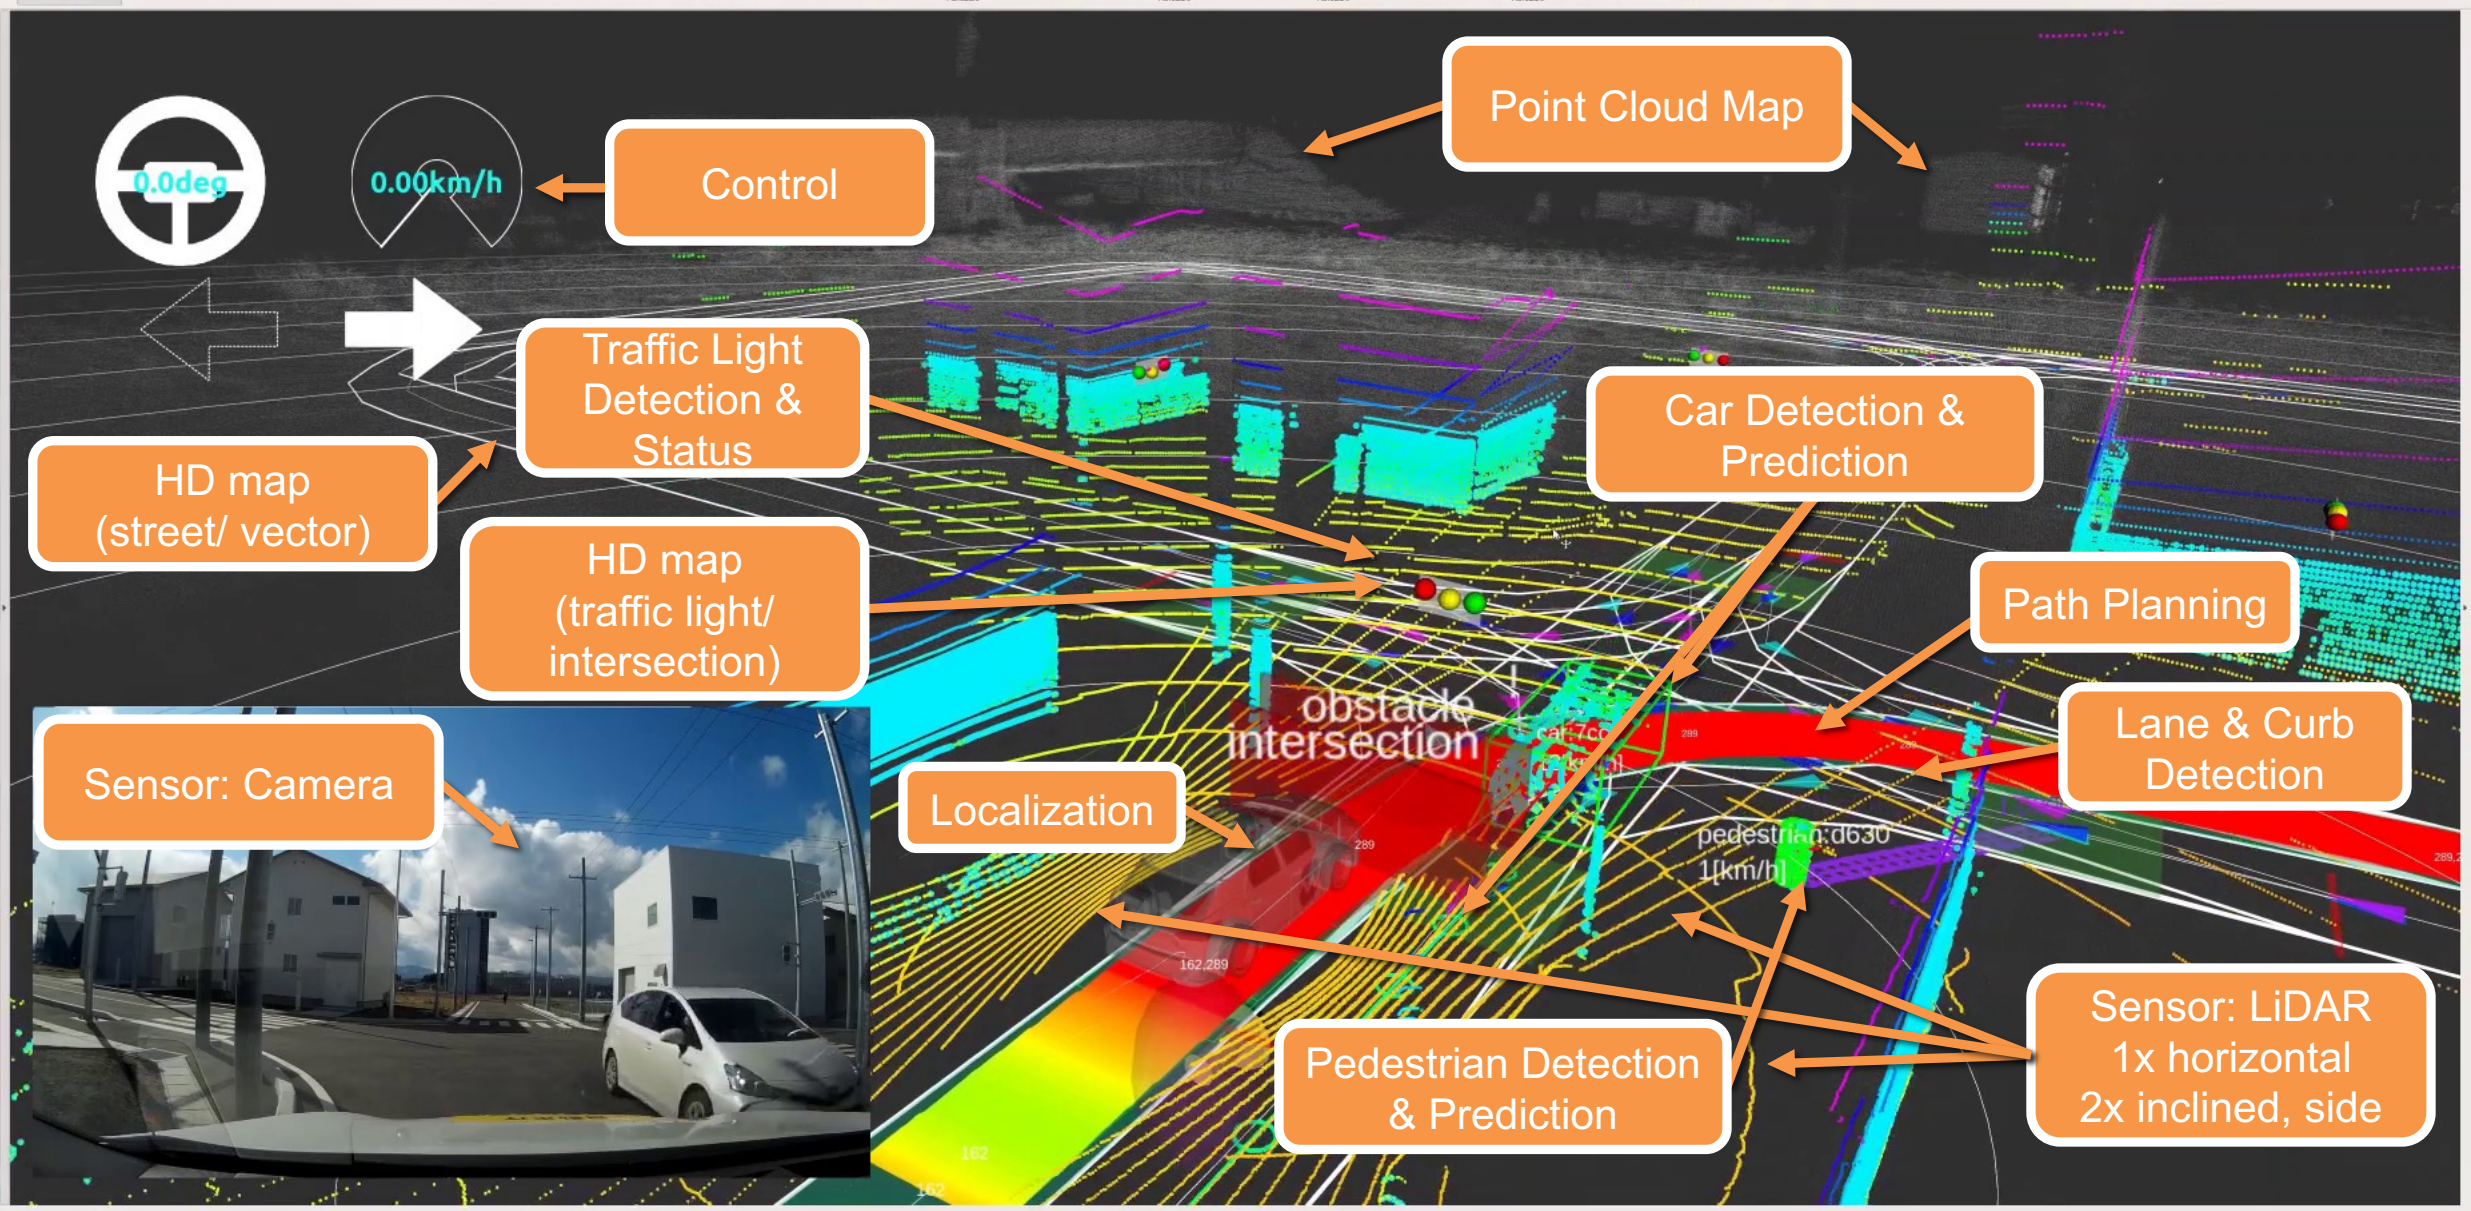


Driverless cars can **ELIMINATE 90% OF TRAFFIC FATALITIES** — Saving 1 million lives every year

HOW WILL DRIVERLESS CARS DETERMINE WHOSE LIFE SHOULD BE SPARED?

File Panels Help

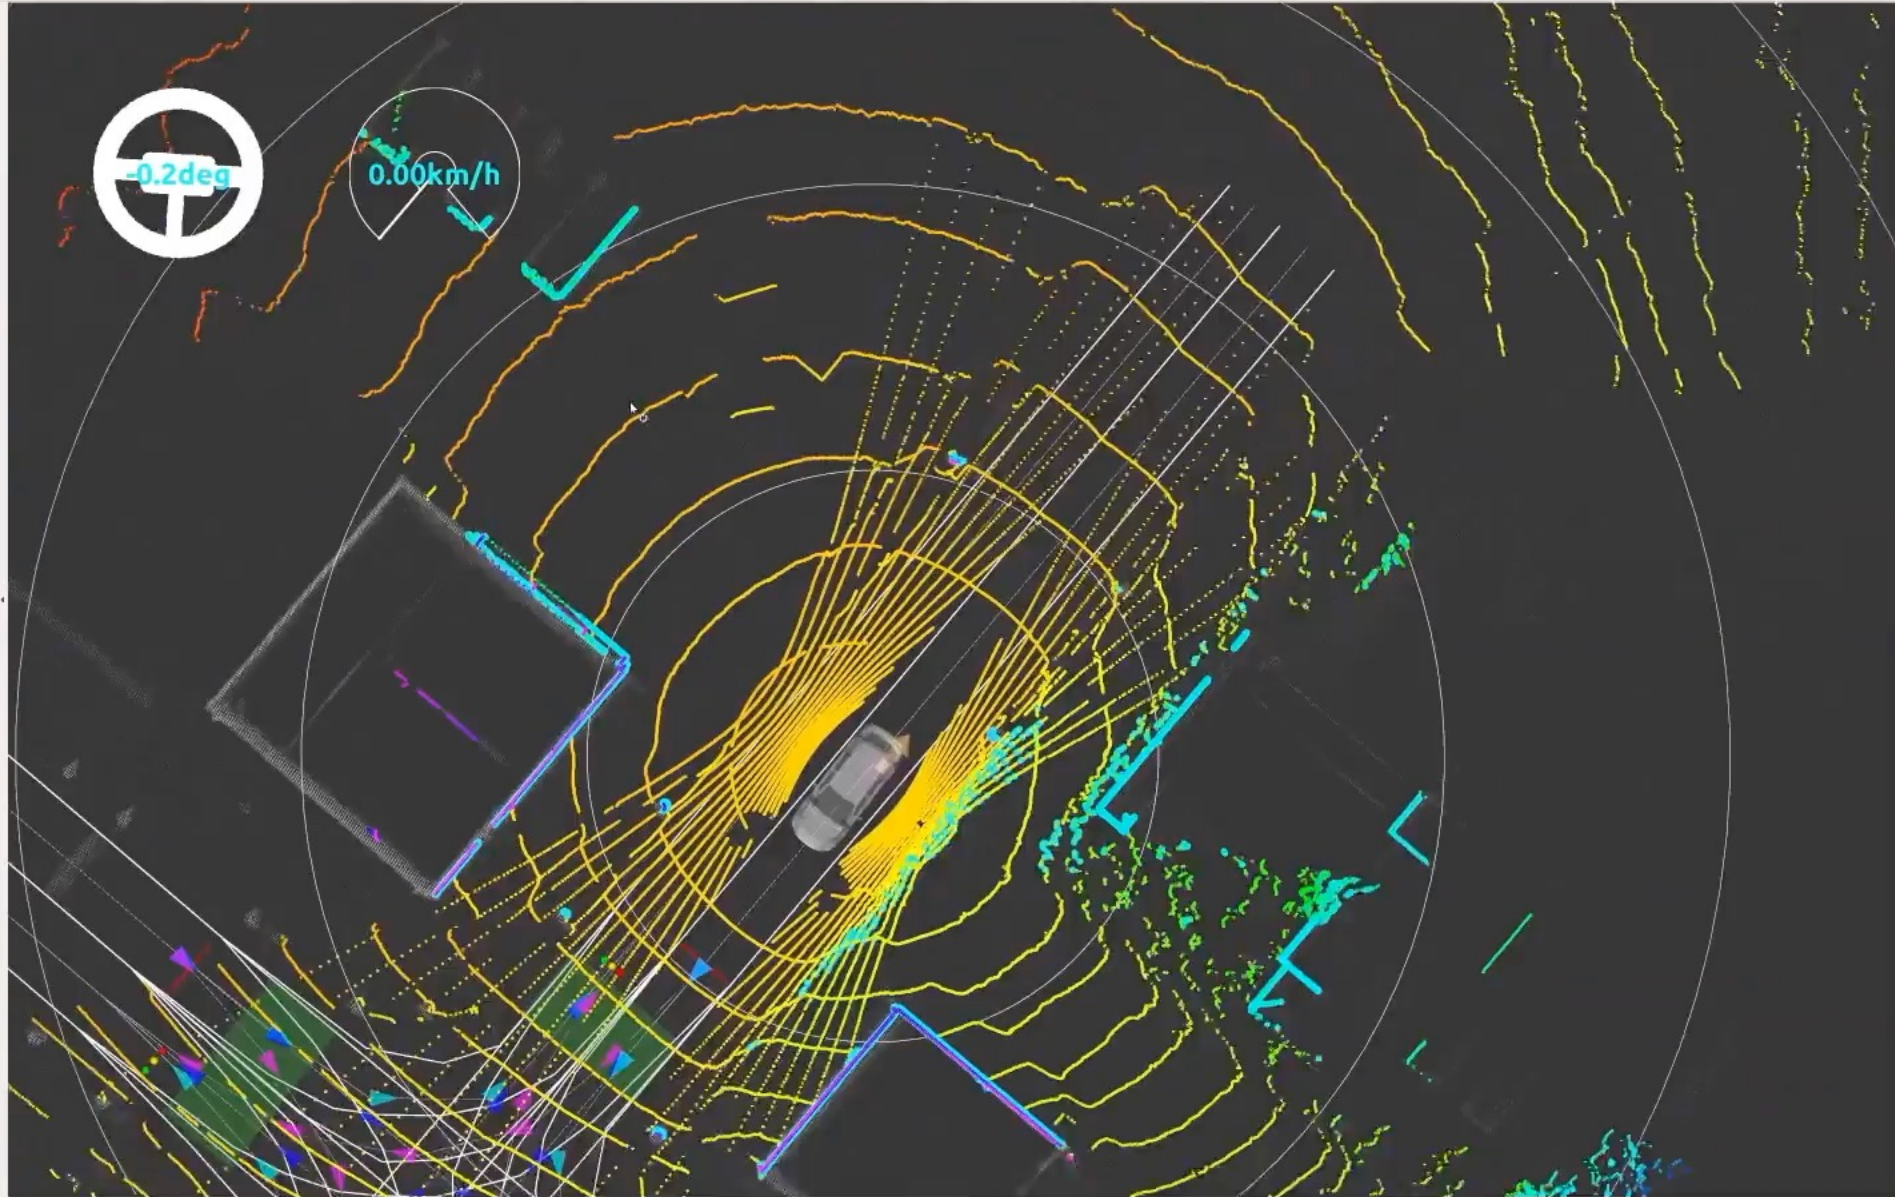
Interact Move Camera Select Focus Camera Measure 2D Pose Estimate 2D Nav Goal 2D Dummy Pedestrian 2D Dummy Car 2D Checkpoint Pose Delete All Objects



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

- Displays
 - Global Options
 - Global Status: Ok
 - System
 - Map
 - Sensing
 - Localization
 - Perception
 - Planning
 - Control
 - Camera



Add Duplicate Remove Rename

Views

Type: TopDownOrtho (rviz) Zero

Current View		TopDownOrtho (rviz)
Near Clip Distance		0.01
Target Frame		base_link
Scale		42.8902
Angle		-0.865
X		5.96944
Y		1.42152
ThirdPersonFollower		ThirdPersonFollower (rviz)

Save Remove Rename

Tool Properties

Interact	
2D Pose Estimate	
Topic	initialpose
X std deviation	0.5

Time

ROS Time: 1582778022.03 ROS Elapsed: 16.83 Wall Time: 1582778022.06 Wall Elapsed: 16.77

Reset Left-Click: Rotate. Middle-Click: Move X/Y. Right-Click: Zoom. Shift: More options.



高速幽灵刹车

TESLA PHANTOM BRAKING



Mobile Manipulation

國1北 268K+410 水上路段

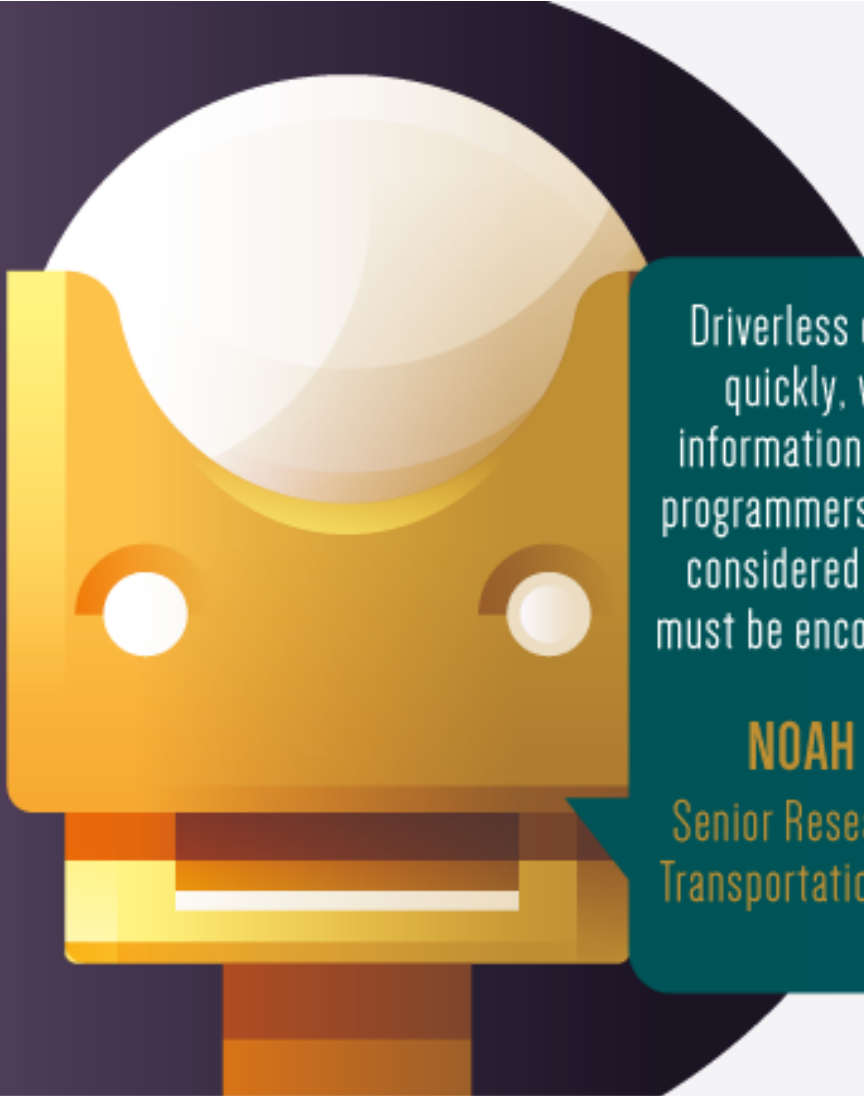
<https://www.youtube.com/watch?v=X3hrKnv0dPQ>

爆料公社



<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?

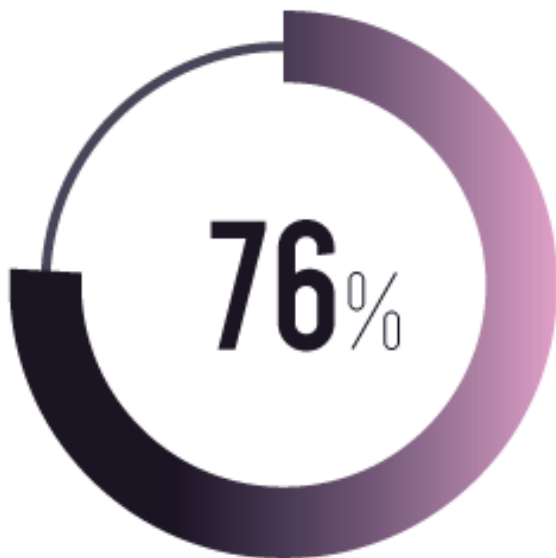
IN A GLOBAL STUDY, MOST PEOPLE PREFERRED

- 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

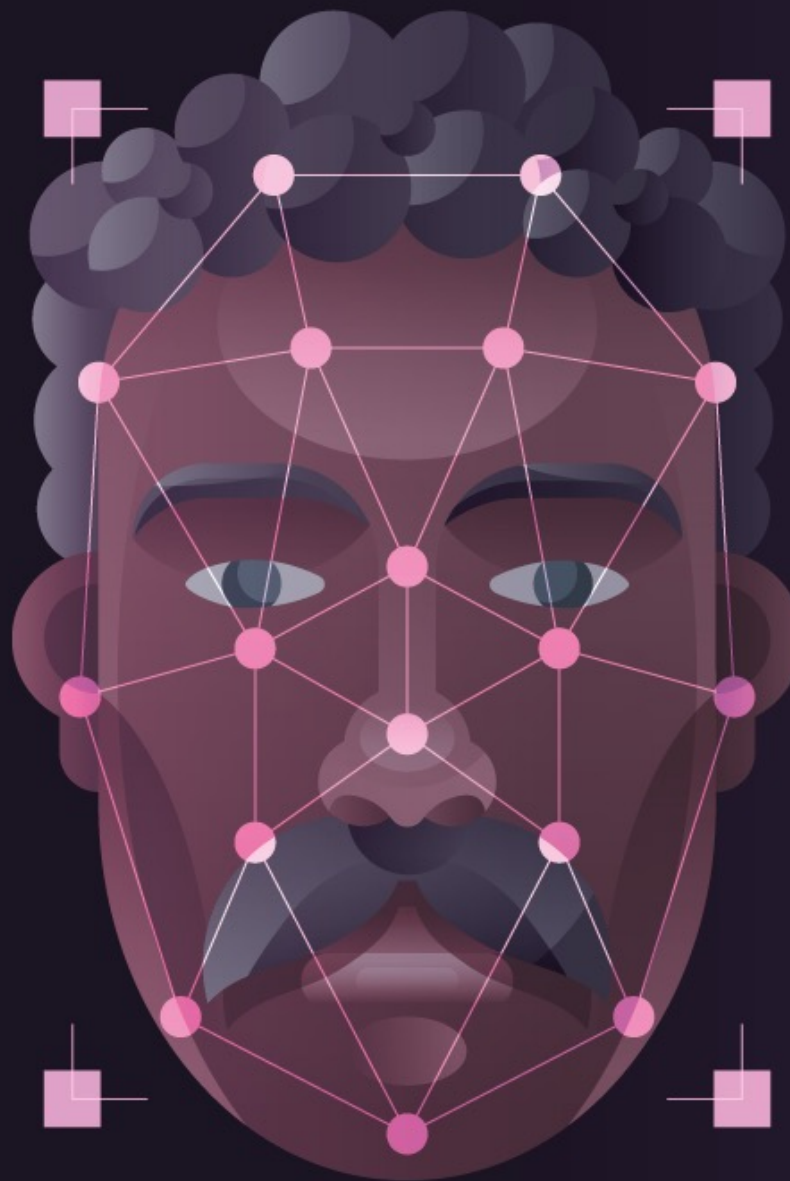


They prefer cars programmed to **PROTECT PASSENGERS AT ALL COSTS**

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

When AI goes wrong

AMAZON'S REKOGNITION



REKOGNITION'S FACIAL RECOGNITION ALGORITHMS CAN

- Identify up to 100 faces in a single image
- Track people in real time through surveillance cameras
- Scan footage from police body cameras

IN 2018, THE ACLU COMPARED 25,000 MUG SHOTS TO PHOTOS OF EVERY MEMBER OF CONGRESS USING REKOGNITION - THEY FOUND

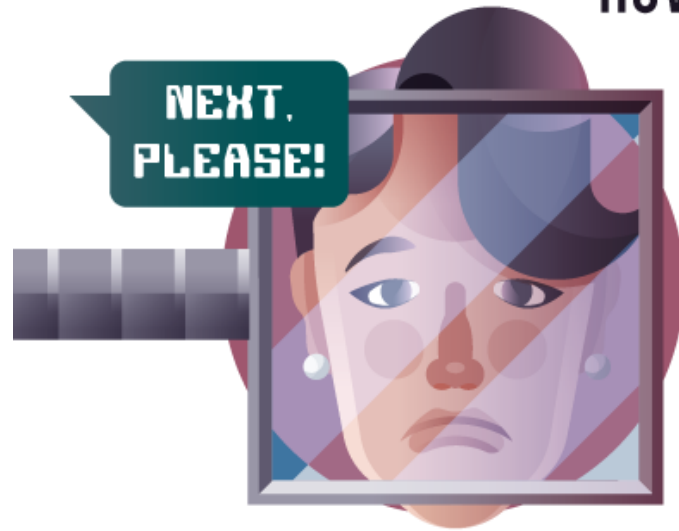
28 False matches

39% Were people of color, who make up just 20% of Congress

LAW ENFORCEMENT AGENCIES ARE ALREADY USING REKOGNITION

- Orlando Police Department (Florida)
- Washington County Sheriff's Office (Oregon)
- In 2016, half of Americans adults were included in a law enforcement facial recognition network

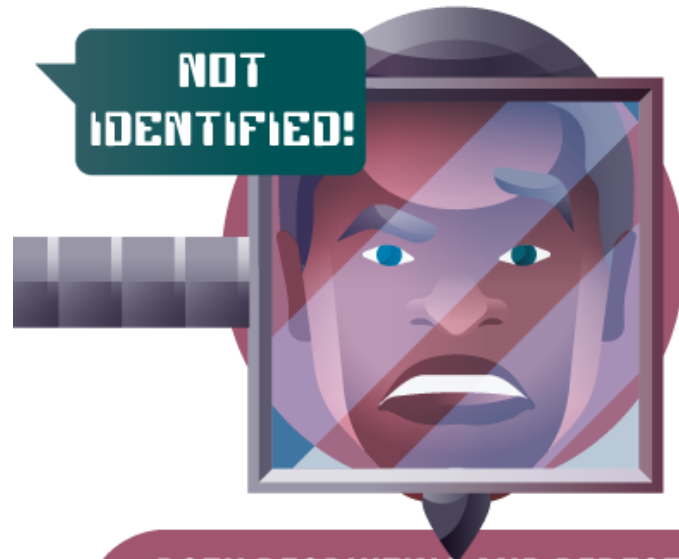
HOW A.I. BIAS HAPPENS



WOMEN NEED NOT APPLY

Starting in 2014, Amazon began training an A.I. to review job candidates

- The system was trained using resumes submitted over 10 years — **MOST CAME FROM MEN**
- The A.I. concluded that “male” was a preferred quality for new hires, and started **FILTERING OUT FEMALE CANDIDATES**



DETECTING DARKER SKIN TONES

In a 2019 study, researchers found that the object detection models used in driverless cars were better at identifying pedestrians with lighter skin

- The study used a standardized set of photos to train their A.I. — but found their **DATASET CONTAINED 3X AS MANY LIGHT SKINNED PEOPLE**
- The A.I. quickly learned to identify light skinned pedestrians, but **STRUGGLED TO IDENTIFY DARKER SKIN TONES**

BOTH RECRUITING AND PEDESTRIAN DETECTION ALGORITHMS FAILED BECAUSE THEY WERE TRAINED ON BAD DATA — A.I. LEARNED BIAS FROM HUMANS

MAKING ETHICAL A.I.

START WITH DATA



A.I. training data must reflect real diversity and control for existing bias

Amazon's recruiting algorithm was trained to eliminate female candidates.
INSTEAD, IT COULD HAVE BEEN PROGRAMMED TO IGNORE GENDER

Pedestrian identification algorithms struggle to identify darker skin tones.
Rather than monitoring success overall, the **TRAINING DATA COULD HAVE WEIGHTED DARK SKIN DATA POINTS MORE HEAVILY**



CONSIDER THE PROCESS

When training A.I., programmers typically split their dataset into 2 parts

Half is used to **TRAIN THE A.I.**

Half is used to **VERIFY AND MEASURE SUCCESS**

If the
initial dataset is
flawed, the test will
have the
same bias



MONITOR FOR UNKNOWNNS

Programmers must monitor for unintentional bias appearing in their A.I.

Subtle patterns can lead A.I. to **PERPETUATE HUMAN BIAS**

Amazon's recruiting algorithm preferred **VERBS LIKE "EXECUTED" AND "CAPTURED" — WHICH TEND TO BE MORE USED BY MALES**

8 Ethical Questions in AI



Bias:

Is AI fair?



Liability:

Who is responsible for AI?



Security:

How do we protect access to AI from bad actors?



Human Interaction:

Will we stop talking to one another?



Employment:

Is AI getting rid of jobs?



Wealth Inequality:

Who benefits from AI?



Power & Control:

Who decides how to deploy AI?



Robot Rights:

Can AI suffer?

<https://www.logikk.com/articles/8-ethical-questions-in-artificial-intelligence/>

Ethical AI: Many open questions and topics:

- **Autonomy and liability**
- **Ethical principles in robotics**
- **Enhancement technologies:** ethical issues
- Defining ethical guidelines for the design, use, and operation of robots
- Privacy & management of personal data
- **Ethical frameworks: universal or region-specific?**
- The role of industry and society in the definition of safety standards
- AI technology to block unethical/mendacious social-media communication
- **Accountability** in autonomous systems
- **Embedding values** and norms into intelligent systems
- Ethics and **standardization**
- Raising ethical awareness among stakeholders
- **Transparency** in autonomous systems
- Political and **legal** frameworks
- **Formal and mathematical** frameworks for robot ethics
- Implementations and engineering studies
- User and HCI/HRI studies at the intersection of the above issues

<https://clawar.org/icres2023/>

Ethical AI: Scientific Discussion

- ICRES 2023 is the 8th edition of the International Conference on Robot Ethics and Standards series
 - <https://clawar.org/icres2023/>
- IEEE Robotics and Automation Society: Technical Committee on Robot Ethics
 - Framework for raising and addressing the urgent ethical questions prompted by and associated with robotics research and technology.
 - <https://www.ieee-ras.org/robot-ethics>
- Conference on Robotics, AI and Humanity, Science, Ethics and Policy organized jointly by the Pontifical Academy of Sciences (PAS) and the Pontifical Academy of Social Sciences (PASS)
 - <http://www.pas.va/content/accademia/en/events/2019/robotics/statementrobotics.html>

Autonomous Weapons are attractive to the military

- Don't endanger the life of your own soldiers.
- Robots don't get tired – they can operate 24/7.
- Superior situational awareness through 360-degree view with advanced sensors & instantaneous communication with various other units in the field
- Stay highly vigilant the whole day
- No fear: perform the most dangerous attacks without degraded performance
- Ease logistics: no need for food, shelter, medical attention or rotation of units.
- Superior performance in battle, e.g. speed of target acquisition, firing accuracy, battle coordination, overall strategy



THeMIS robot from the Estonia company MILREM Robotics.
The 2t semi-autonomous robot can drive with a speed of up to 20km/h.

Ukraine's Future Vision and Current Capabilities for Waging AI-Enabled Autonomous Warfare



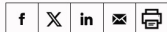
Photo: Anton Petrus/GETTY IMAGES

Table of Contents

• Executive Summary

Introduction

The Ukrainian Military's Strategic Vision and Technological Road Map



Report by **Kateryna Bondar**
Published March 6, 2025

↓ Available Downloads



Kateryna Bondar
Fellow, Wadhvani AI Center

<https://www.csis.org/analysis/ukraines-future-vision-and-current-capabilities-waging-ai-enabled-autonomous-warfare>

March 2025

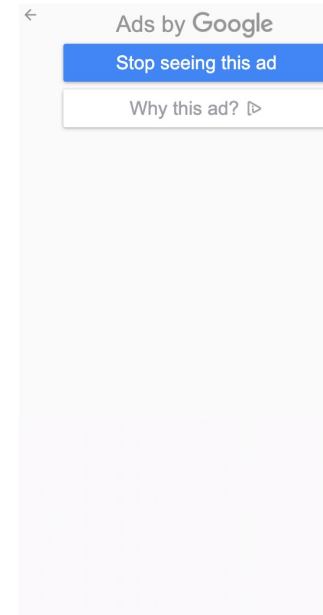
Key Findings:

- The Ukrainian military's objective is to remove warfighters from direct combat and replace them with autonomous unmanned systems.
- Autonomy is not yet present on the battlefield in the war in Ukraine.
- The current deployment of AI is partial in scope, enhancing certain functions and addressing some operational challenges rather than enabling full system autonomy.
- Ukrainian forces have widely adopted small and medium first-person-view (FPV) drones as platforms that may be quickly adapted for diverse missions through modular design and interchangeable equipment.
- Ukraine's defense industry is developing standalone AI-driven software that can be integrated across various platforms to expand battlefield autonomy.
- Delegating target recognition to AI-enabled automatic target recognition (ATR) systems onboard unmanned platforms reduces human limitations and allows locking on to targets up to 2 km away.
- Autonomous navigation makes drones strikes three to four times more likely to succeed.
- Human oversight remains pivotal—particularly for engagement decisions—reflecting a human-in-the-loop approach that could shift toward higher-level supervision in the future while still maintaining human control of the system.

WORLD > EUROPE • 6 MIN READ

'Robots don't bleed': Ukraine sends machines into the battlefield in place of human soldiers

APR 20, 2026

By [Ivana Kottasová](#), [Daria Tarasova-Markina](#), [Victoria Butenko](#)

- Ukraine claims to have taken Russian position just using ground robots in summer 2025
- “first time in history that an enemy position was stormed and prisoners taken by ground robots and drones without the involvement of infantry”
- Maybe propaganda or exaggeration – but the trend is clear...

THANKS FOR YOUR INTEREST IN
ROBOTICS 😊
