iGibson

Simulation for Visual Embodied Agents

Presenter: Fei Xia

The Gibson Team

Fei Xia, William B. Shen, Chengshu Li, Micael Tchapmi, Lyne Tchapmi, Noriaki Hirose

Kevin Chen, JunYoung Gwak, Yongpu Jin, Alexander Toshev, Amir R. Zamir, Roberto Martín-Martín

Li Fei-Fei, Silvio Savarese

A glimpse into Gibson

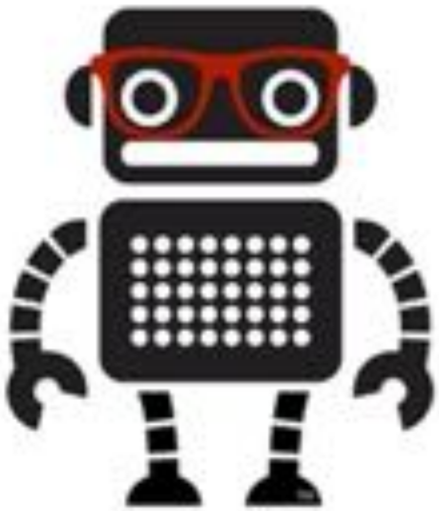

Gibson is a simulation environment
to train visually controlled robotic agents
in navigation and manipulation tasks

A glimpse into Gibson

3D reconstructed real environments

A glimpse into Gibson

3D reconstructed real environments

A glimpse into Gibson

3D reconstructed real environments

A glimpse into Gibson

3D environments

realistic virtual images

A glimpse into Gibson

3D environments

realistic virtual images

A glimpse into Gibson

3D environments

realistic images

to train embodied agents using vision

A glimpse into Gibson

3D environments

to train robots

realistic images

in navigation and manipulation tasks

A glimpse into Gibson

3D environments

to train robots

realistic images

in navigation and manipulation tasks

A glimpse into Gibson

3D environments

to train robots

realistic images

in navigation and manipulation

Gibson is a simulation environment...

with hundreds of 3D reconstructed real models,
that generates realistic virtual images
to train embodied agents
in navigation and manipulation tasks

In this presentation…

1. Why do we need simulators in robotics?

2. Requirements for a good robotics simulator

3. Gibson V1

4. Interactive Gibson

Outline

Why do we need simulators in robotics?

Requirements for a good robotics simulator

Gibson v1

Interactive Gibson


Robots break…


IRPLEX
FAIRPLEX
FAIRPLEX
6:16:34 05/06/2015


Robots are slow….


20x
POLARIS
01:02:09 06/06/2015


3

Reinforcement learning

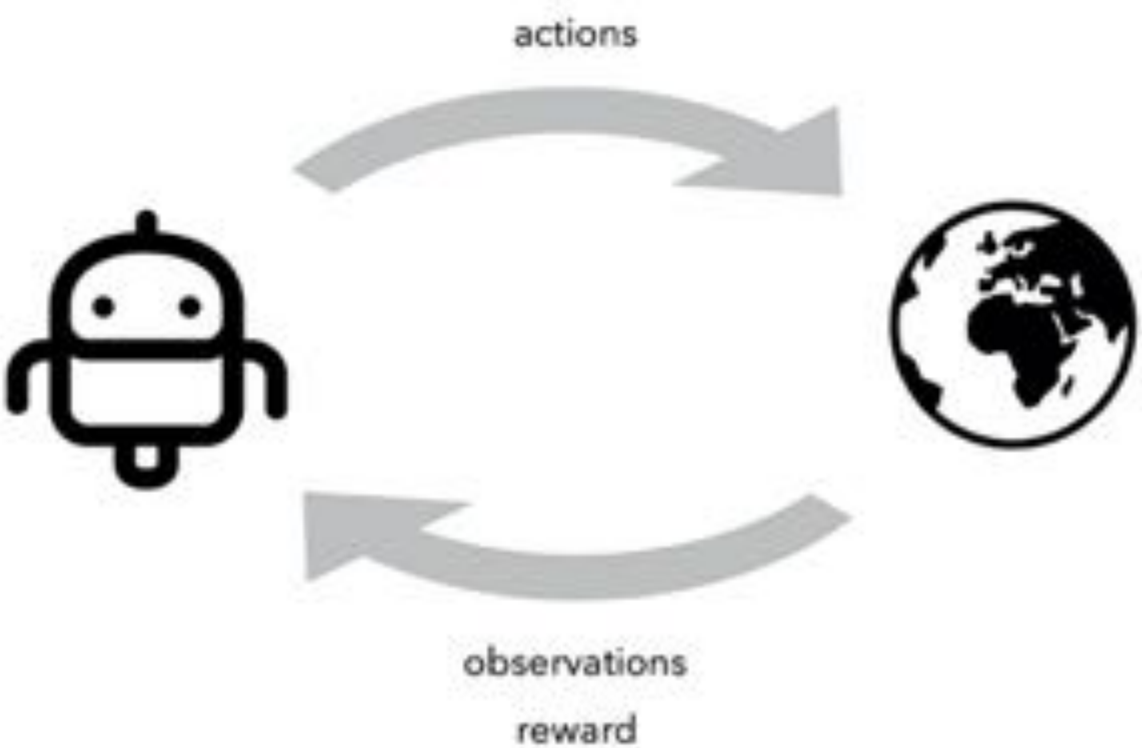

- Robot learns from own experiences
- Current visual RL algorithms require MANY (~1e6) interactions

How to train your robot?

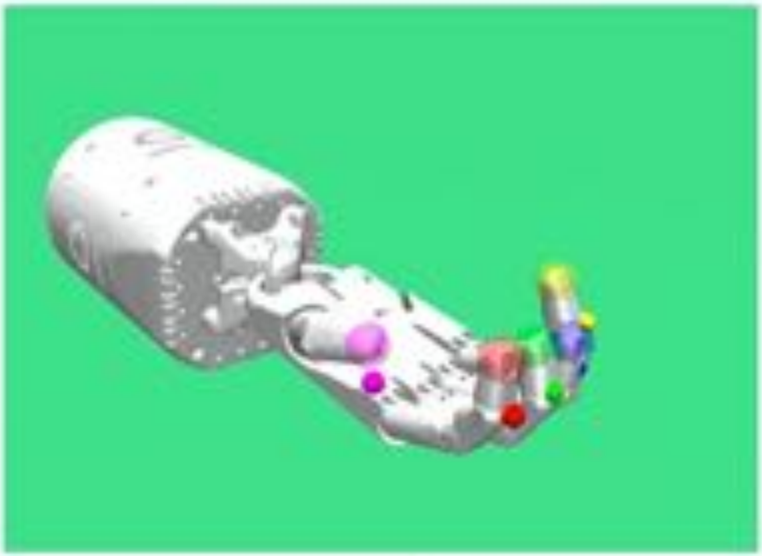

Train in simulation

Deploy in real world

Images by OpenAI

100 years of experiences in just few days!!!

[Learning dexterous in-hand manipulation, Andrychowicz et al., IJRR20]

Requirements for a robotics simulator

Fast

Realistic
Physics

Realistic
Virtual Sensor
Signals

Realistic
Semantic
Distribution

Outline

- Why do we need simulators in robotics?
- Requirements for a good robotics simulator
- Gibson v1
- Interactive Gibson

Gibson V1

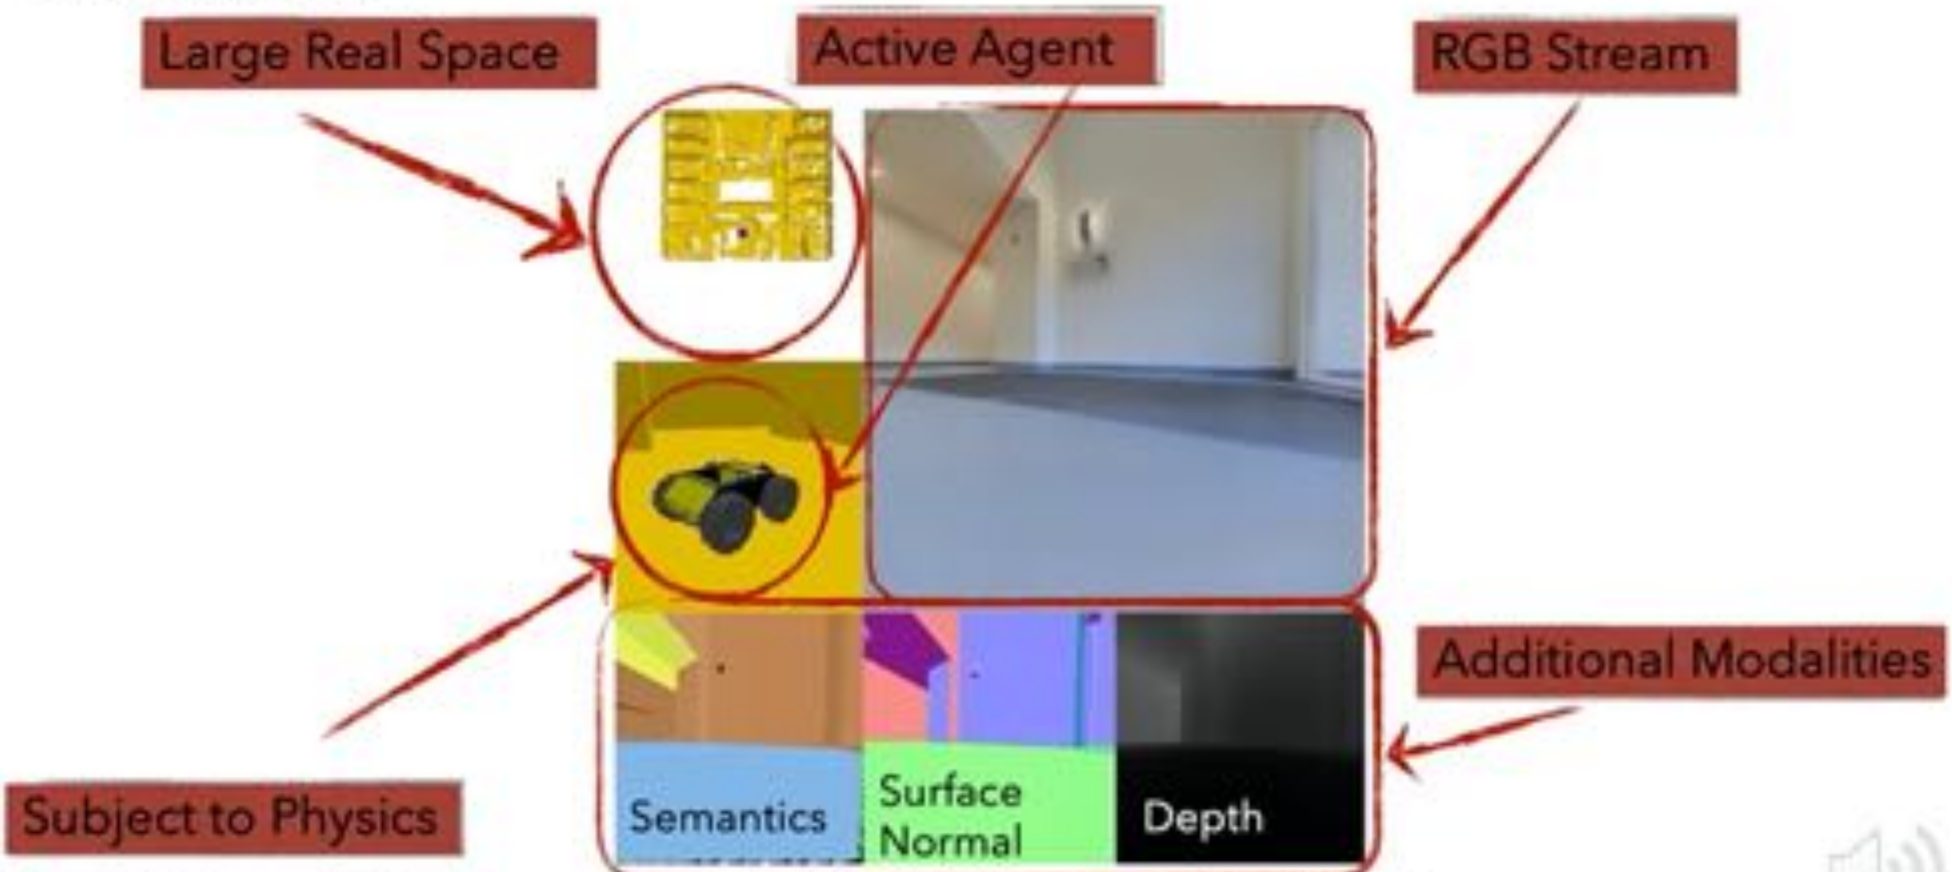

Xia, Fei, et al. "Gibson env: Real-world perception for embodied agents." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018.

Gibson V1 - Assets

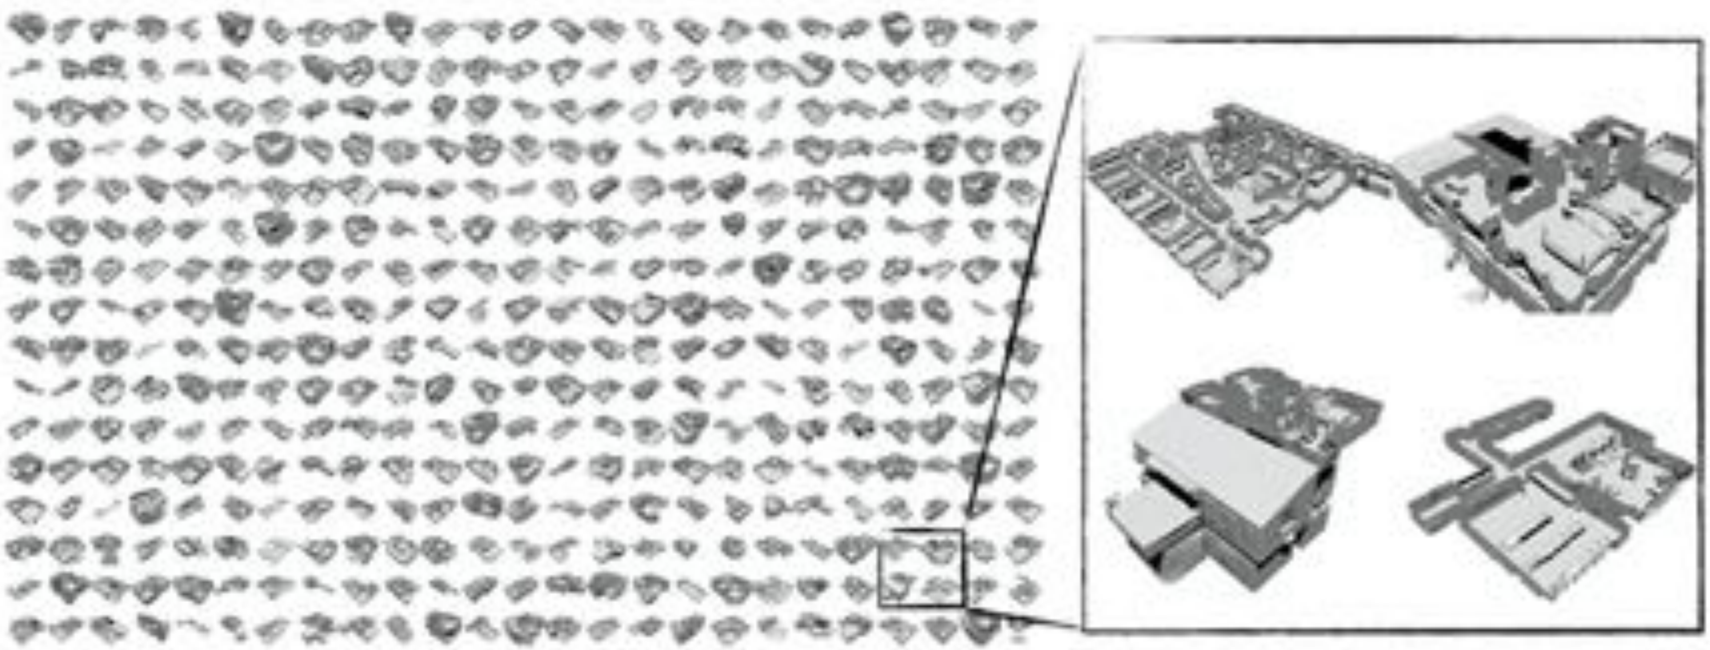

572 full buildings. Real spaces, scanned with 3D scanners.
211,000 m2. 1400+ floors.
Browse data at: http://gibson.vision/database/

Gibson V1 - Assets

572 full buildings. Real spaces, scanned with 3D scanners.
211,000 m2. 1400+ floors.
Browse data at: http://gibson.vision/database/

Gibson V1 - Assets

Multiple RGBD Images per Building

High Quality Reconstruction

Gibson V1 – Active Agents

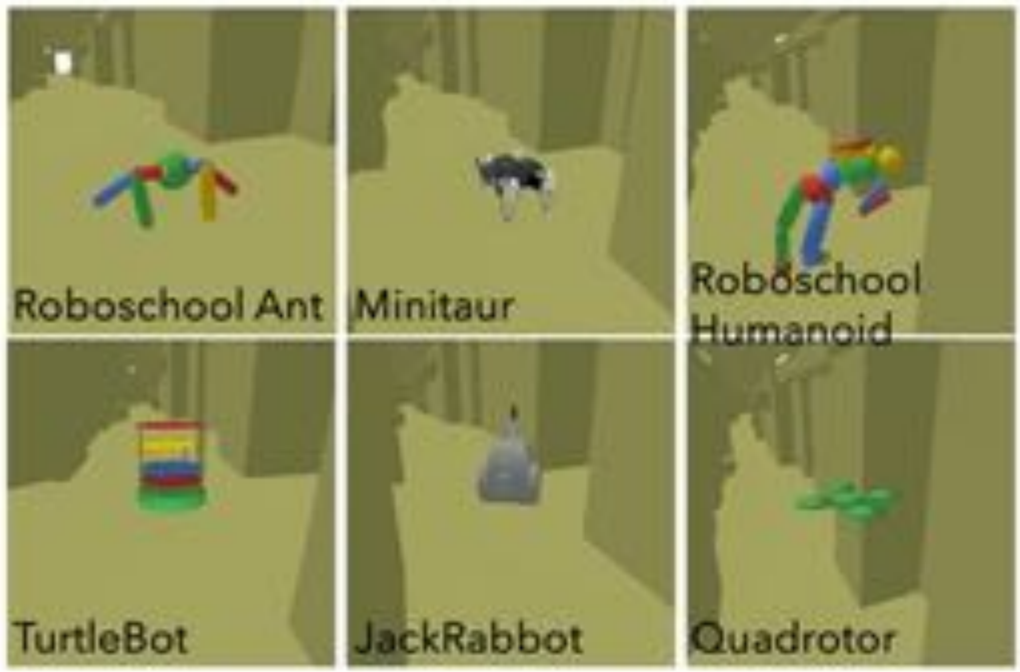

Arbitrary agents can be improved using their URDF

Gibson V1 – Agents are subject to Physics

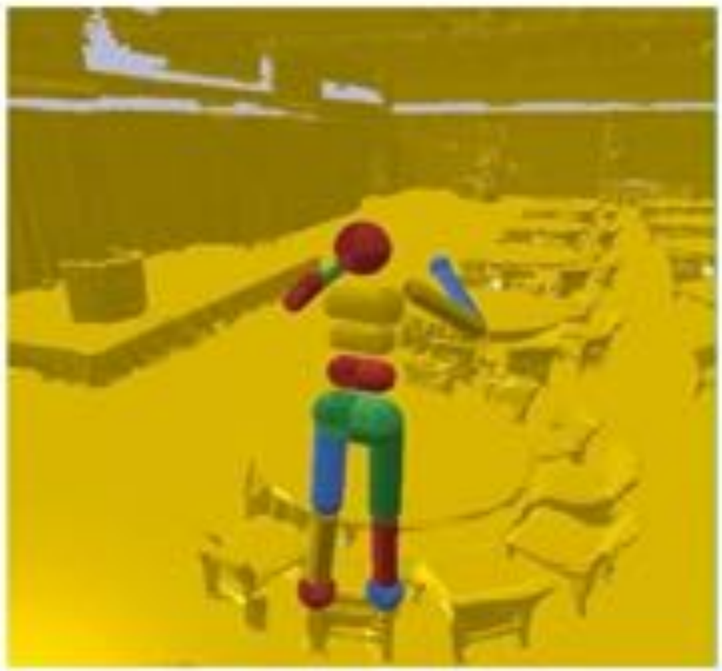

Integrated with physics engine, PyBullet3D [Coumans2016]

Gibson V1 – RGB Rendering

Neural network based view synthesis engine

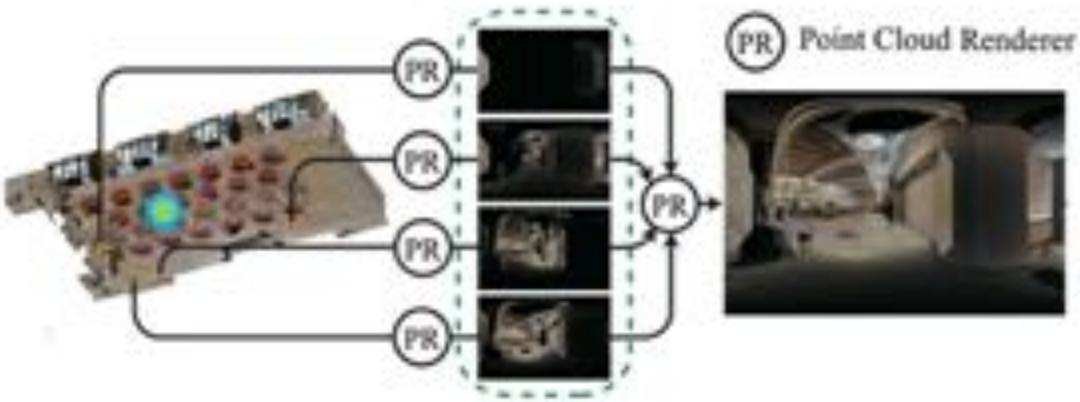

Gibson V1 – RGB Rendering

Neural network based view synthesis engine

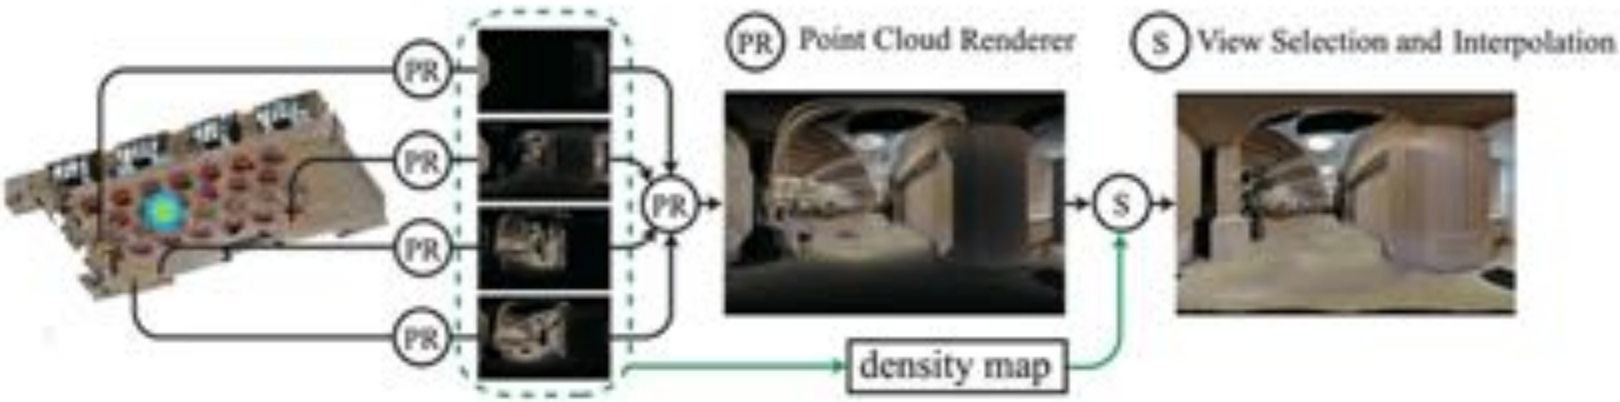

Gibson V1 – RGB Rendering

Neural network based view synthesis engine

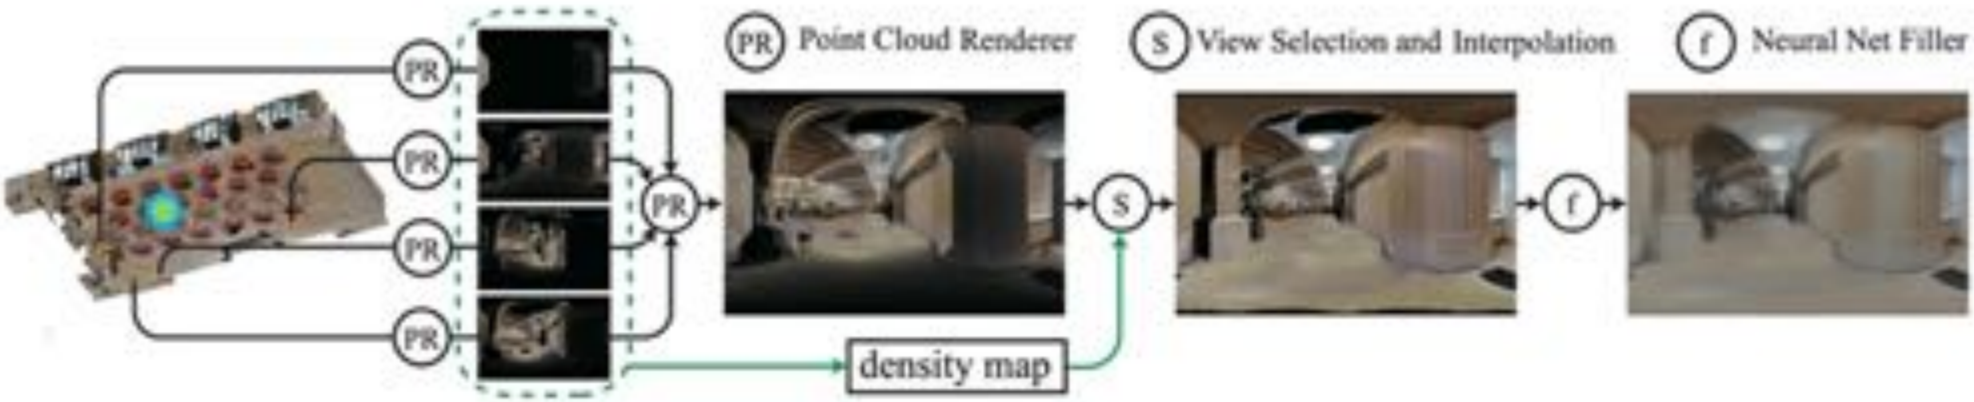

Gibson V1 – RGB Rendering

Neural network based view synthesis engine

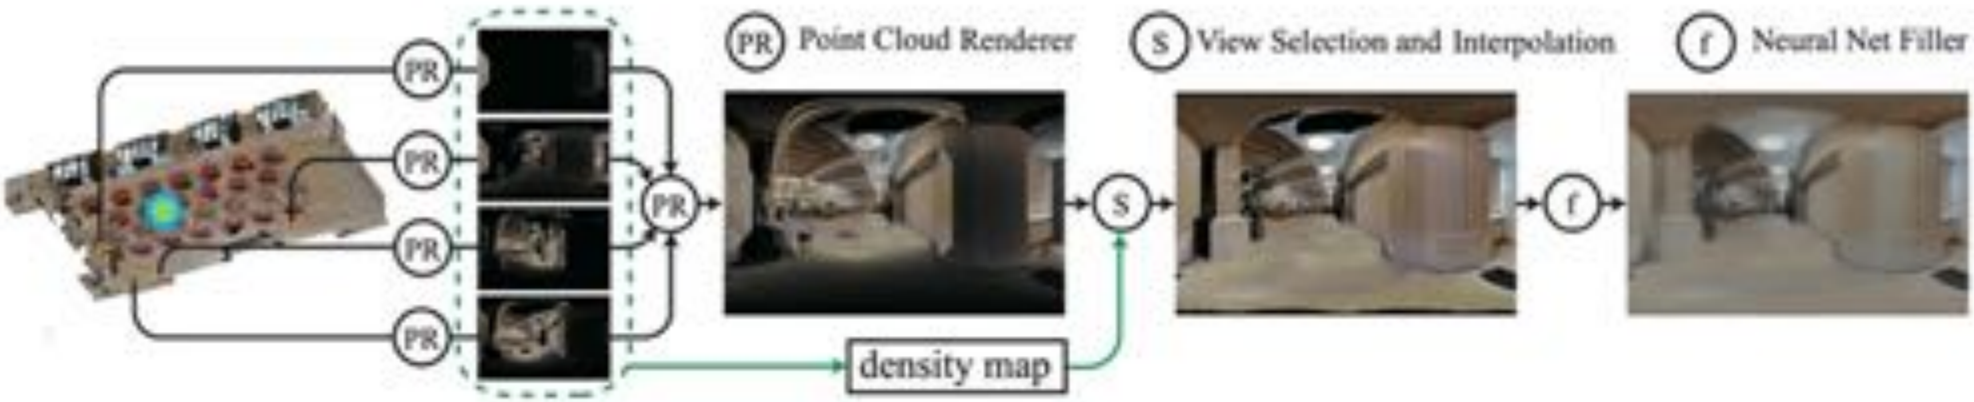

Given sparse RGB-D images, renders the scene from arbitrary viewpoints
Faster than real-time
Neural Net filler to correct "holes"

Gibson V1 – RGB Rendering Examples

Gibson V1 – RGB Transfer → Goggles!

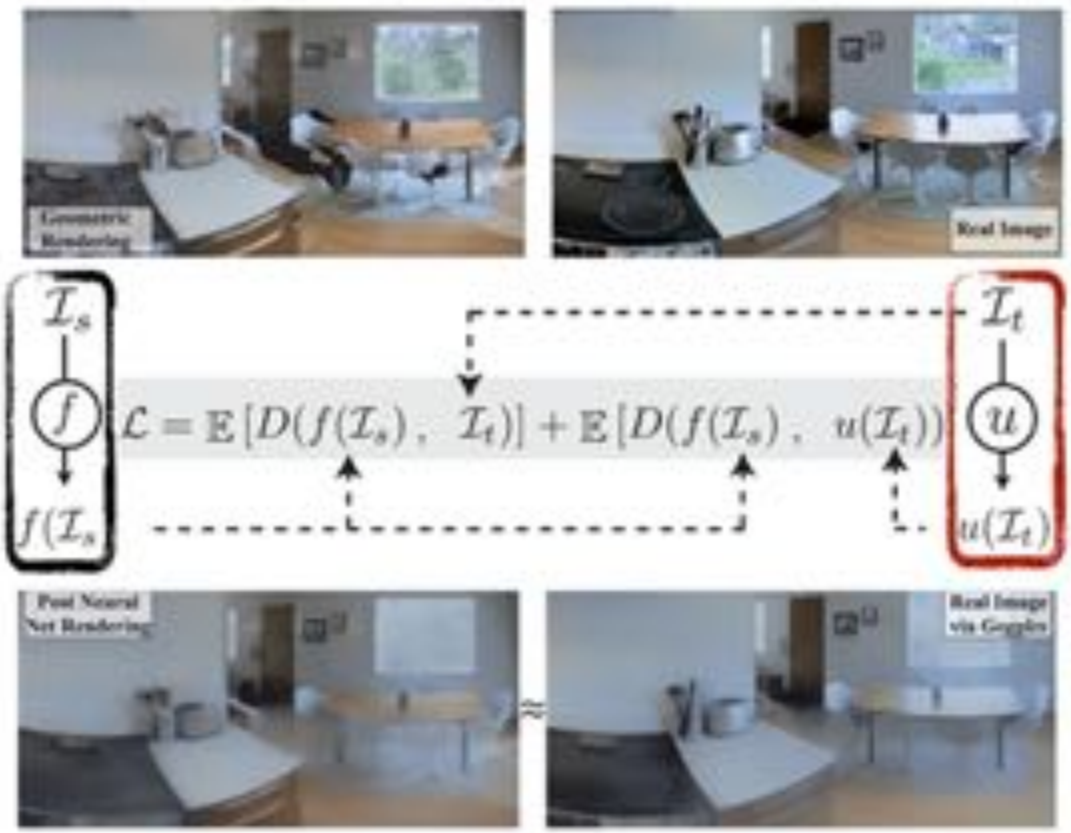

Gibson V1 – Examples of Trained Visual Agents

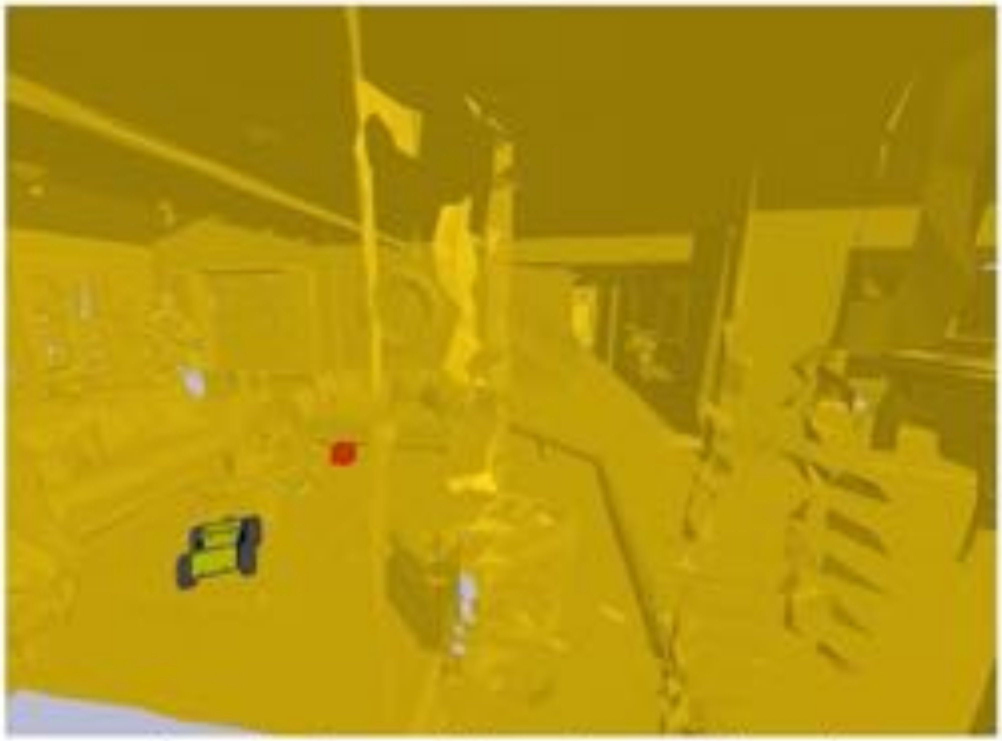

Local planner (obstacle avoidance)

Gibson V1 – Examples of Trained Visual Agents

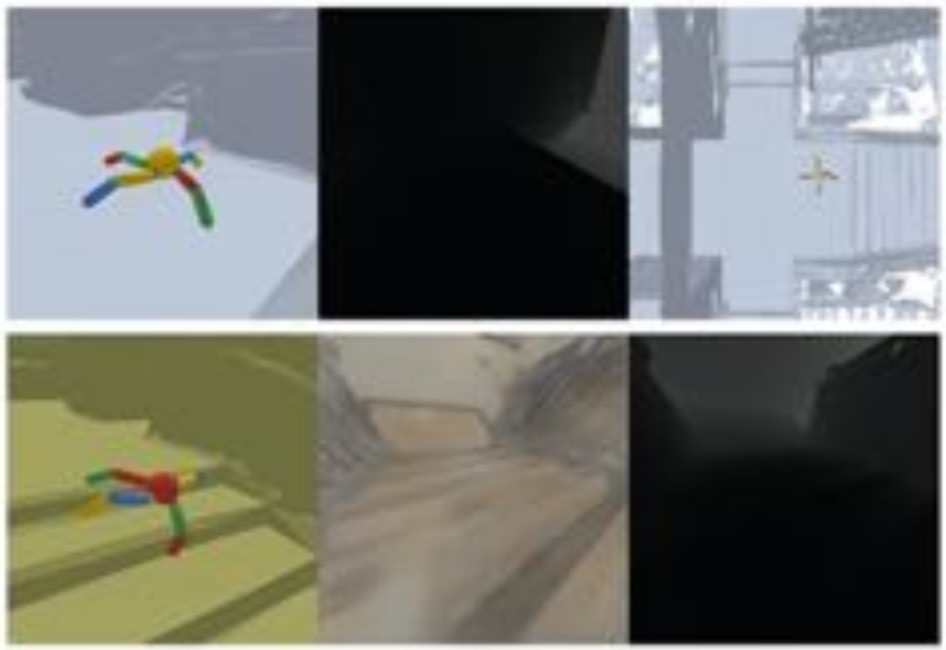

Stair climb

Gibson V1 – Examples of Trained Visual Agents

Distant navigation

What did we achieve with Gibson V1

- A perceptual environment
- Real spaces
- Neural network synthesis. "Goggles".
- +physics, +OpenAI GYM, +ROS

What can be done in Gibson V1?

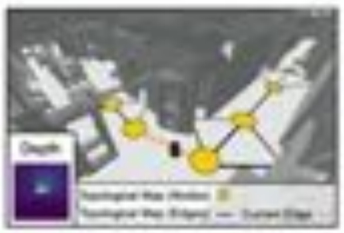

[A behavioral approach to visual navigation with graph localization networks, Chen et al., RSS19]

[Neural Autonomous Navigation with Riemannian Motion Policy, Meng et al., ICRA19]

[Mid-Level Visual Representations Improve Generalization and Sample Efficiency for Learning Visuomotor Policies, Sax et al., 2018]

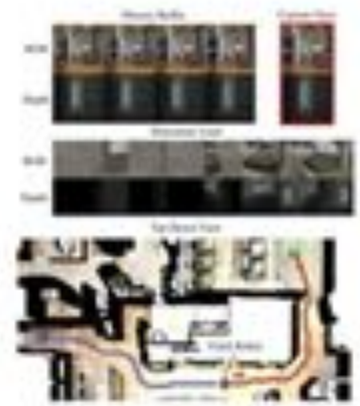

[Learning Your Way Without Map or Compass: Panoramic Target Driven Visual Navigation, Watkins-Valls et al., 2019]

[Scaling Local Control to Large-Scale Topological Navigation, Meng et al., 2019]

[Generalization through Simulation: Integrating Simulated and Real Data into Deep Reinforcement Learning for Vision-Based Autonomous Flight, Kang et al., ICRA19]

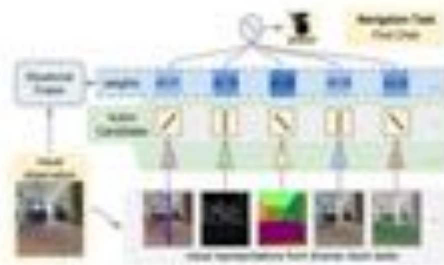

[Situational Fusion of Visual Representation for Visual Navigation, Shen et al., CVPR19]

Controlling a Simulated Robot with Brain Signals

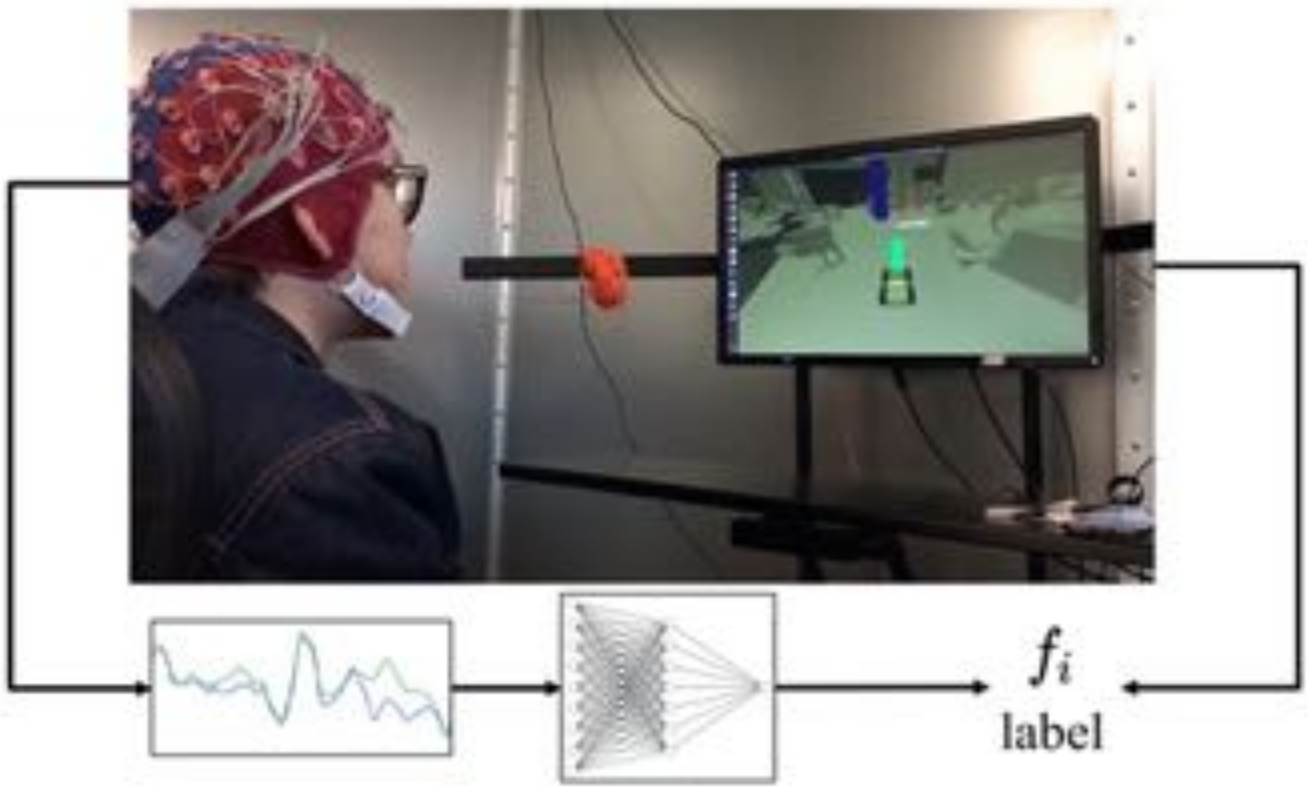

[Accelerated Robot Learning via Human Brain Signals, Akinola et al., 2019]

Learning to Follow Visual Trajectories

[Deep Visual MPC-Policy Learning for Navigation, Hirose et al., RAL2019]

Learning to Follow Visual Trajectories


Visual Trajectory

[Deep Visual MPC-Policy Learning for Navigation, Hirose et al., RAL2019]

Learning to Follow Visual Trajectories

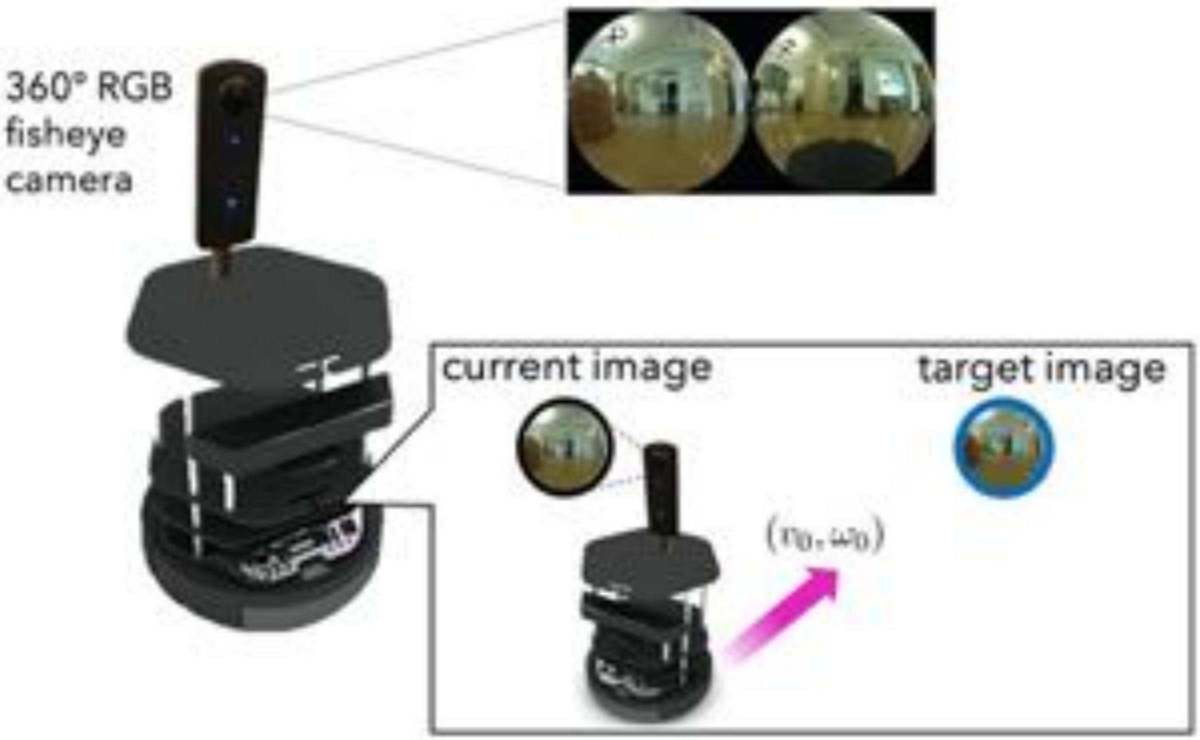

Visual Trajectory

MPC Policy

[Deep Visual MPC-Policy Learning for Navigation, Hirose et al., RAL2019]

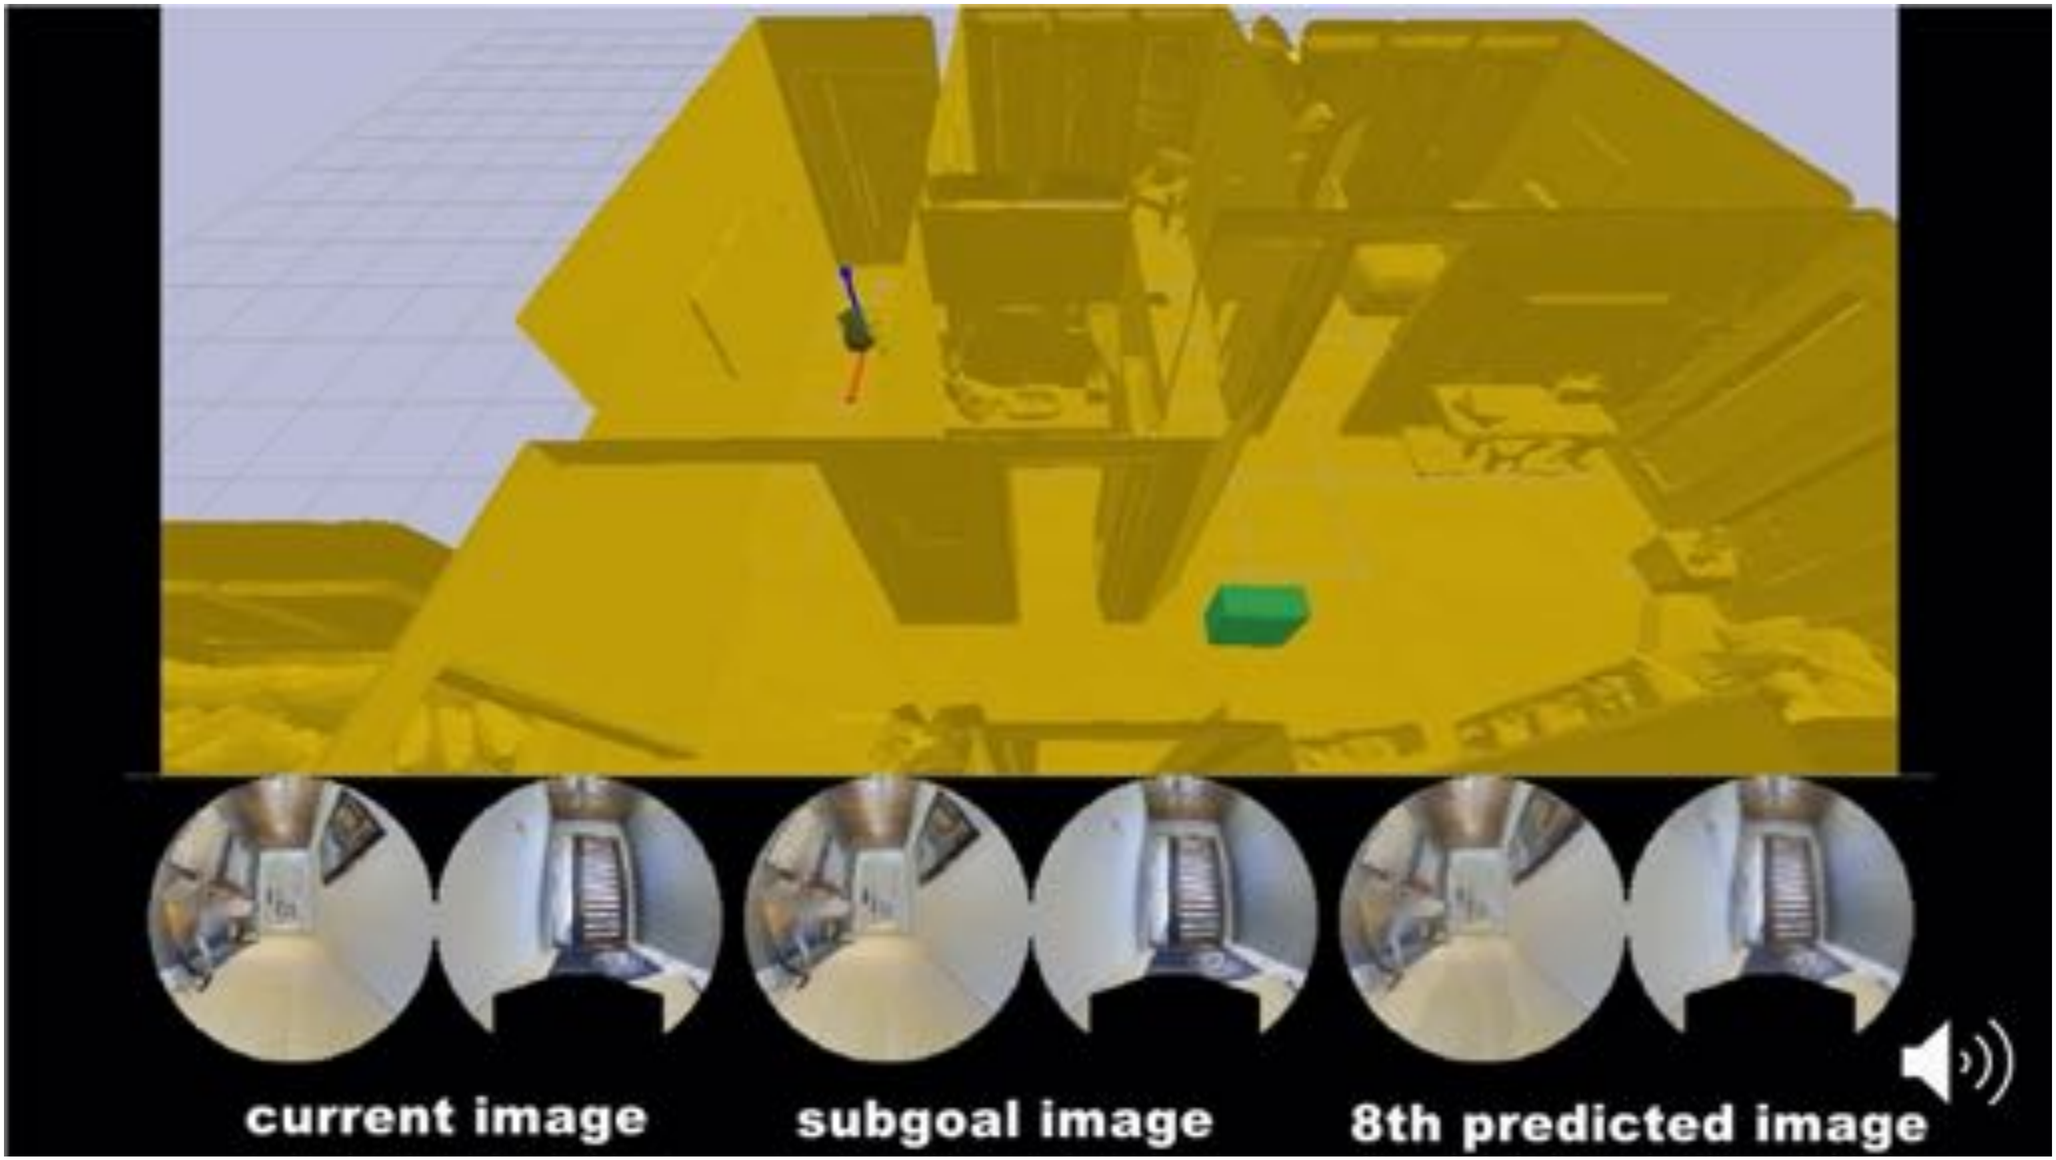
current image
subgoal image
8th predicted image

Limitations of Gibson v1: Rendering Speed

128x128
@ 73fps

256x256
@ 42fps

512x512
@ 18fps

Limitations of Gibson v1: Interactivity

Requirements for a robotics simulator

Fast

Realistic Physics

Realistic Virtual Sensor Signals

Realistic Semantic Distribution

Outline

Why do we need simulators in robotics?

Requirements for a good robotics simulator

Gibson v1

Interactive Gibson

Interactive Gibson

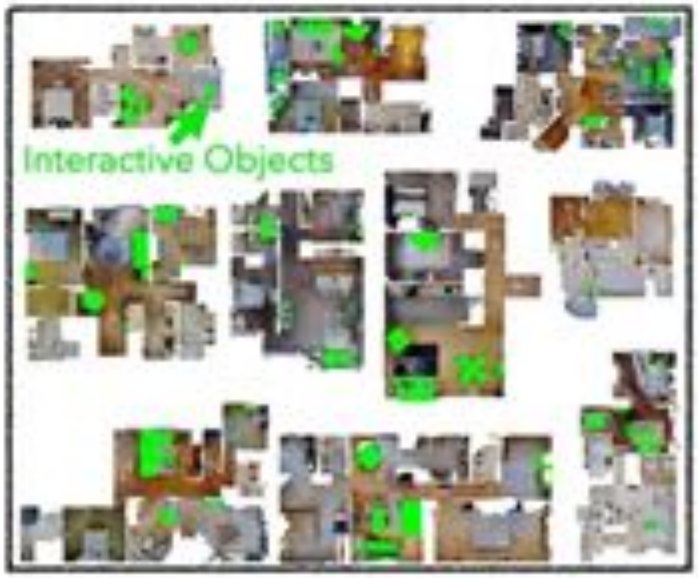

A new version of Gibson with:

- hundreds of real-world reconfigurable scenes
- hundreds of interactive objects
- faster rendering

From Gibson to Interactive Gibson

- Interactive parts of the environment have to be selected

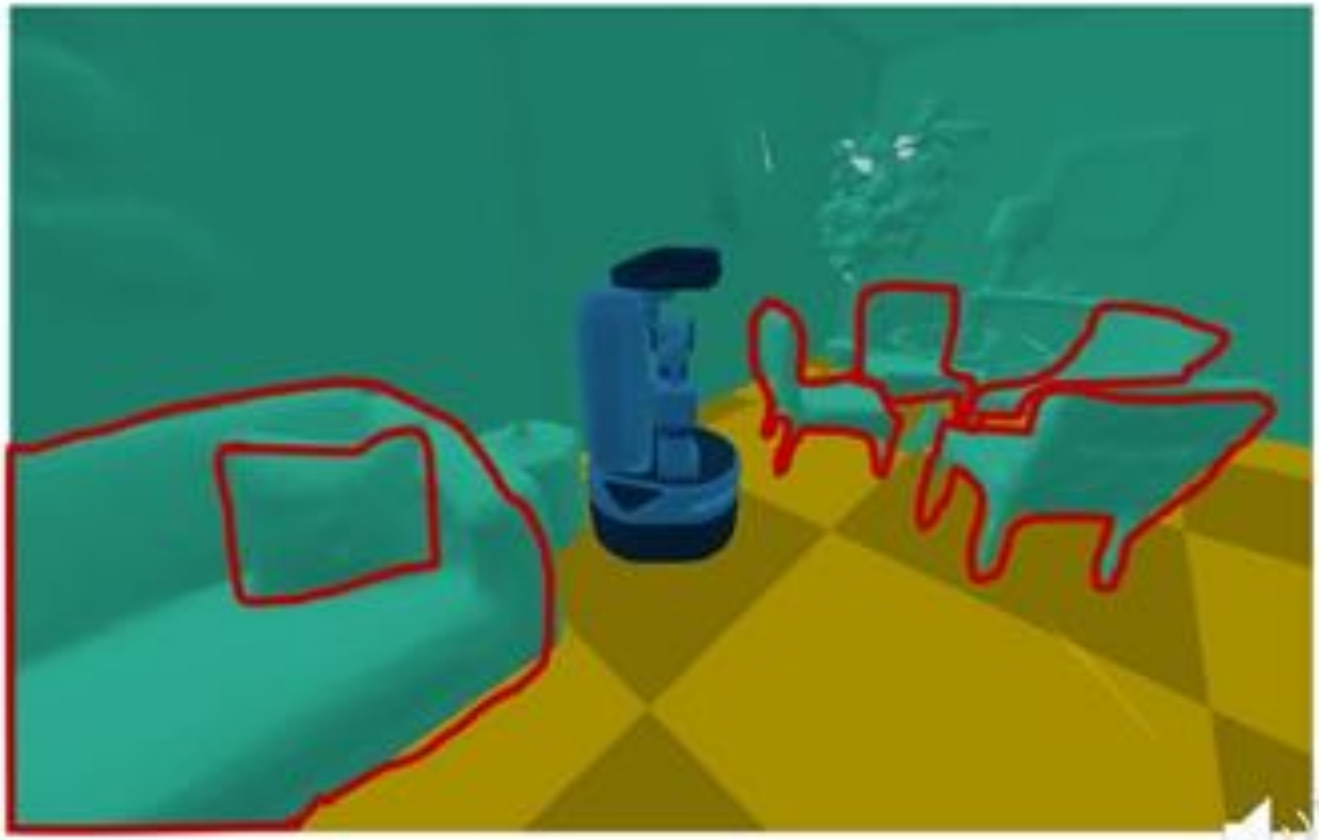

From Gibson to Interactive Gibson

- Interactive parts of the environment have to be selected

- They have to be segmented

From Gibson to Interactive Gibson

- Interactive parts of the environment have to be selected
- They have to be segmented
- Then replaced by interactive models

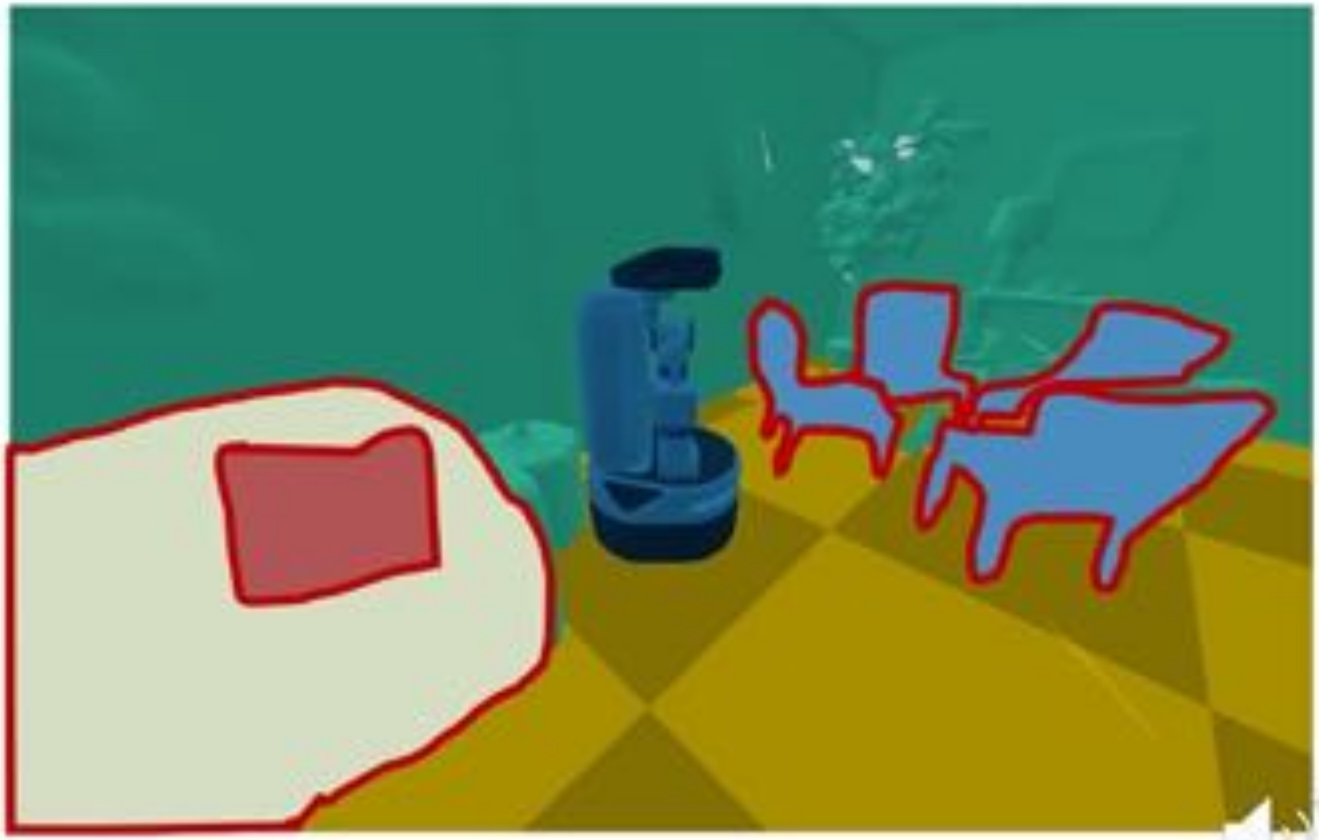

From Gibson to Interactive Gibson

From Gibson to Interactive Gibson

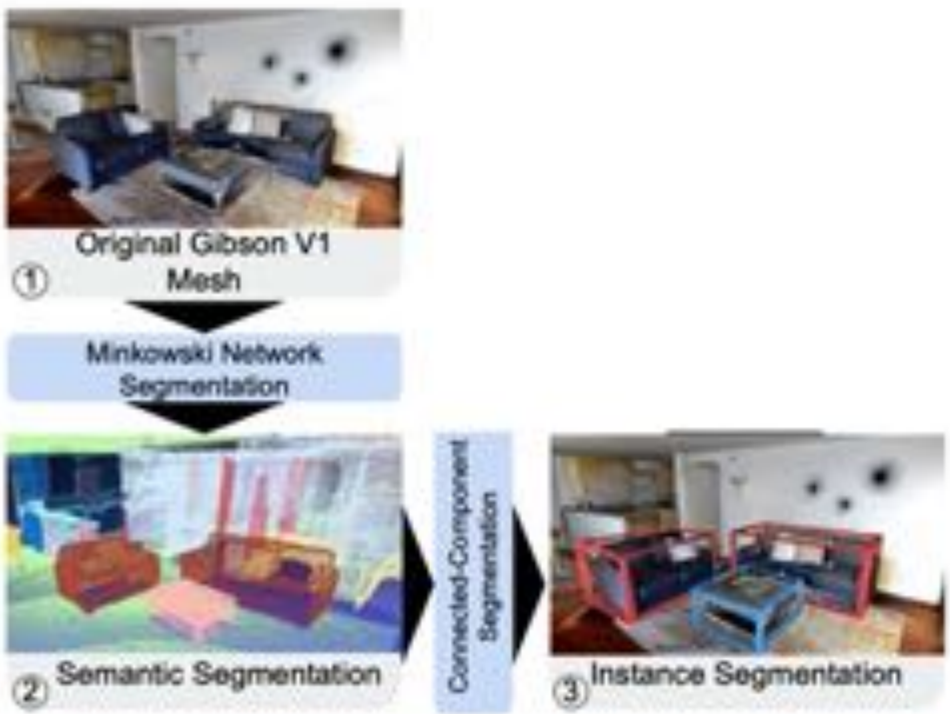

From Gibson to Interactive Gibson

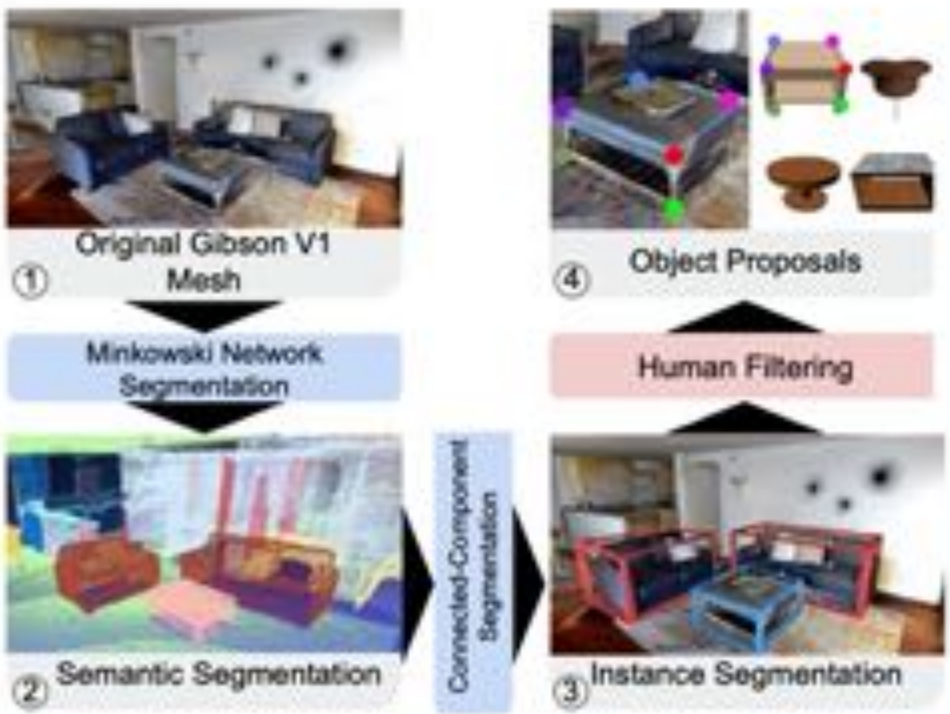

From Gibson to Interactive Gibson

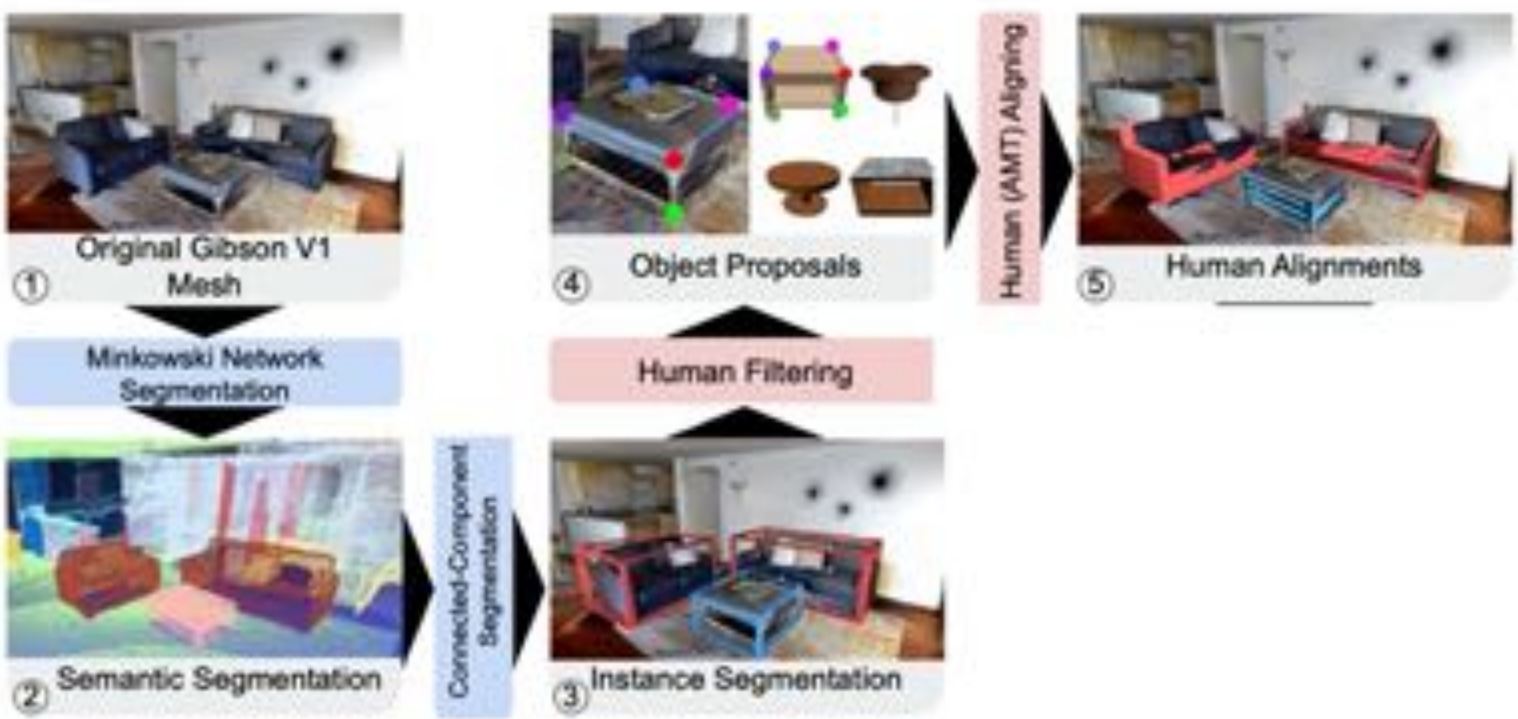

From Gibson to Interactive Gibson

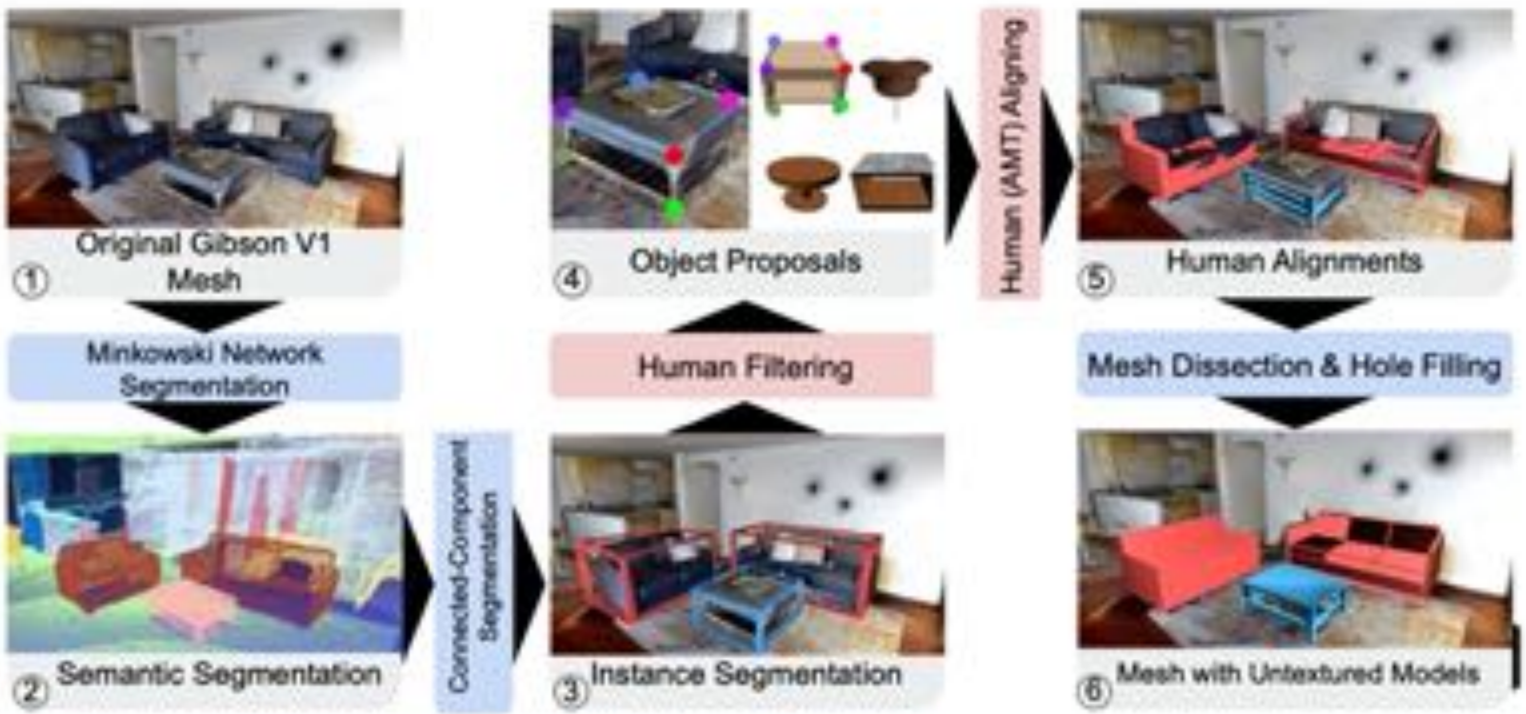

From Gibson to Interactive Gibson

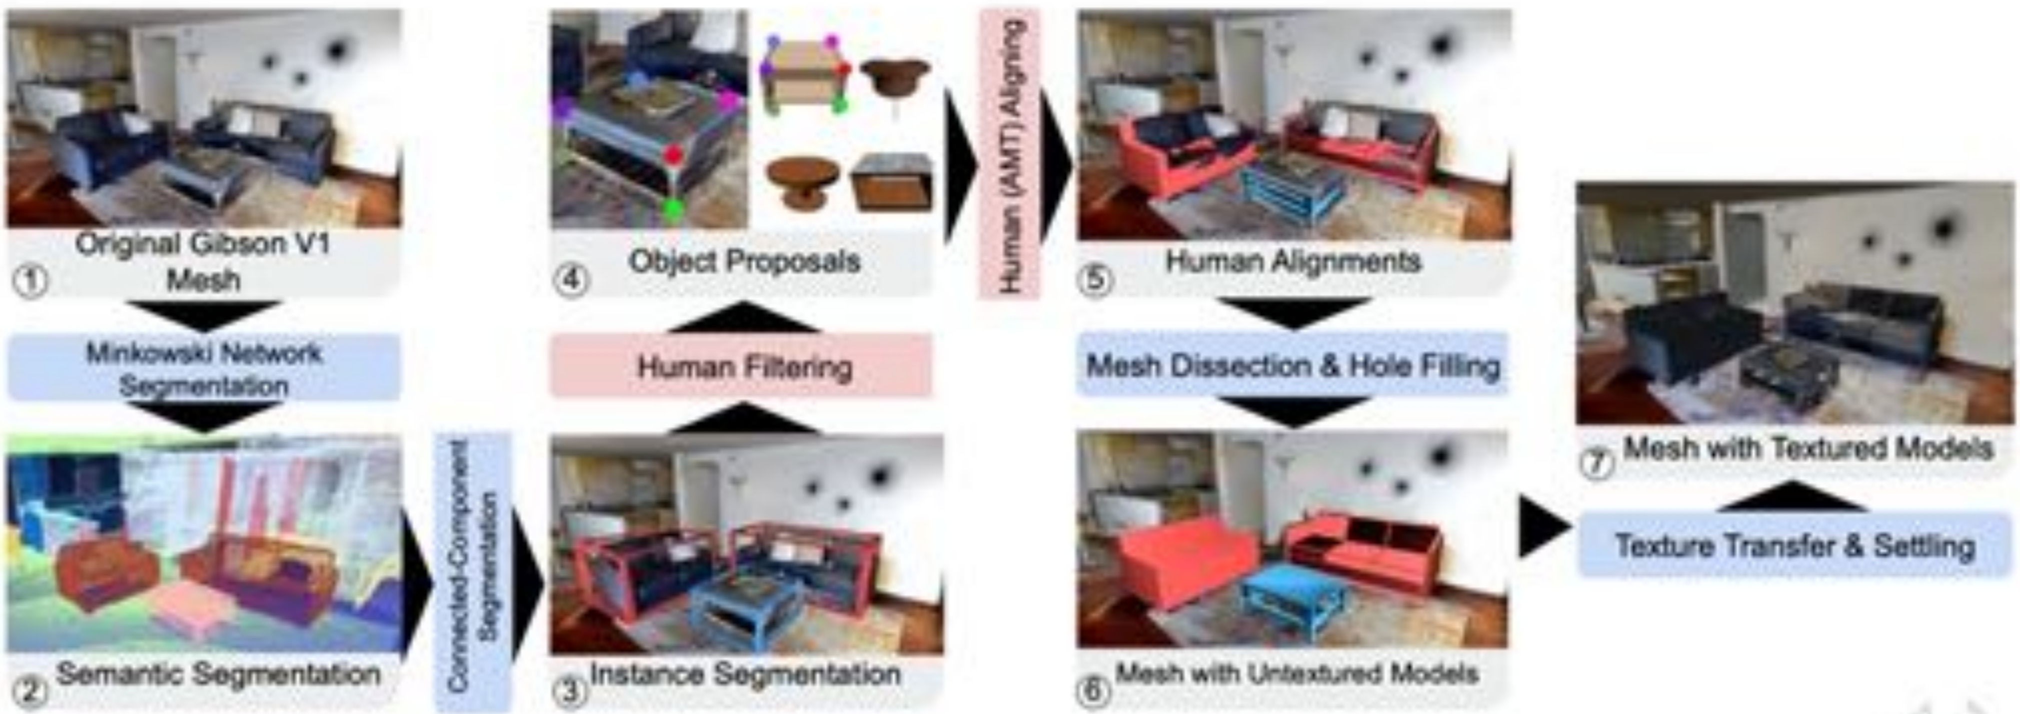

From Gibson to Interactive Gibson

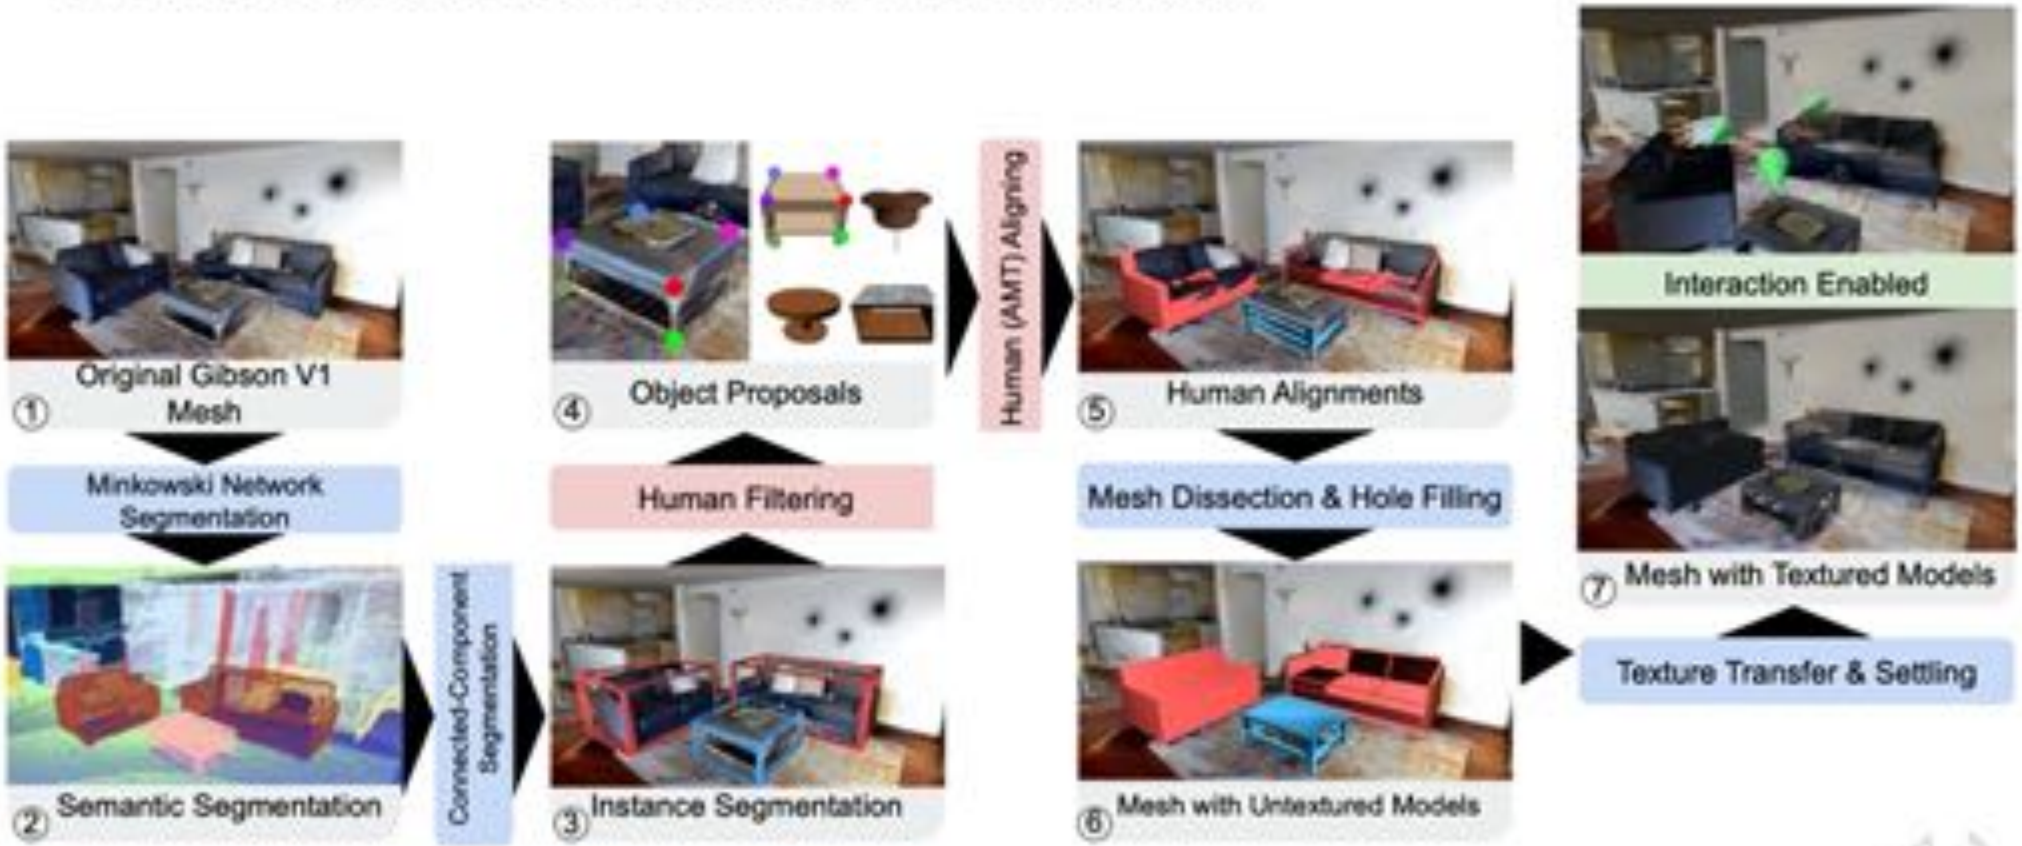

Xia, Fei, et al. "Interactive Gibson: A Benchmark for Interactive Navigation in Cluttered Environments." *arXiv preprint arXiv:1910.14442* (2019).

Interactive Objects in Interactive Gibson

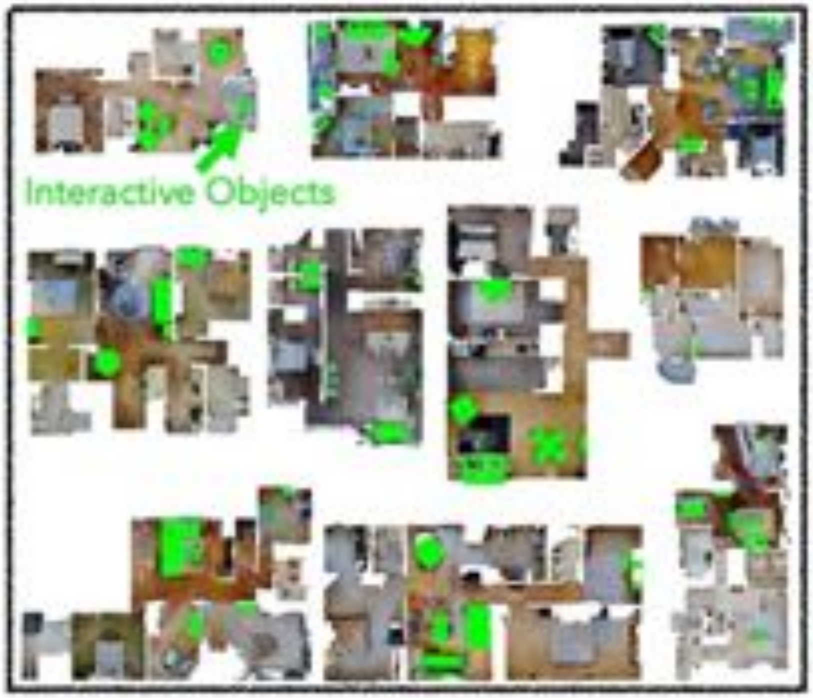

- doors
- chairs
- tables
- sofas
- beds

- toys
- shoes
- baskets
- ...

Gibson vs. Interactive Gibson

Gibson vs. Interactive Gibson: Speed

	Gibson V1	**Interactive Gibson**
RGBD, pre network f	58.5	264.1
RGBD, post network f	30.6	-
Surface Normal only	129.7	271.1
Semantic only	144.2	279.1
Non-Visual Sensory	396.1	1017.4

256x256 with full physical simulation
Nvidia GTX 1080ti

Hillsdale	620.4
Albertville	422.0

640x480 without physical simulation
Nvidia GTX 1080ti

What can be done in Interactive Gibson?

Interactive Search

Tidying a Room

Interactive Navigation

Interactive Navigation

Interactive Navigation tasks are navigation problems that require (or benefit from) physical interaction with the environment

Interactive Navigation with Reinforcement Learning

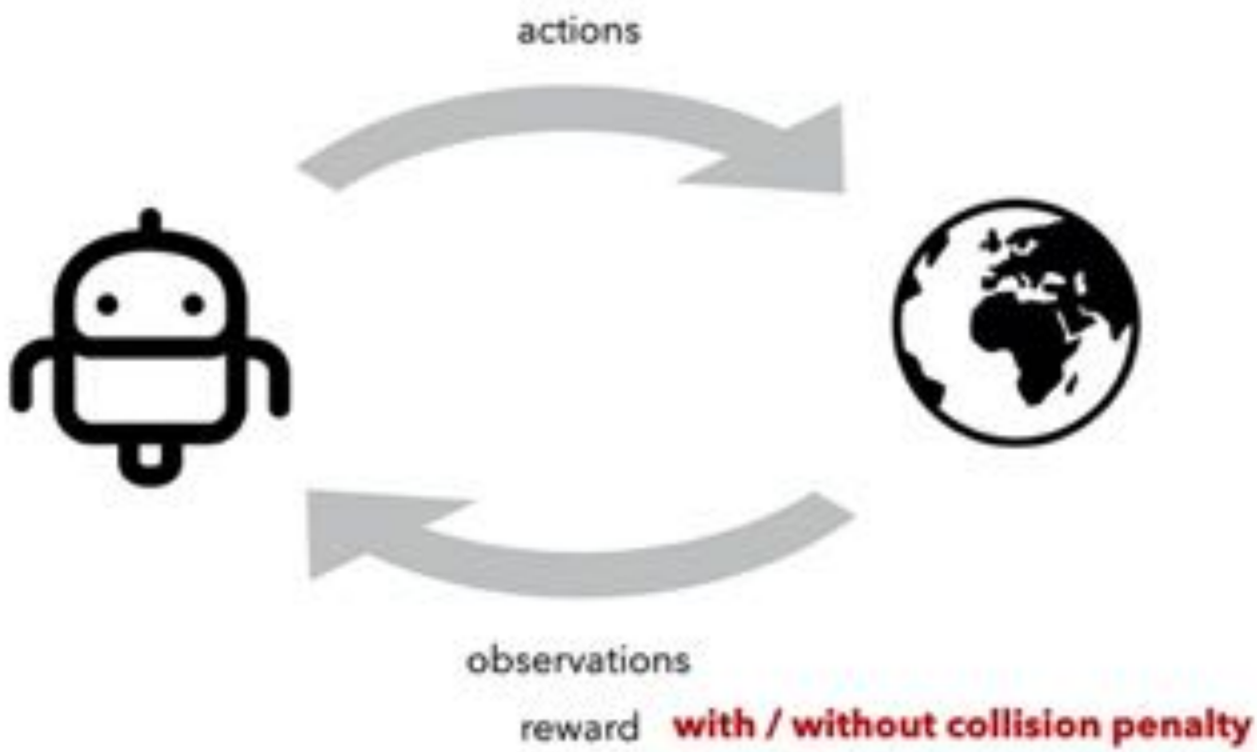

"Aggressive" Interactive Navigation Behavior

When there is no interaction penalty, agent doesn't avoid collision.

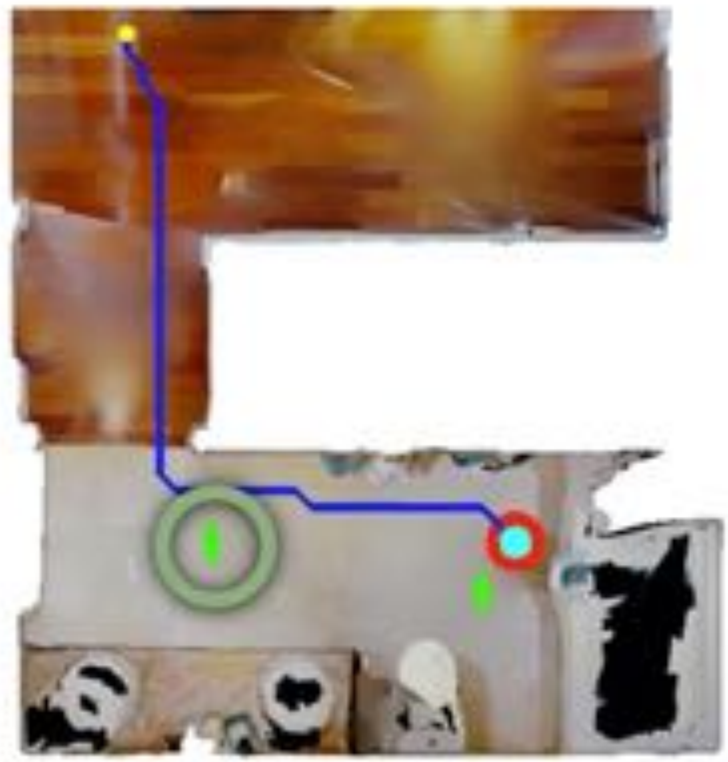

“Conservative” Interactive Navigation Behavior

As interaction penalty increase, agent learns to avoid collision (left). But can get blocked by simple objects.

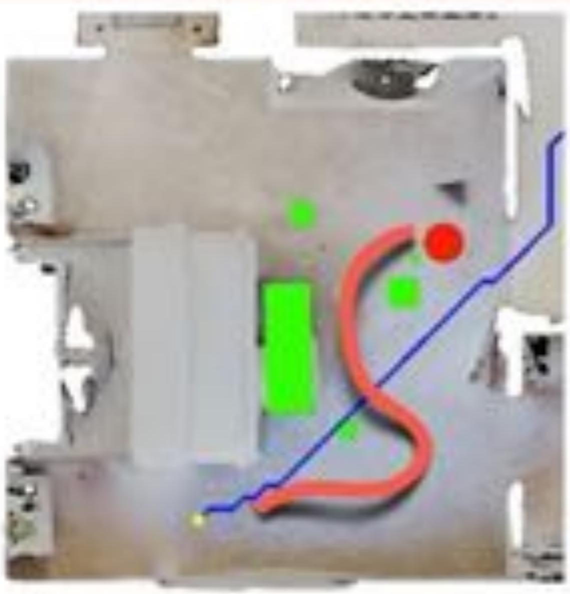

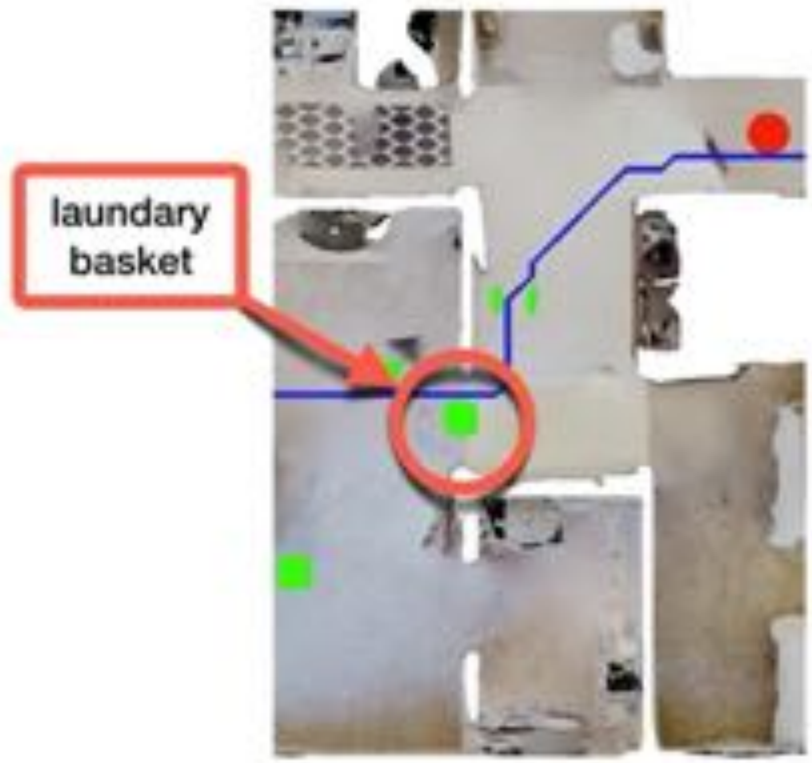

Interactive Navigation Score

$$INS = \alpha P_{eff} + (1 - \alpha)E_{eff}$$

$\alpha = 0.5$ balanced score between path and effort efficiency

Interactive Navigation with Mobile Manipulators

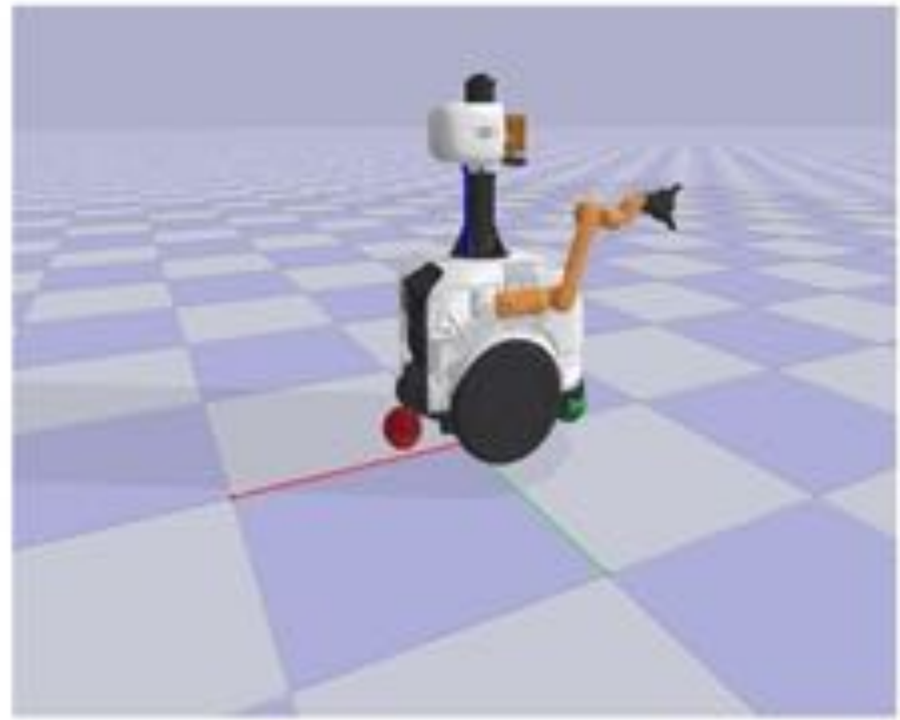

Opening a Door: a Heterogeneous IN Task

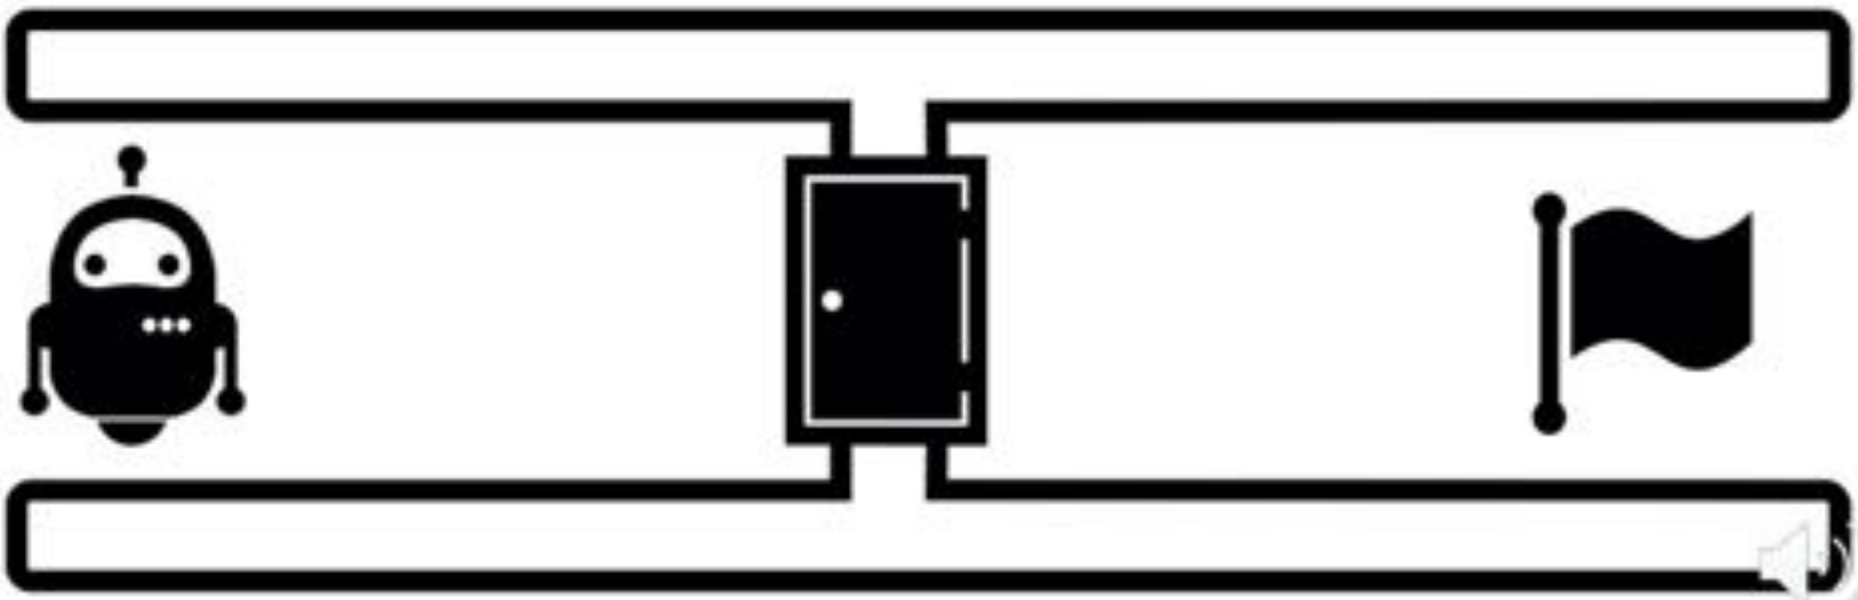

Opening a Door: a Heterogeneous IN Task

Subtask 1
Get to the door

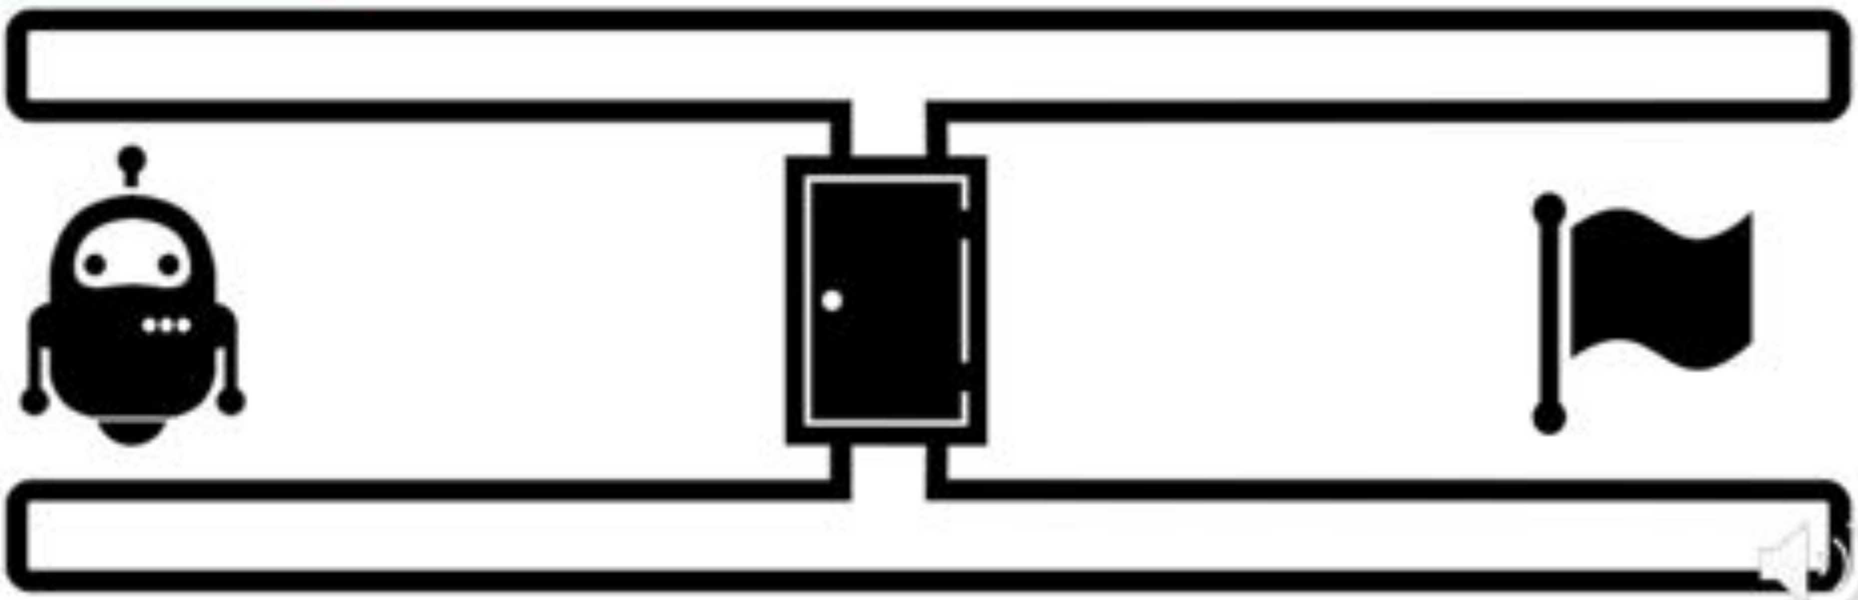

Opening a Door: a Heterogeneous IN Task

Subtask 1
Get to the door

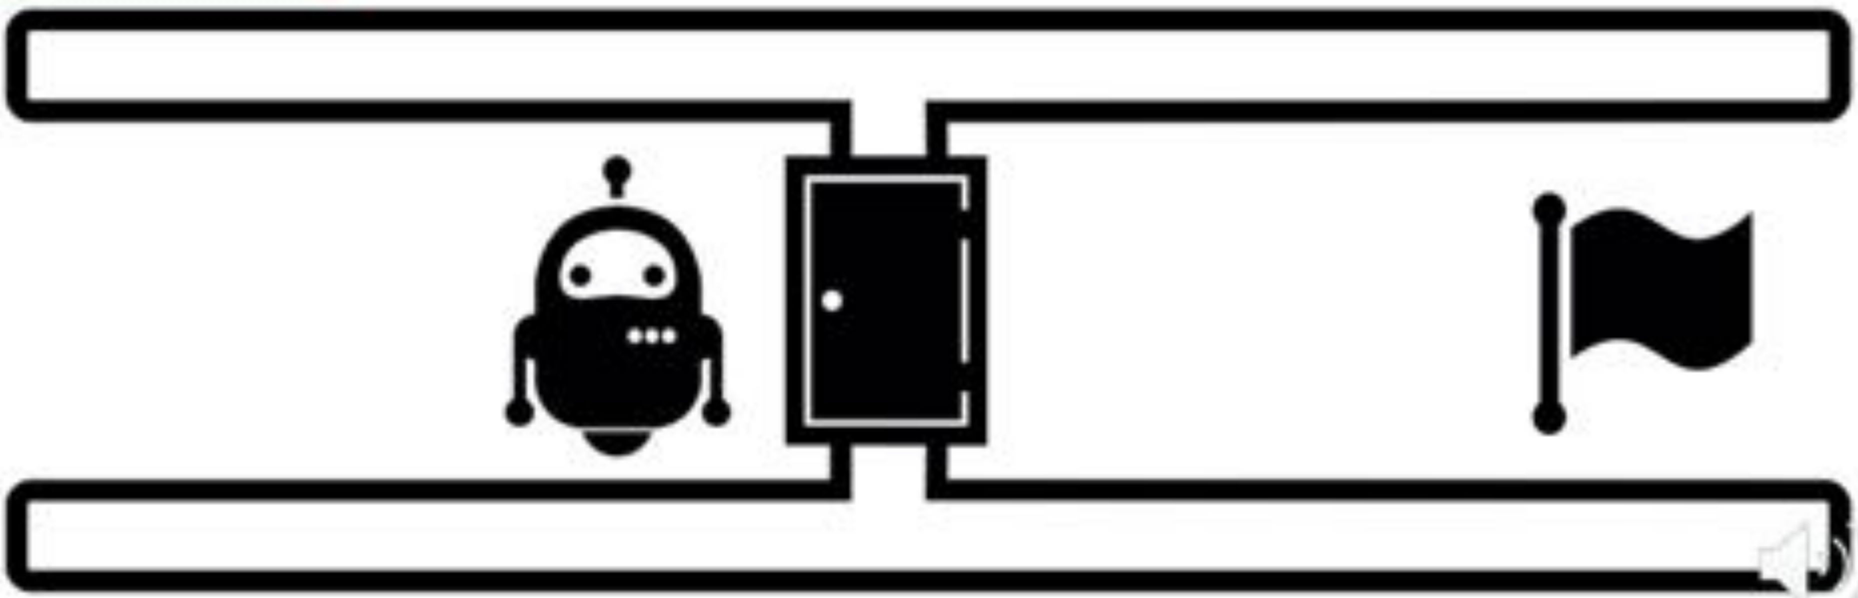

Opening a Door: a Heterogeneous IN Task

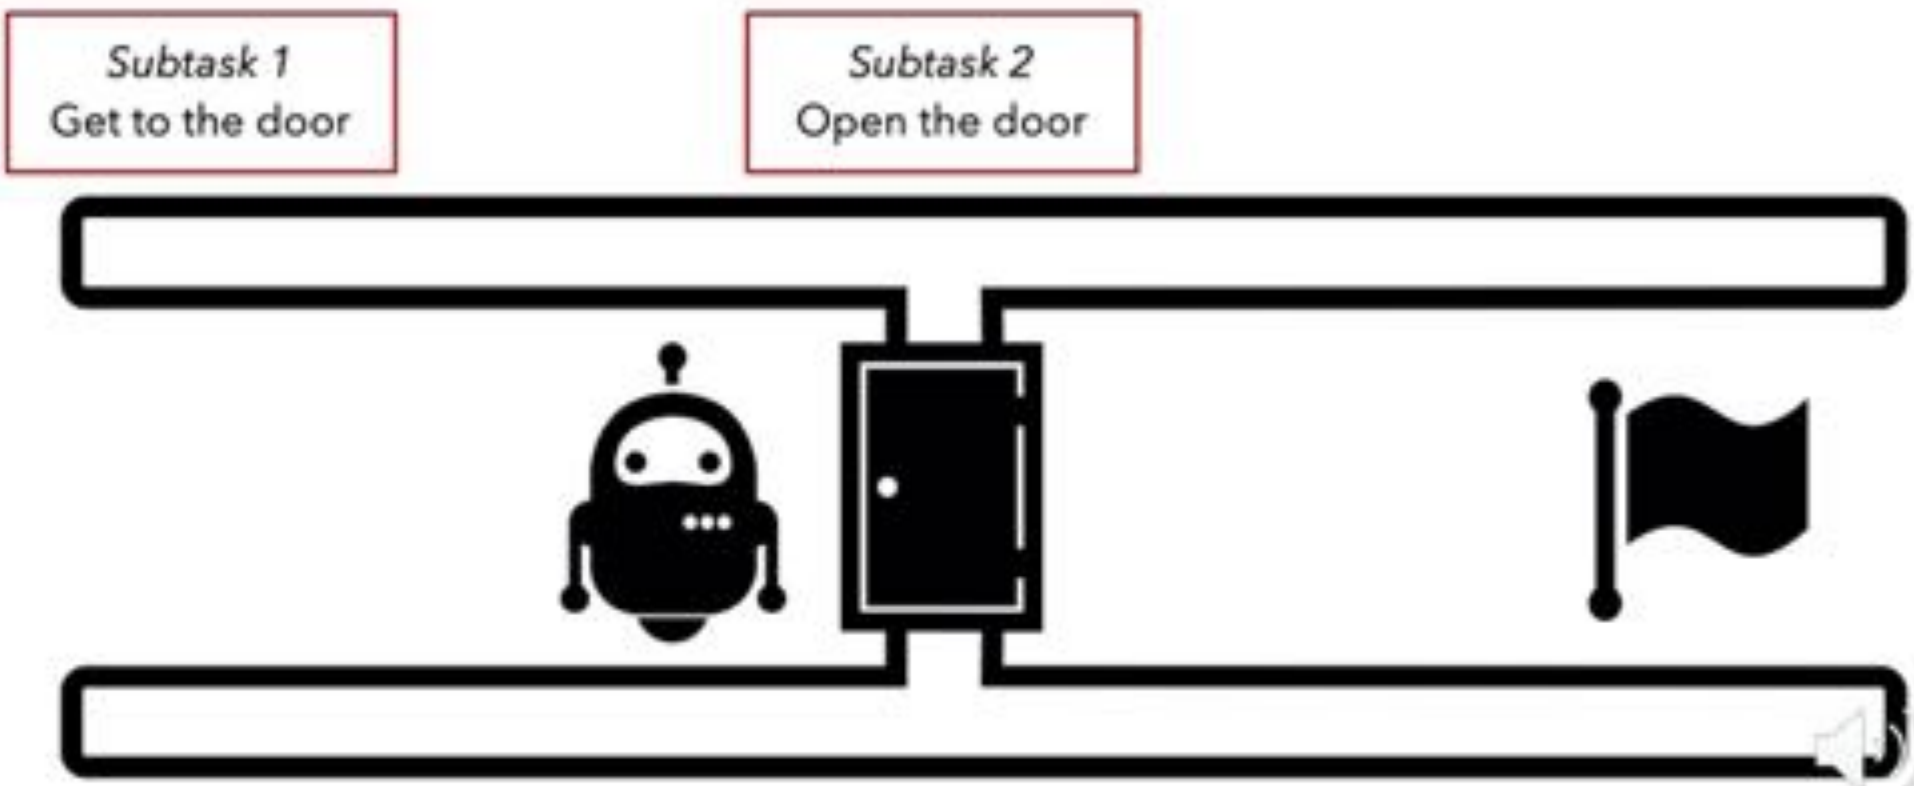

Opening a Door: a Heterogeneous IN Task

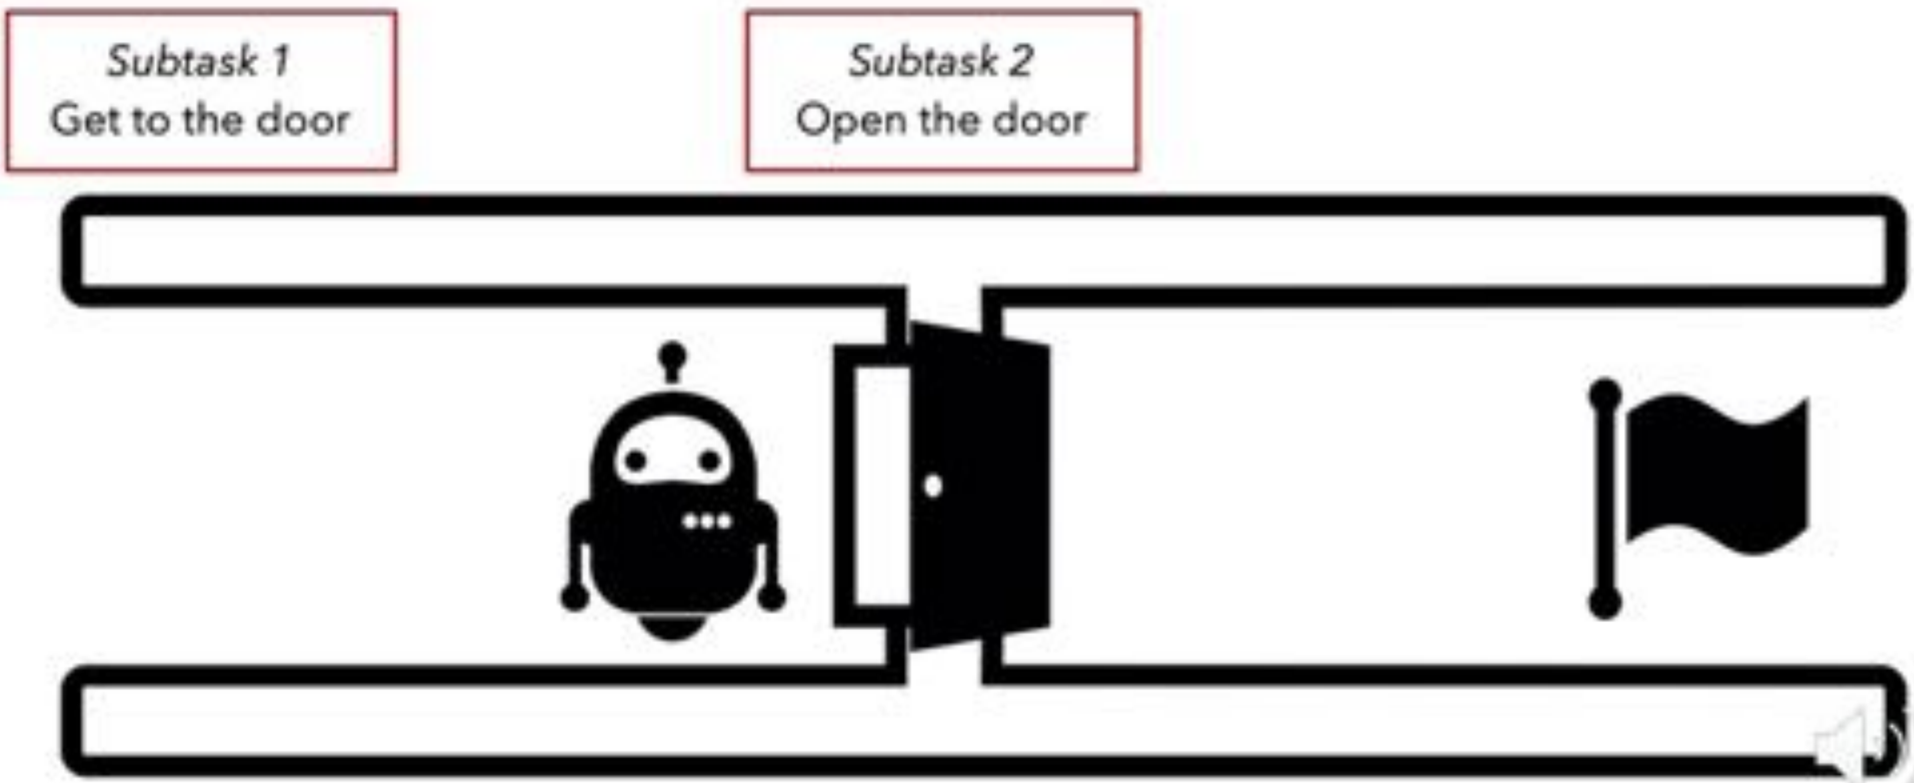

Opening a Door: a Heterogeneous IN Task

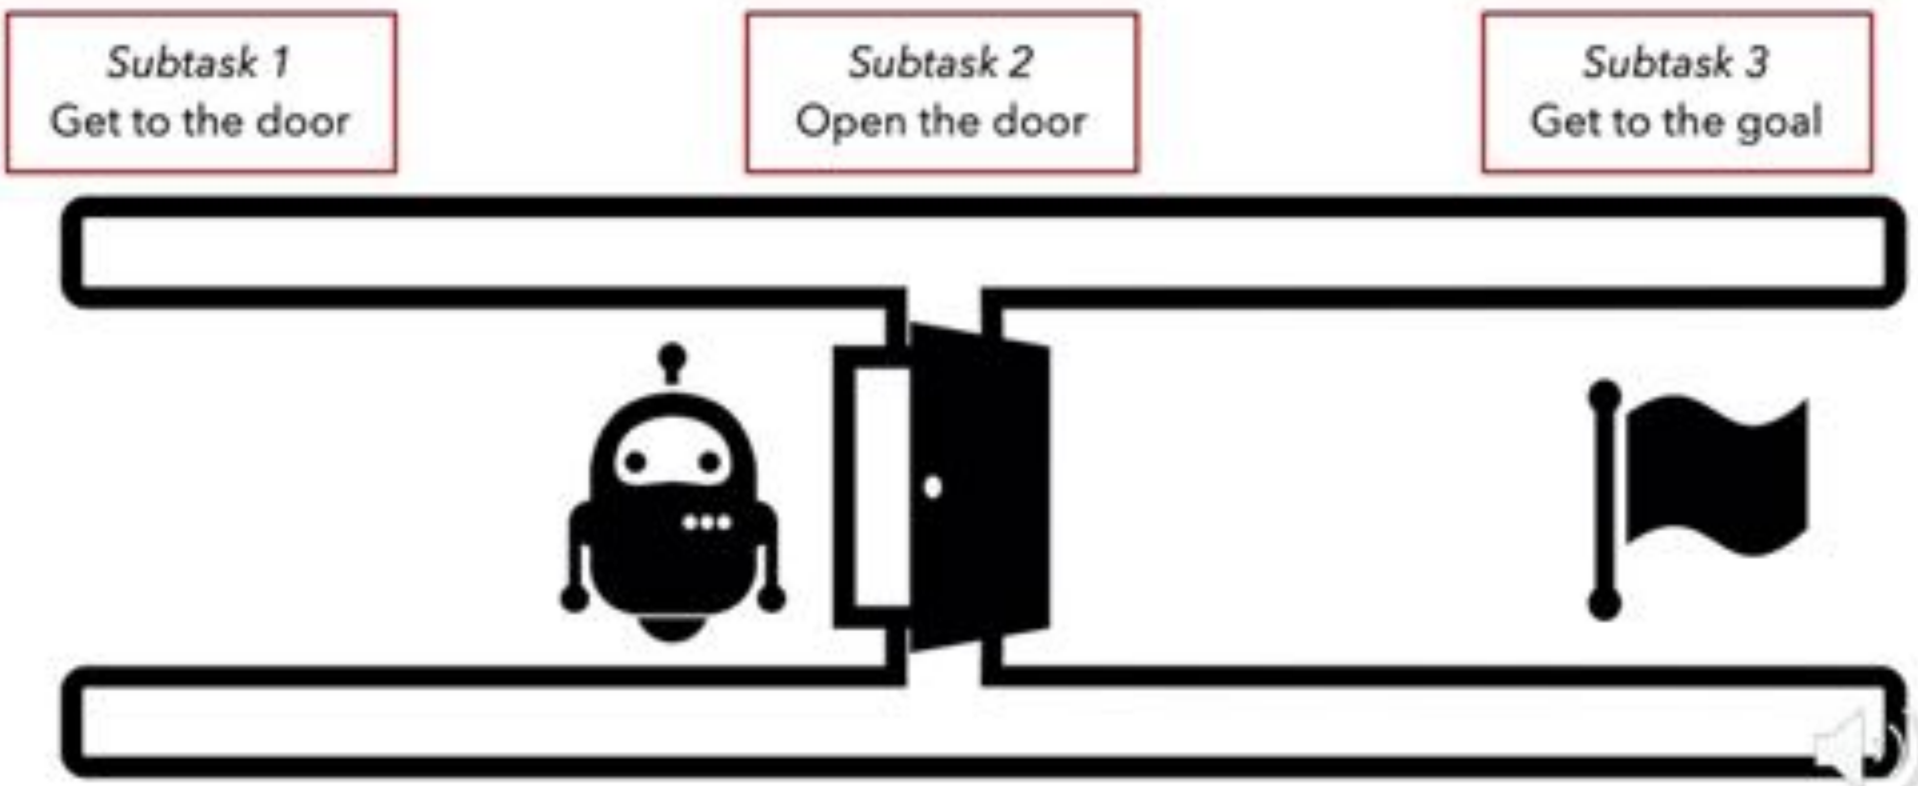

Opening a Door: a Heterogeneous IN Task

Subtask 1
Get to the door

Subtask 2
Open the door

Subtask 3
Get to the goal

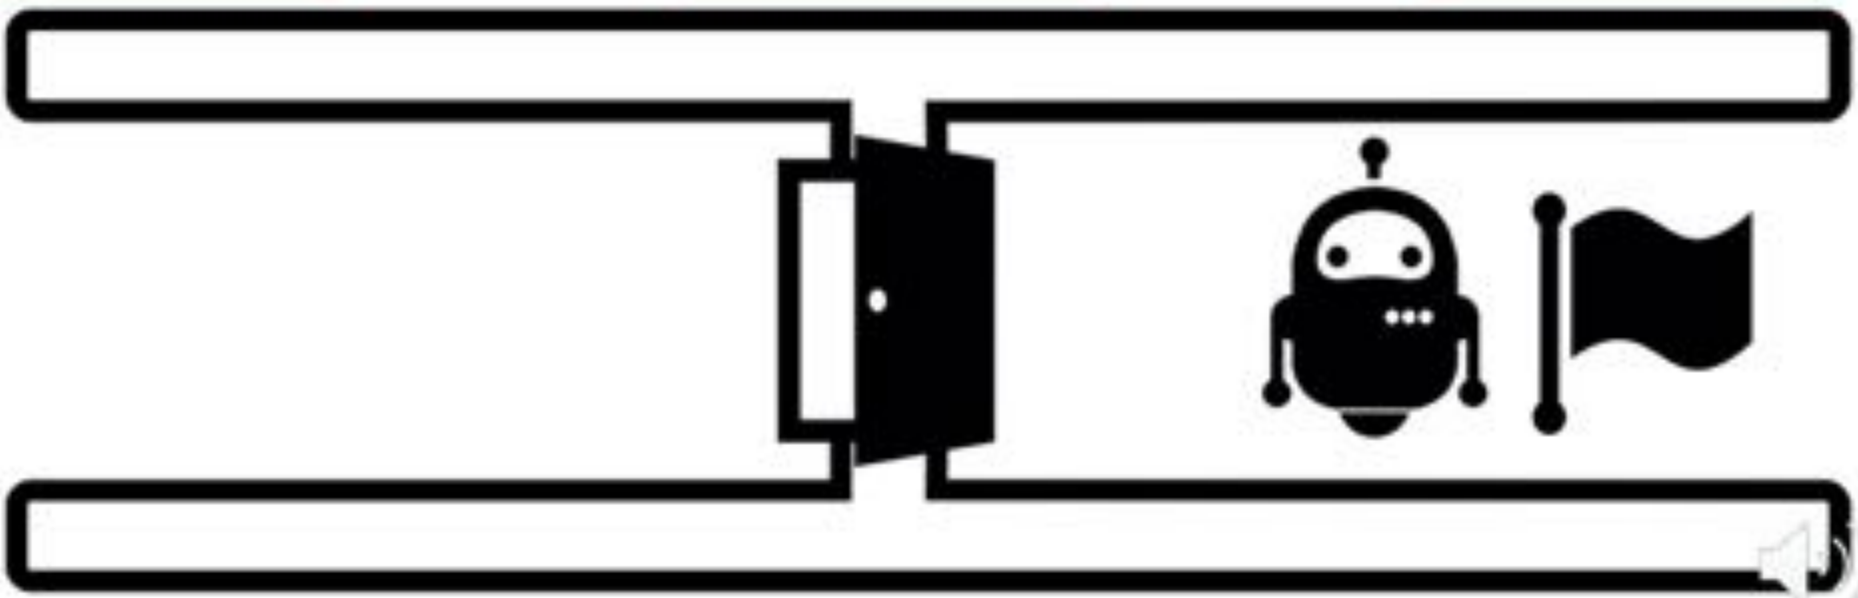

HRL4IN: Hierarchical RL for Interactive Navigation

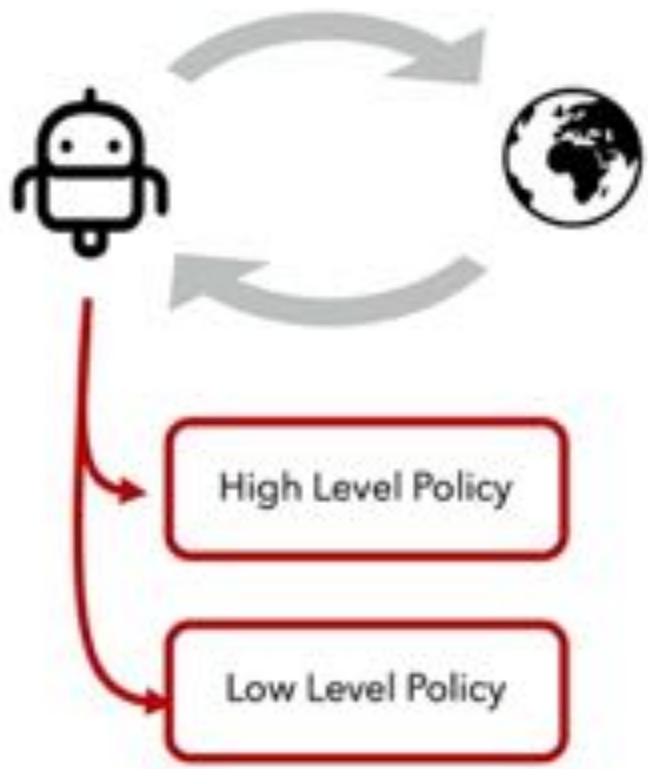

HRL4IN: Hierarchical RL for Interactive Navigation

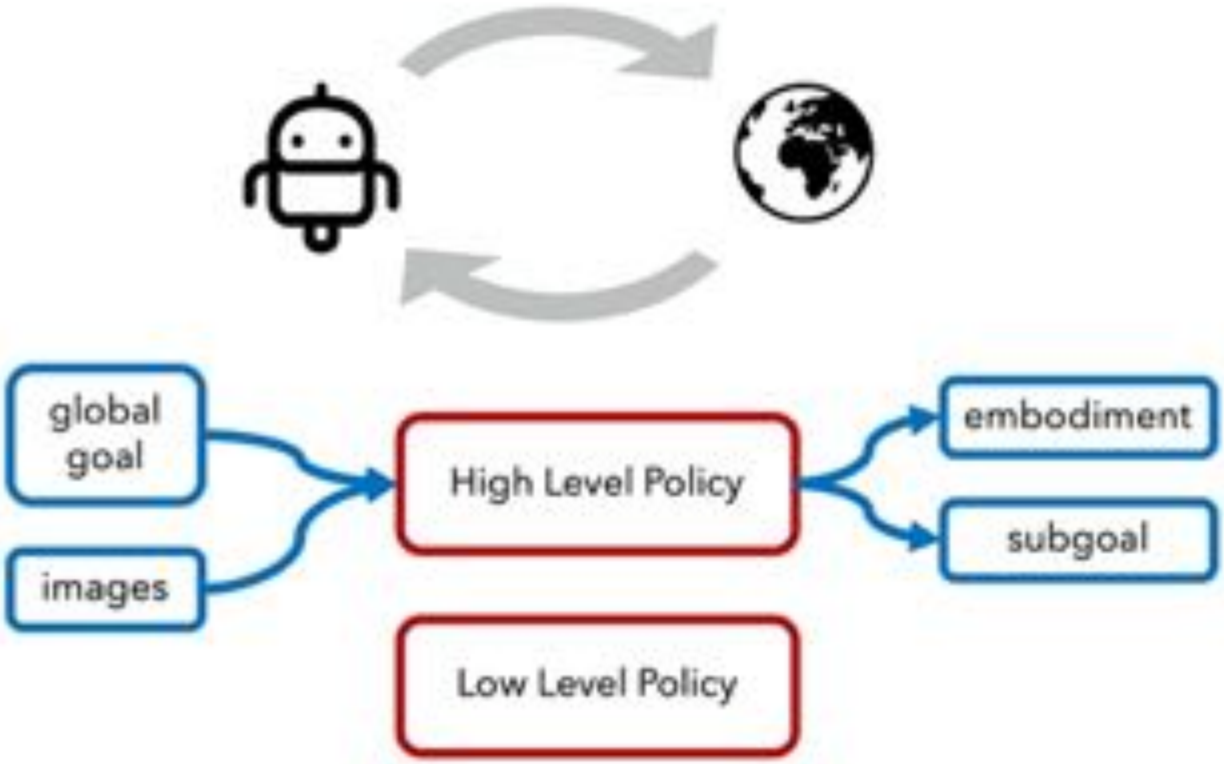

HRL4IN: Hierarchical RL for Interactive Navigation

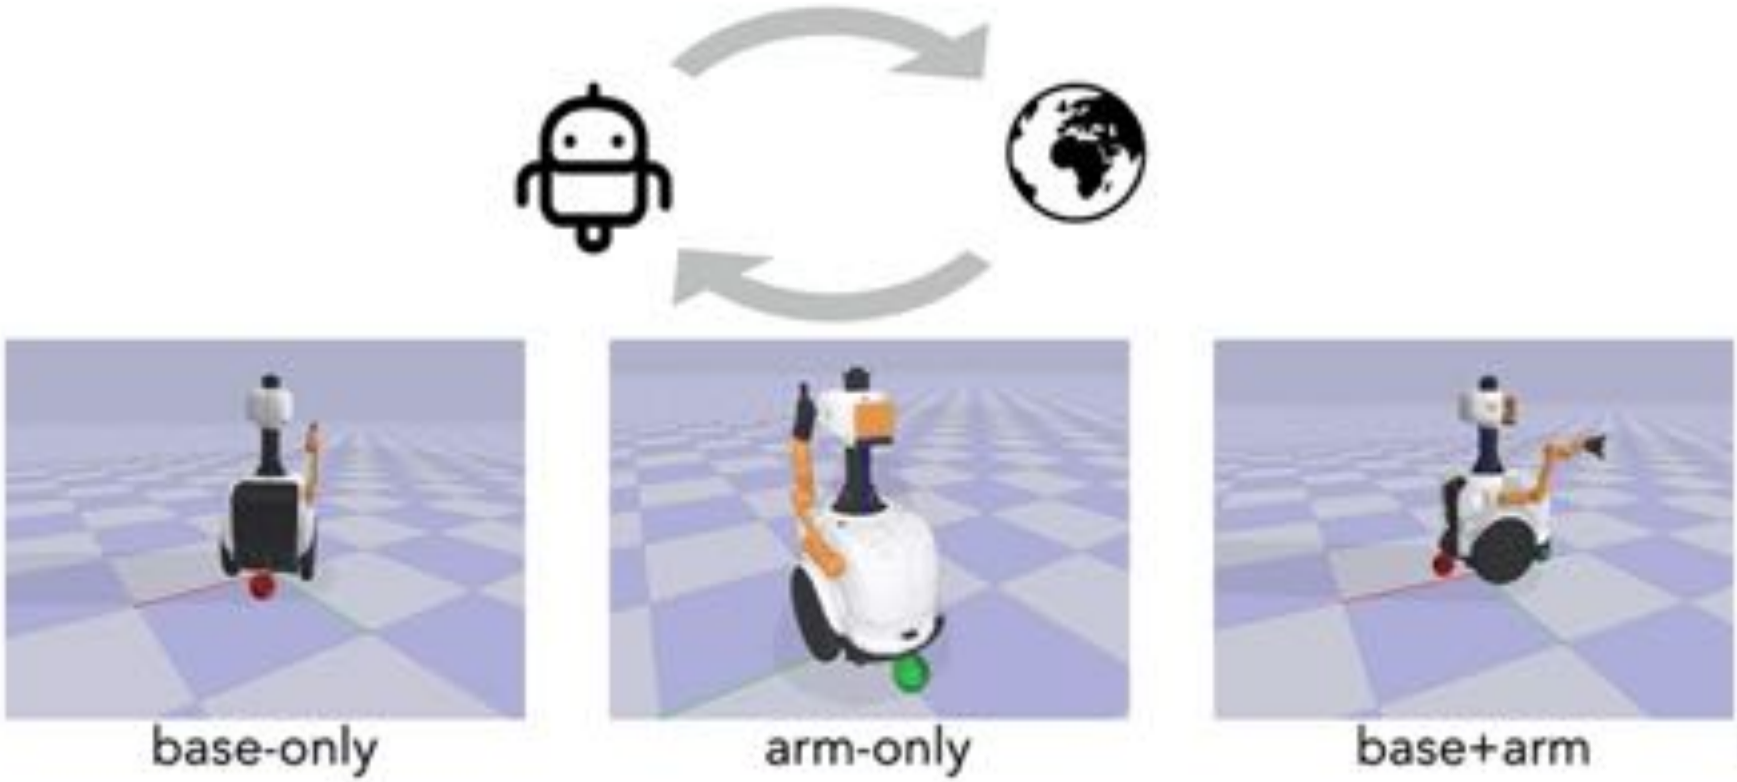

HRL4IN: Hierarchical RL for Interactive Navigation

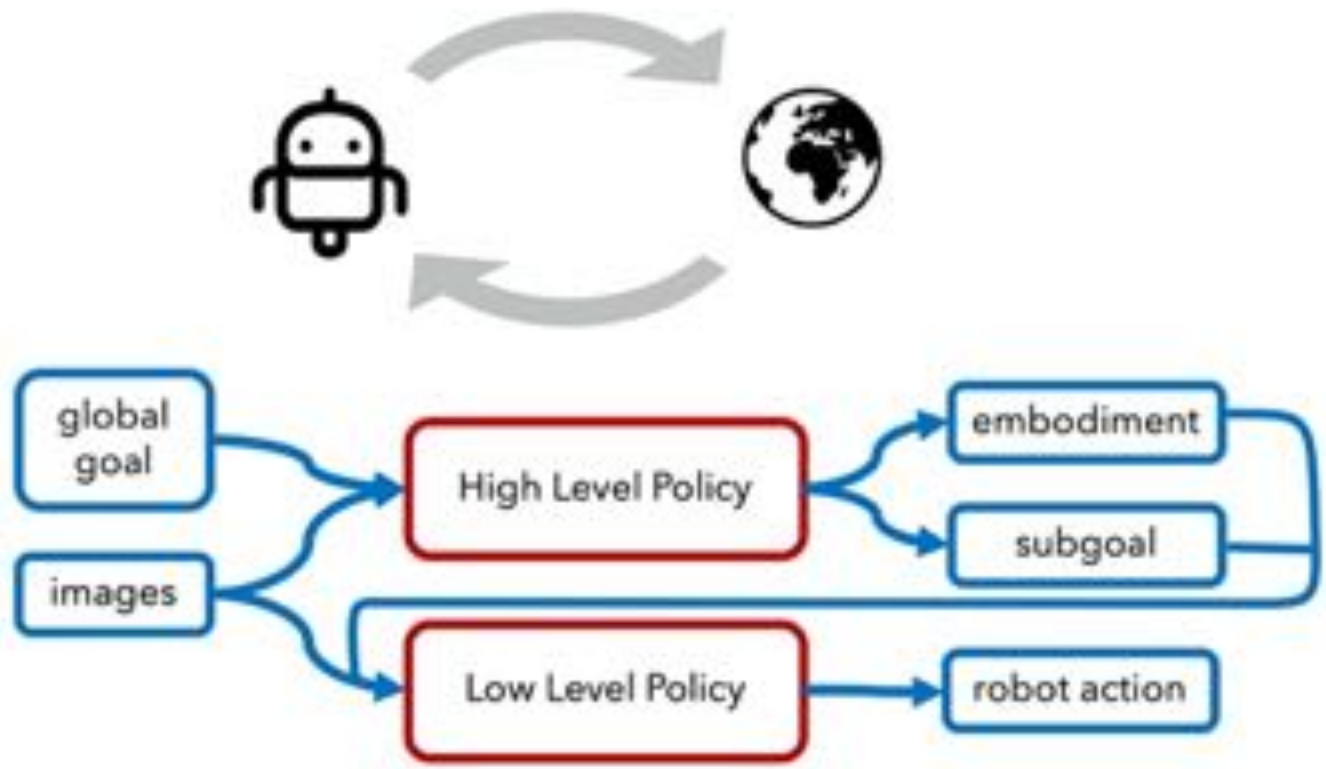

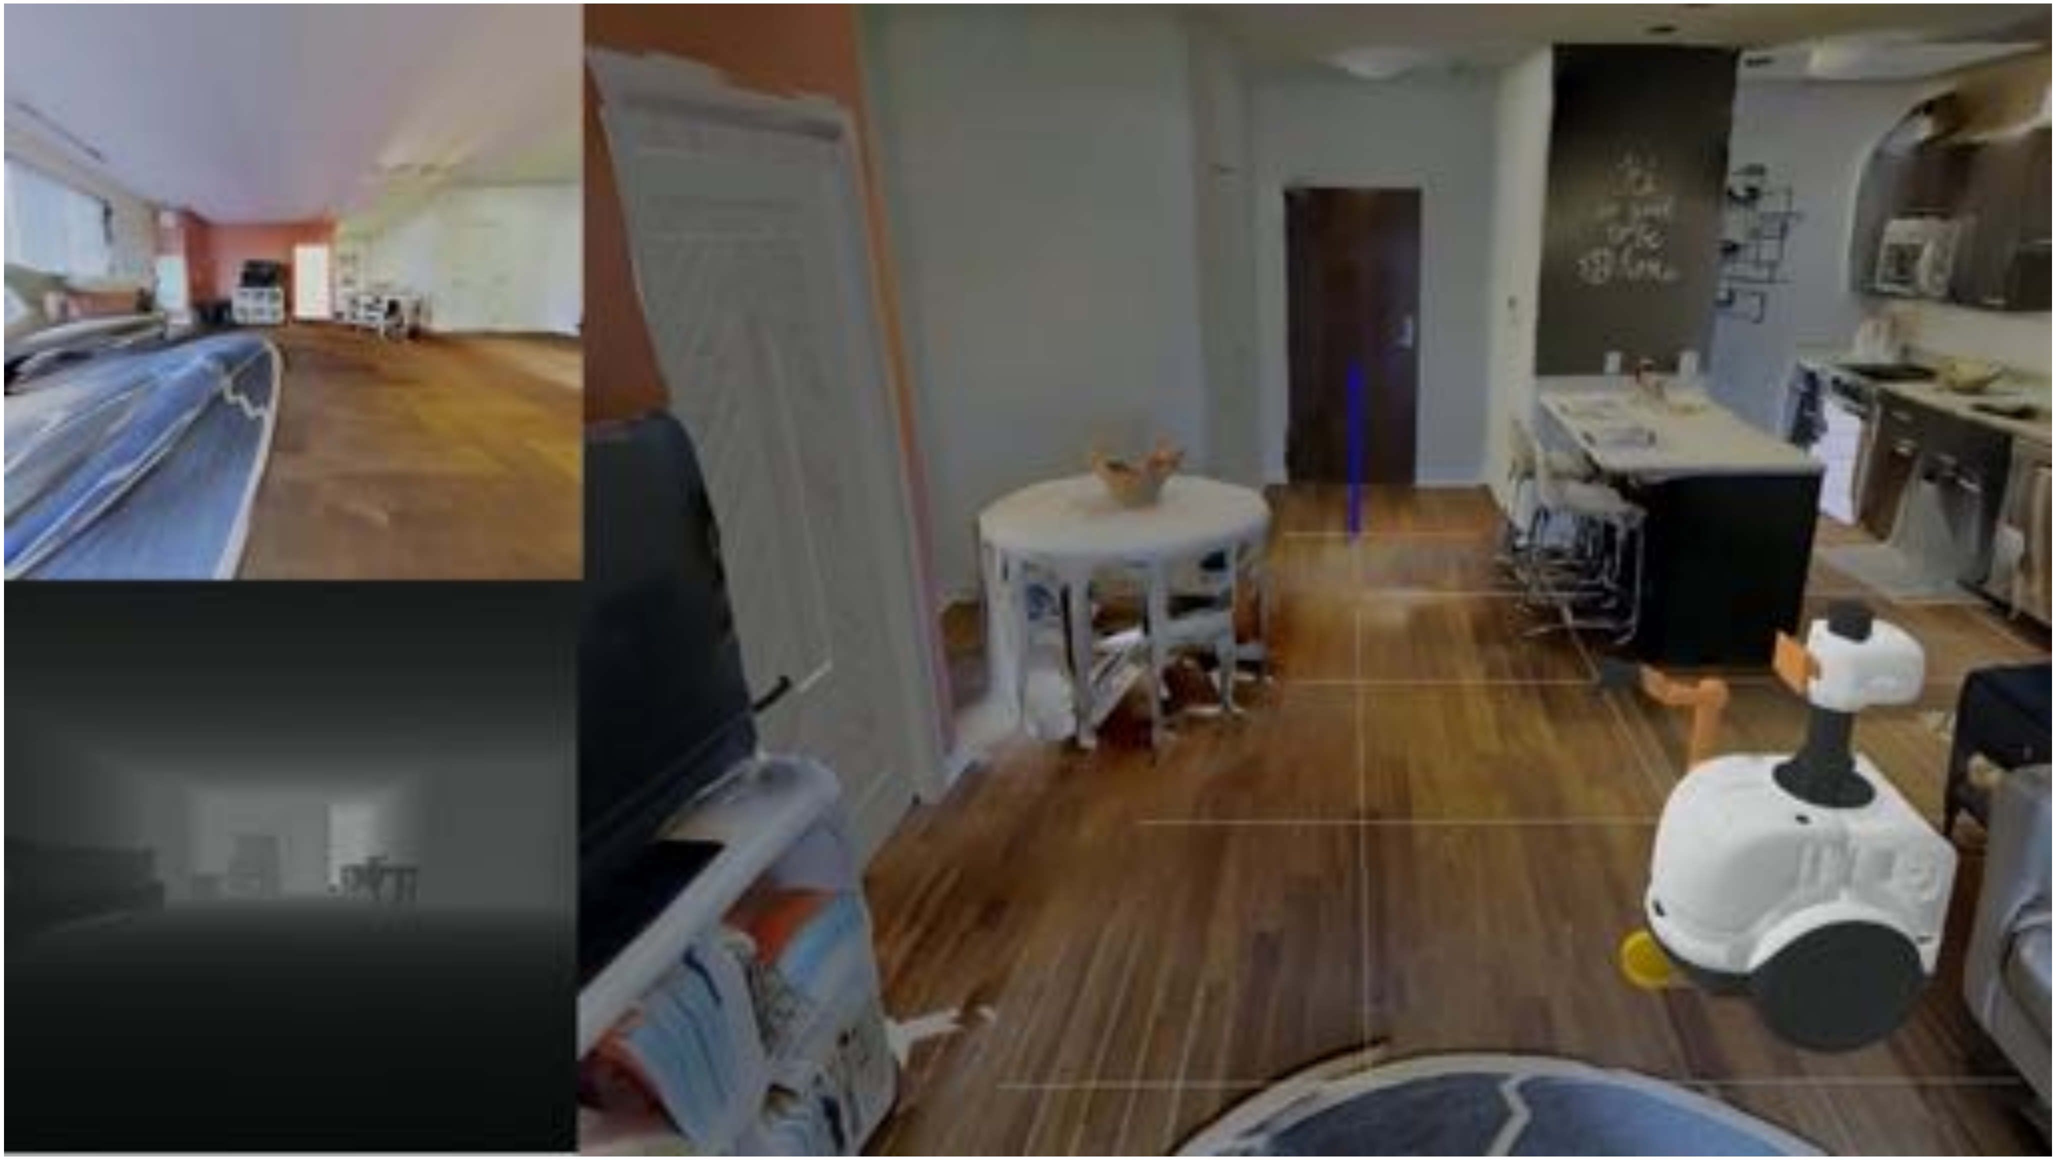

What did we achieve with Interactive Gibson

- Make objects interactive
- Keep real-world semantic distribution
- Visuo-control learned solutions for Interactive Navigation
- Proposed Interactive Navigation Score

Summary

- Gibson is a state-of-the-art simulator to train robots for visuo-motor tasks: navigation and manipulation
- Includes hundreds of model of real-world large environments with interactive objects
- Enables easier sim2real transference of learned strategies via realistic virtual images

- We continue improving Gibson in multiple fronts. Check it out!

Download Gibson and try it yourself!

Gibson2 Code

https://github.com/StanfordVL/iGibson

Gibson2 Website

http://svl.stanford.edu/igibson/

Install it with "pip install gibson2"

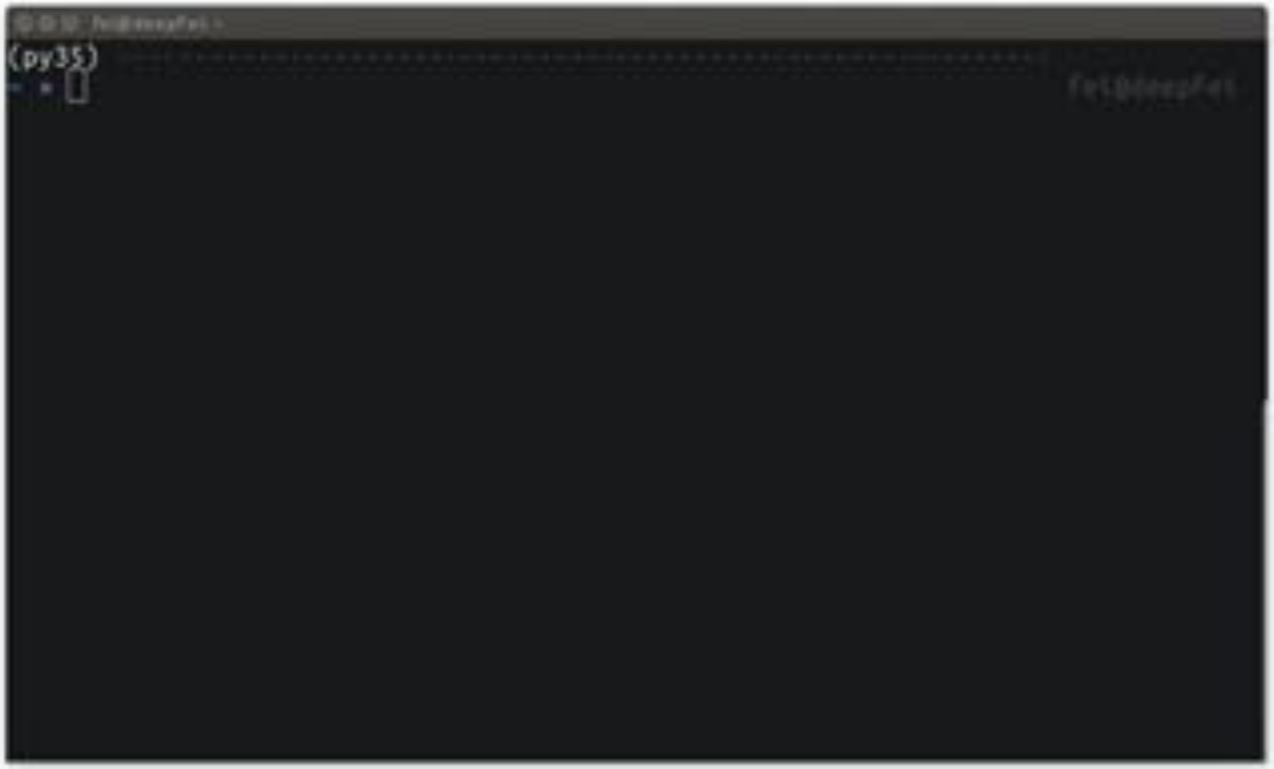

Run demo with "python -m gibson2.envs.demo"

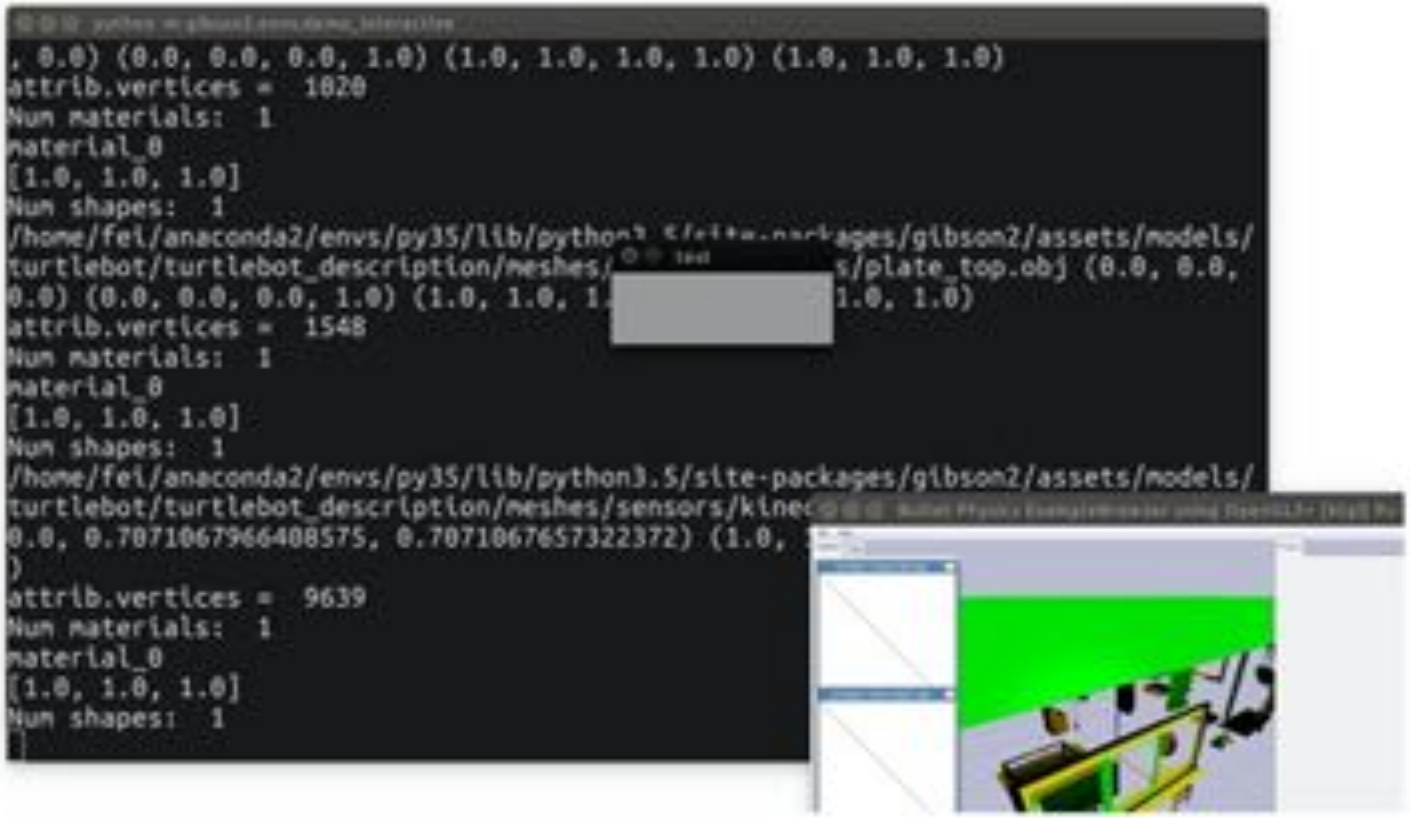

The Gibson Team

Tribute to James J. Gibson [1904-1979]

- American psychologist
- Coined the term "ecological perception"
- Defended the idea that perception should be studied as an active process of embodied agents in their environments

"We must perceive in order to move, but we must also move in order to perceive"

iGibson

Thank you for listening!

Contact us: robertom@stanford.edu
feixia@stanford.edu