Emergence of Maps in the Memories of Blind Navigation Agents

Erik Wijmans

Manolis Savva

Irfan Essa

Stefan Lee

Ari Morcos

Dhruv Batra

Georgia Tech

Meta AI

SFU

Google

Oregon State University

Animals build maps

Animals build maps

Tolman, 1948

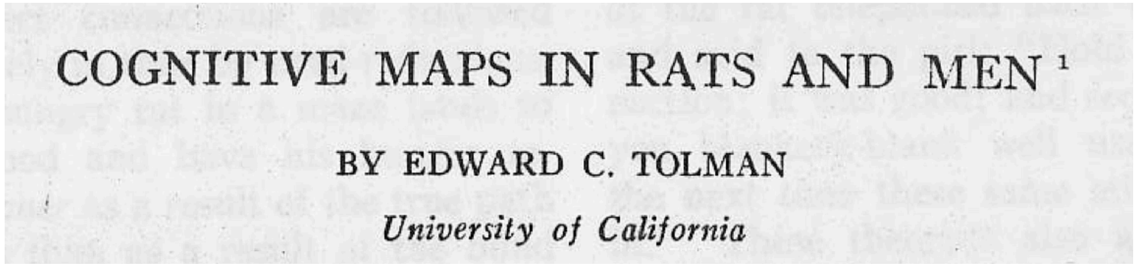

COGNITIVE MAPS IN RATS AND MEN[1]

BY EDWARD C. TOLMAN

University of California

Toledo et al (2020) and Harten et al (2020)

Bats navigate with cognitive maps

O'Keefe and Nadel, 1978

THE

HIPPOCAMPUS

AS A COGNITIVE MAP

O'Keefe, Moser, and Moser, 2014

The Nobel Prize in Physiology or Medicine 2014

Moser "for their discoveries of cells that constitute a positioning system in the brain"

Animals build maps

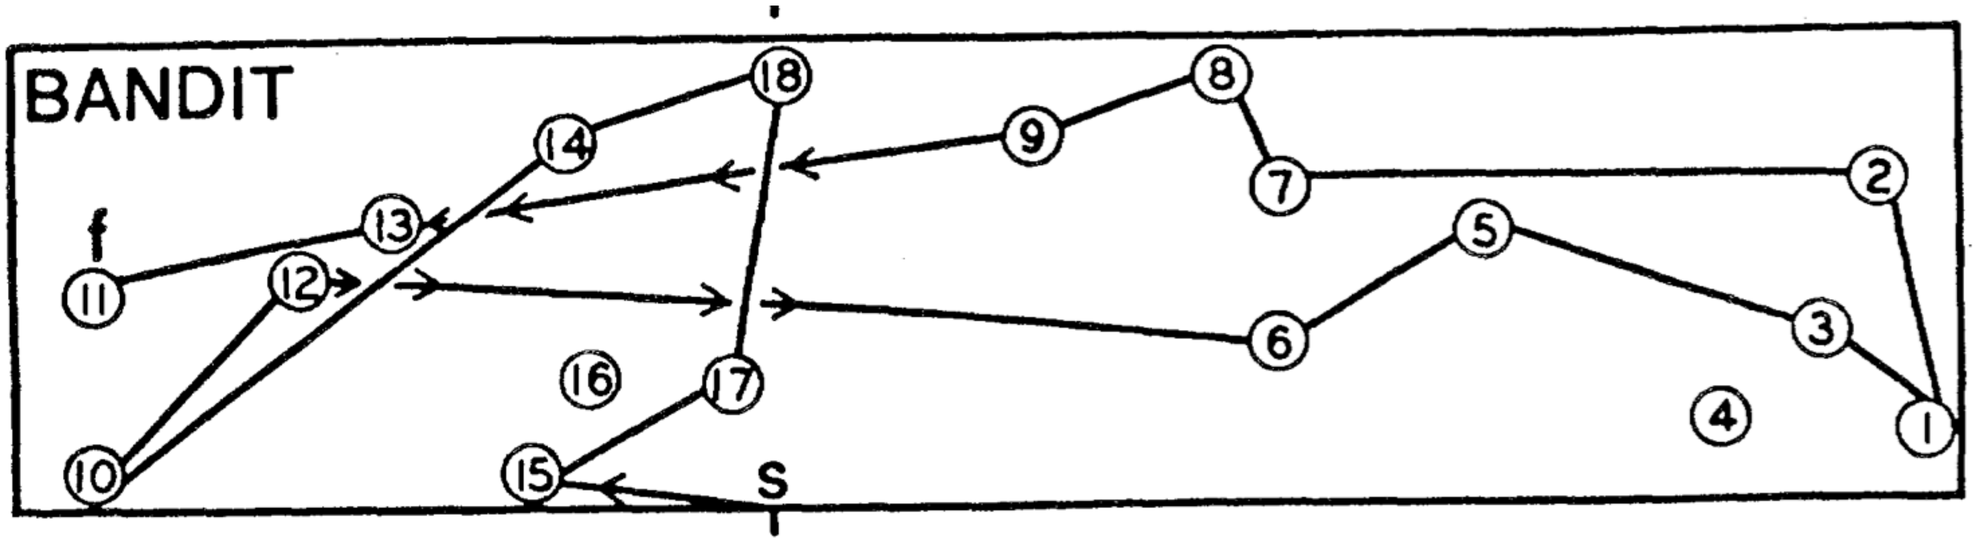

Chimpanzees

Menzel 1973

Animals build maps

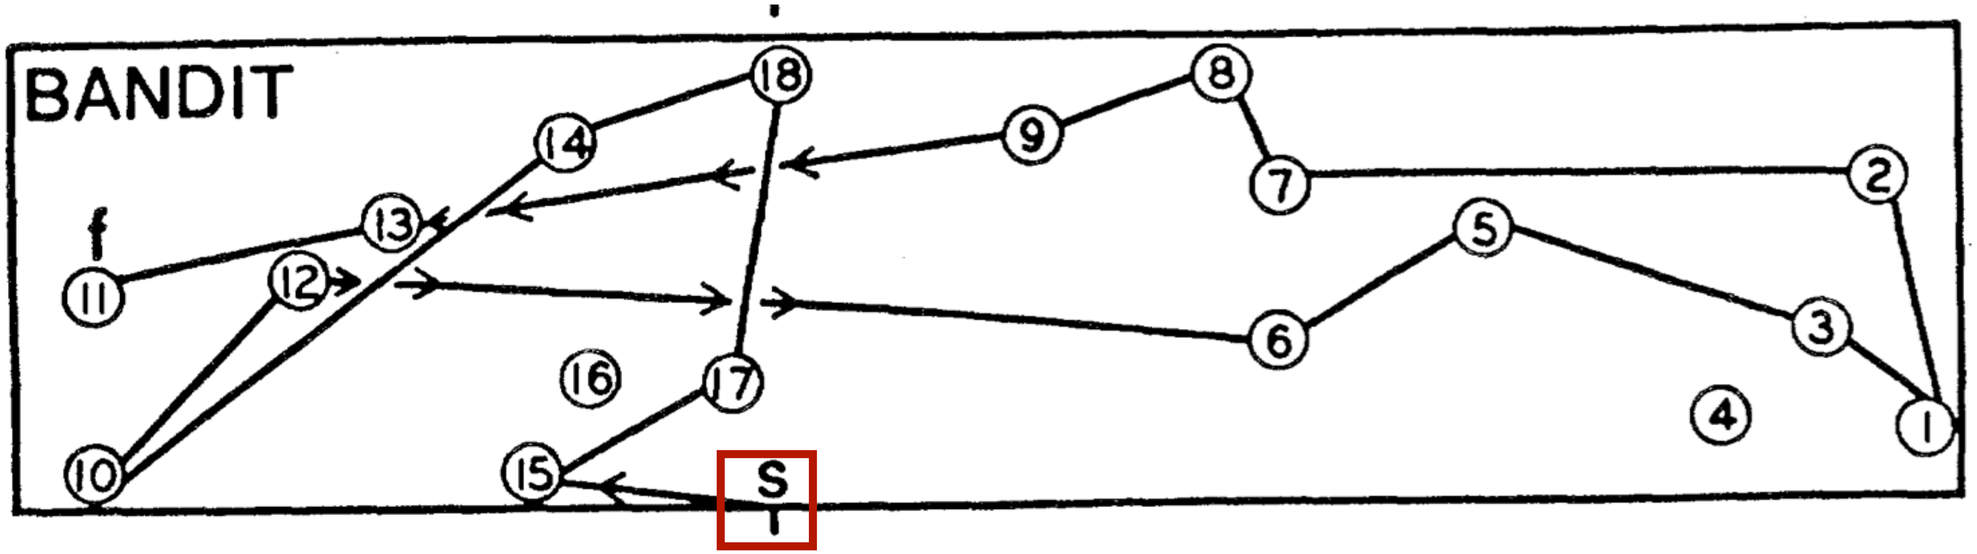

Chimpanzees

Menzel 1973
Image Credit: Frans B. M. de Waal.
A century of getting to know the chimpanzee. Nature 2005

Animals build maps

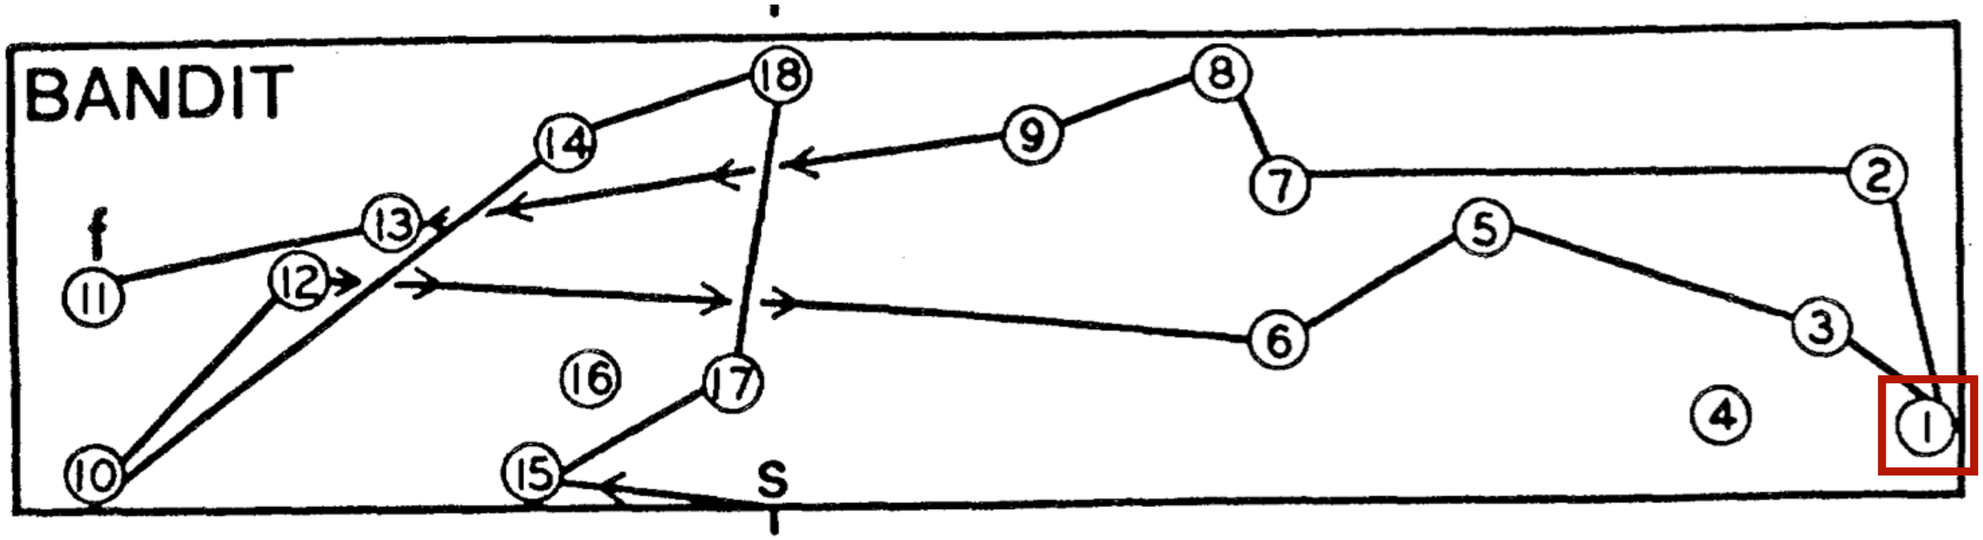

Chimpanzees

Menzel 1973
Image Credit: Frans B. M. de Waal.
A century of getting to know the chimpanzee. Nature 2005

Animals build maps

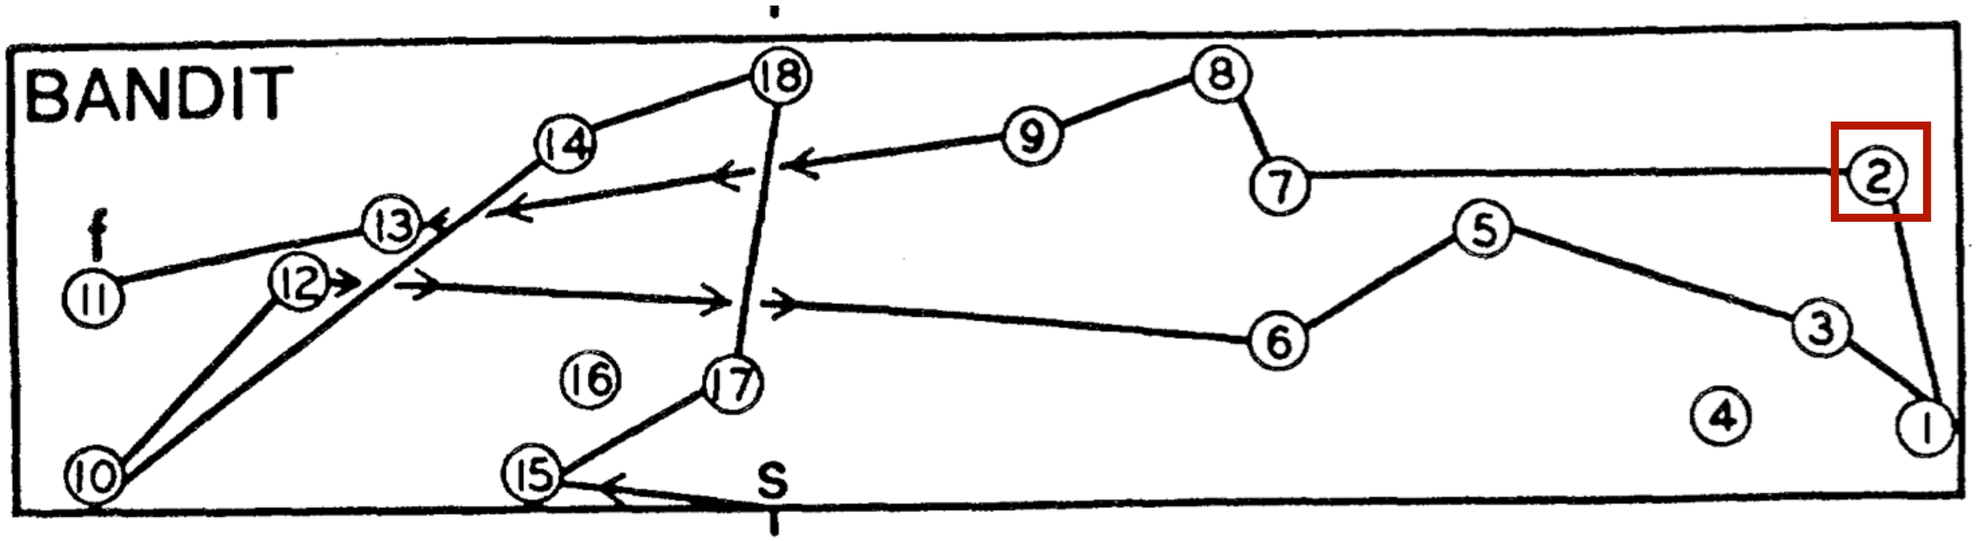

Chimpanzees

Menzel 1973
Image Credit: Frans B. M. de Waal.
A century of getting to know the chimpanzee. Nature 2005

Animals build maps

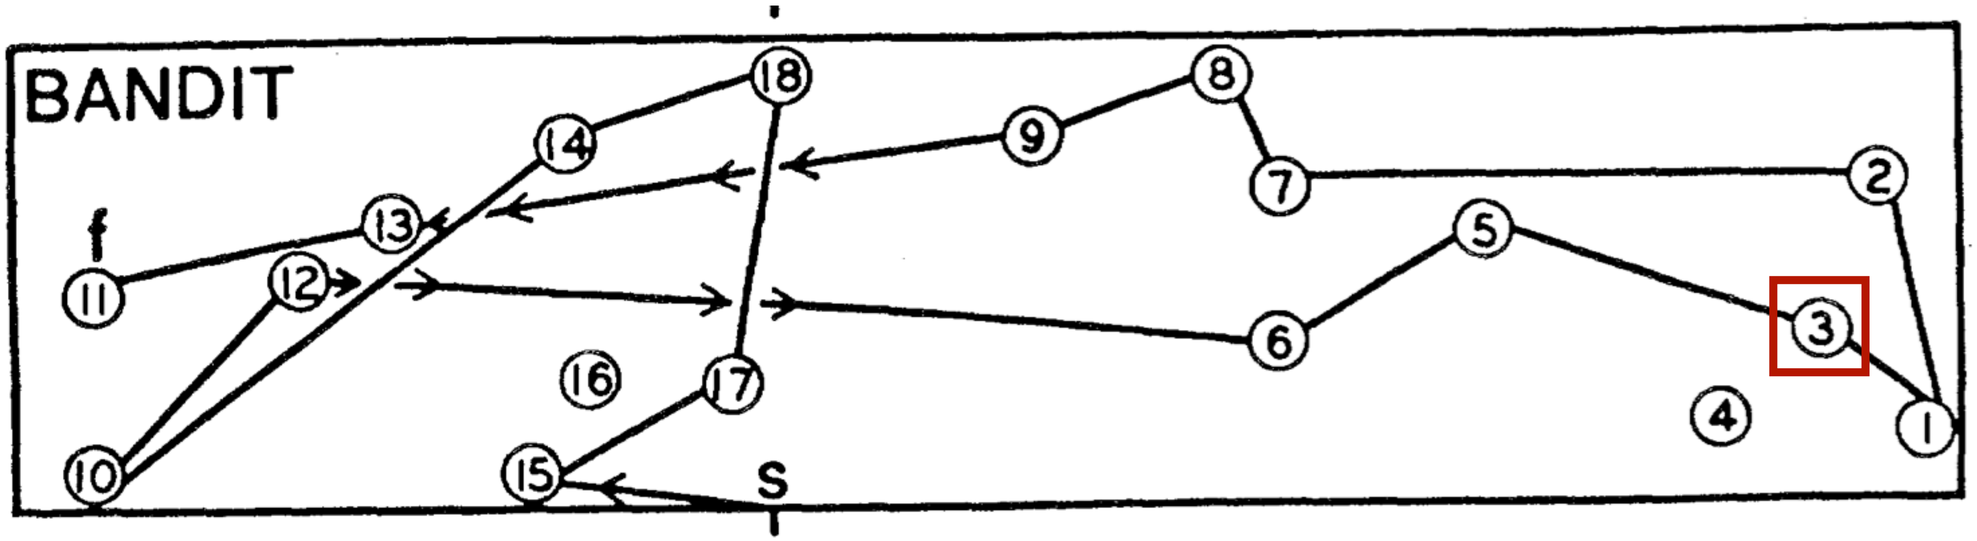

Chimpanzees

Menzel 1973
Image Credit: Frans B. M. de Waal.
A century of getting to know the chimpanzee. Nature 2005

Animals build maps

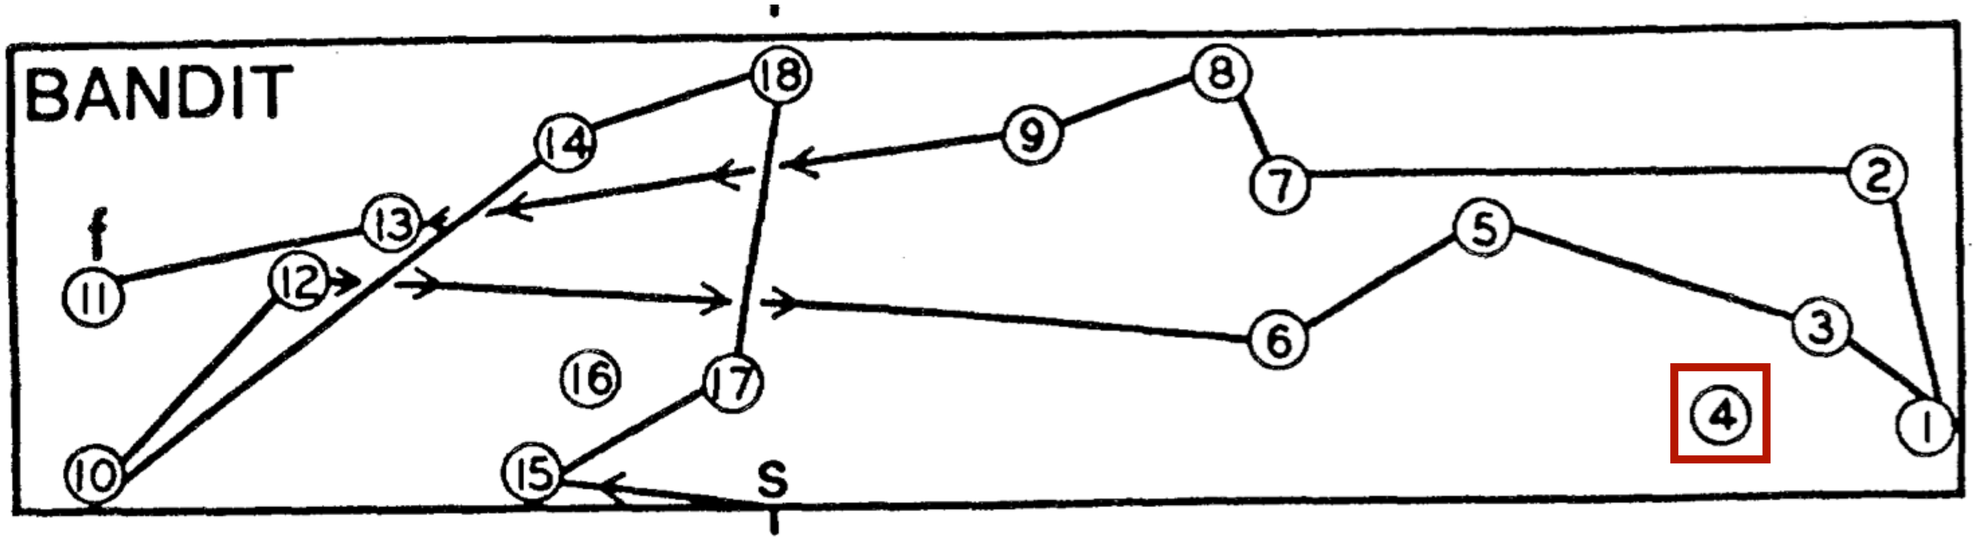

Chimpanzees

Menzel 1973
Image Credit: Frans B. M. de Waal.
A century of getting to know the chimpanzee. Nature 2005

Animals build maps

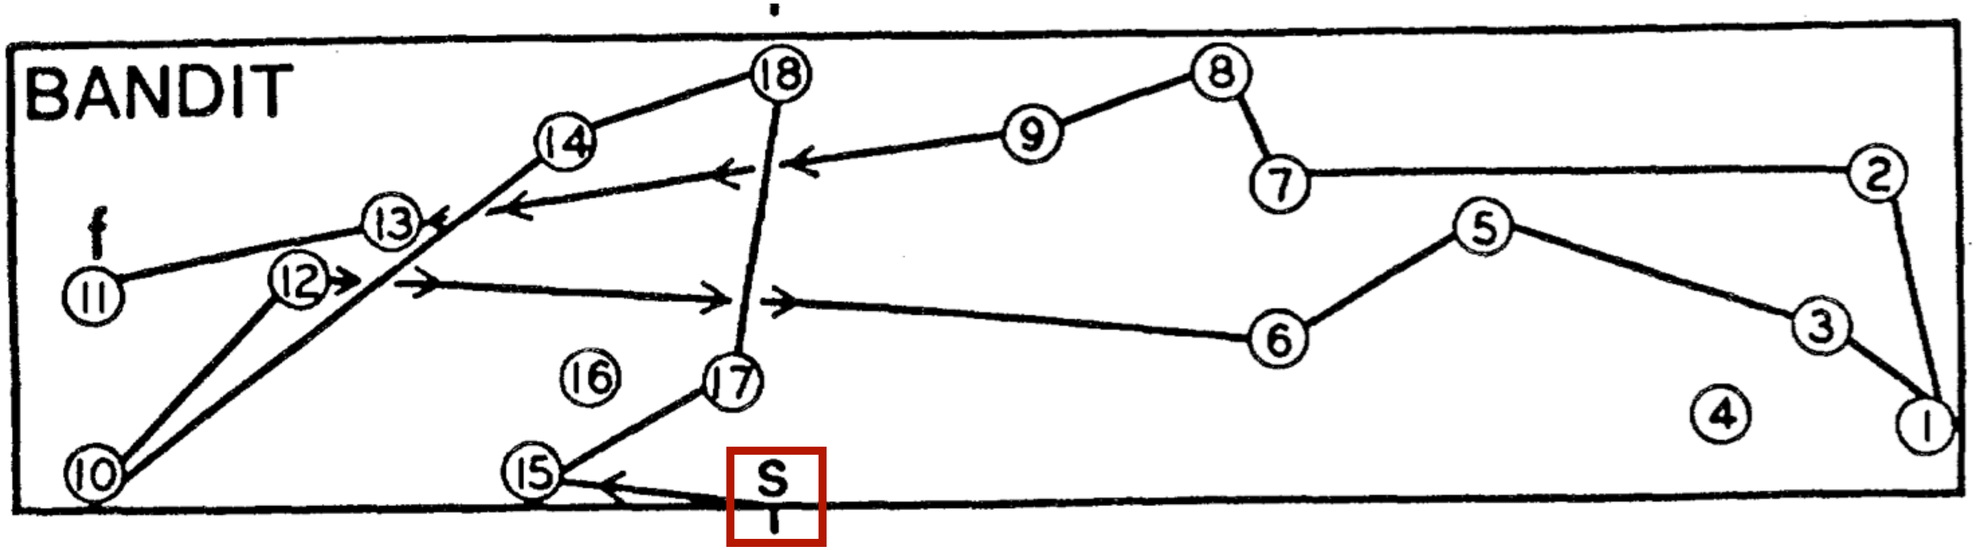

Chimpanzees

Menzel 1973

Animals build maps

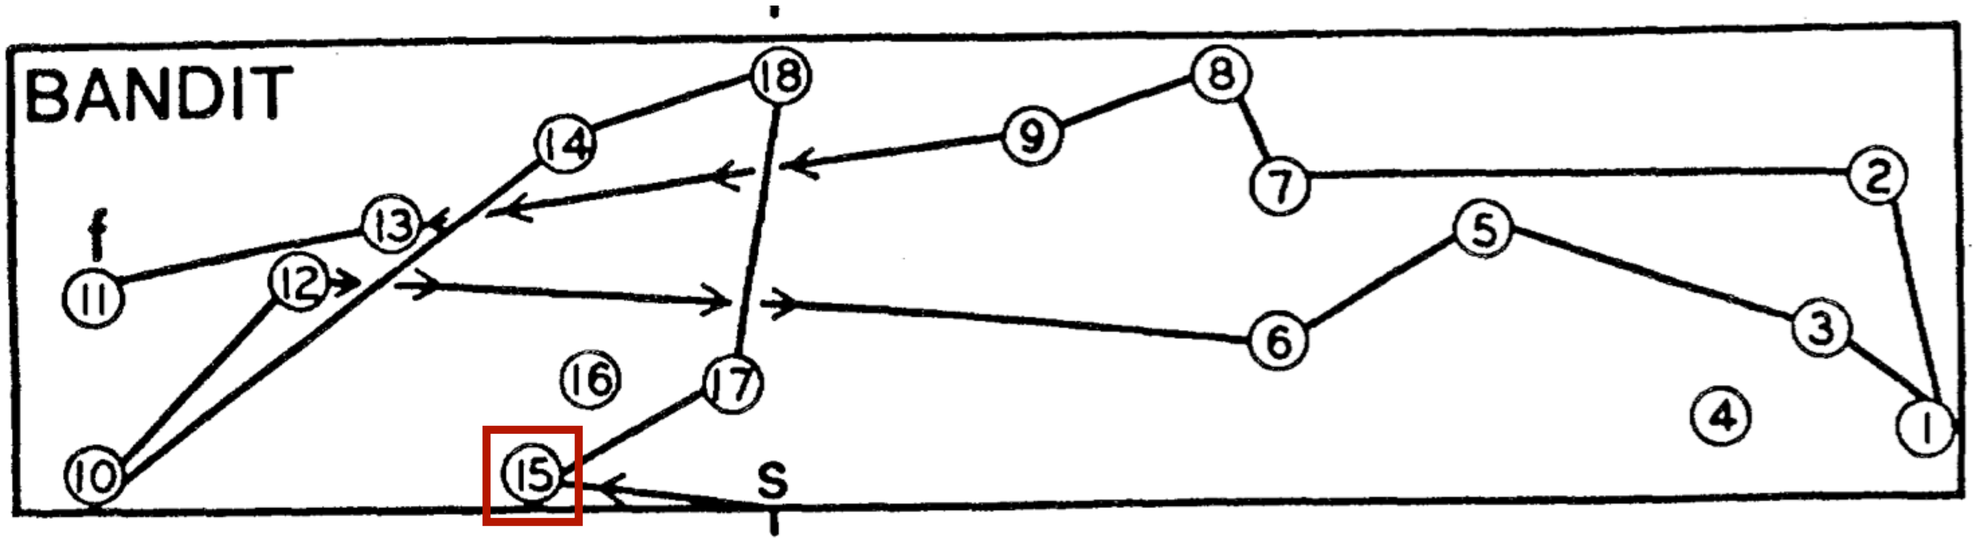

Chimpanzees

Menzel 1973

Animals build maps

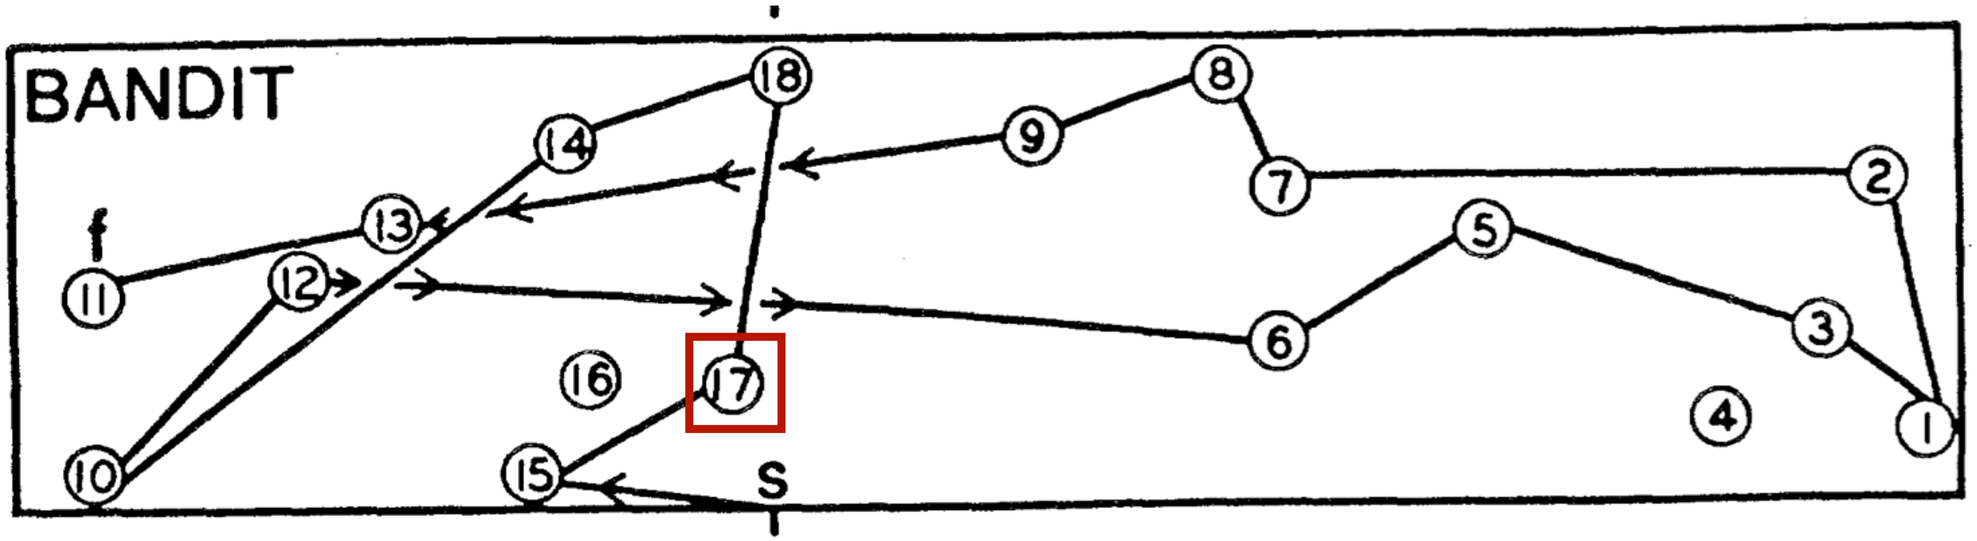

Chimpanzees

Menzel 1973

Animals build maps

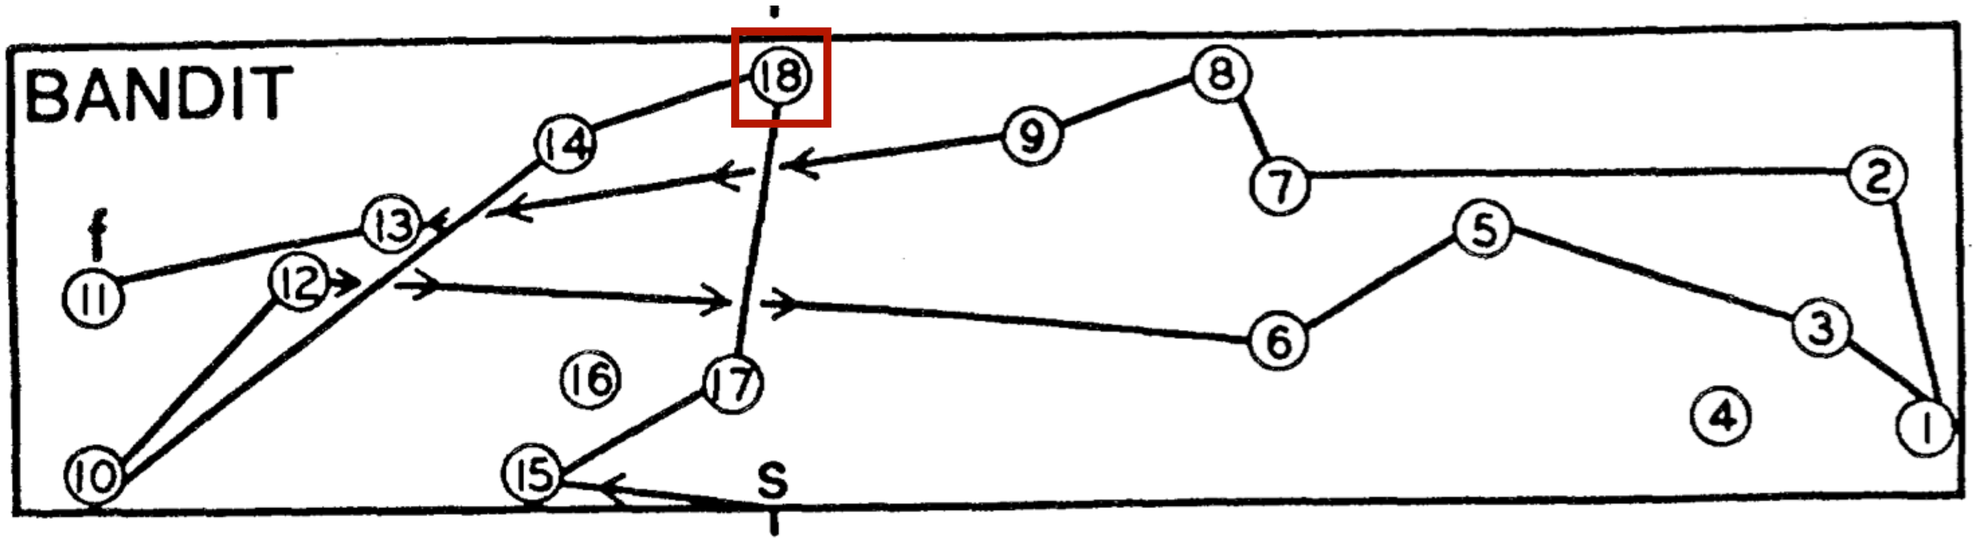

Chimpanzees

Menzel 1973

Animals build maps

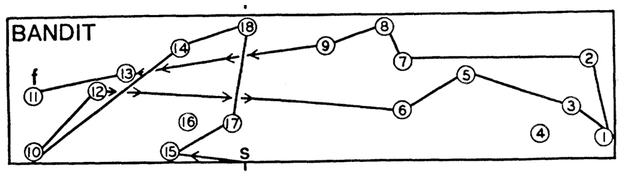

Chimpanzees

Animals build maps

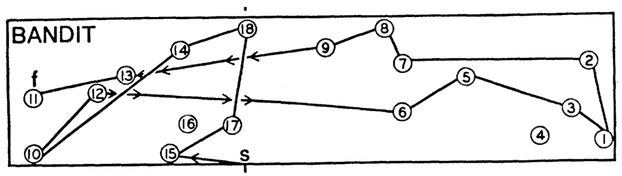

Chimpanzees

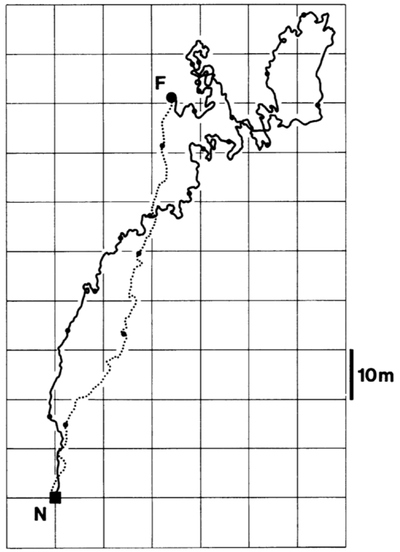

Ants

Animals build maps

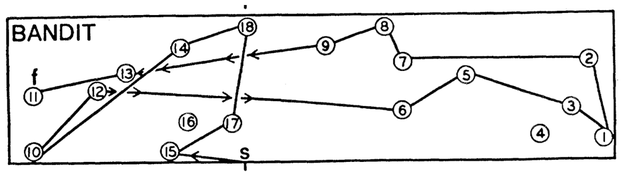

Chimpanzees

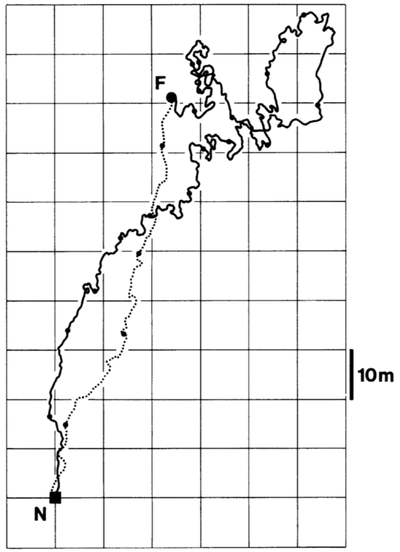

Ants

Bats

Animals build maps

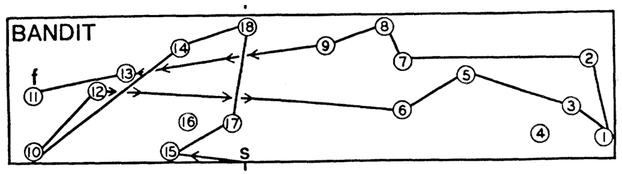

Chimpanzees

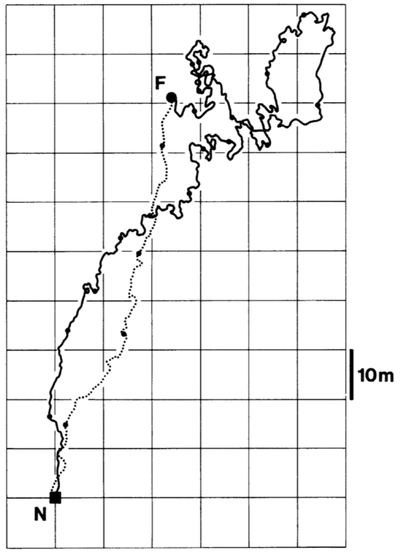

Ants

Bats

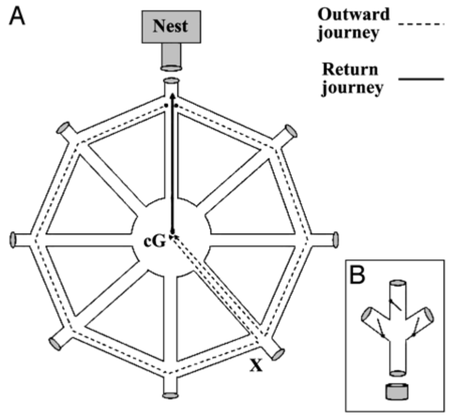

Blind mole rats

Machines use maps

Shakey the Robot, 1972

Machines use maps

Shakey the Robot, 1972

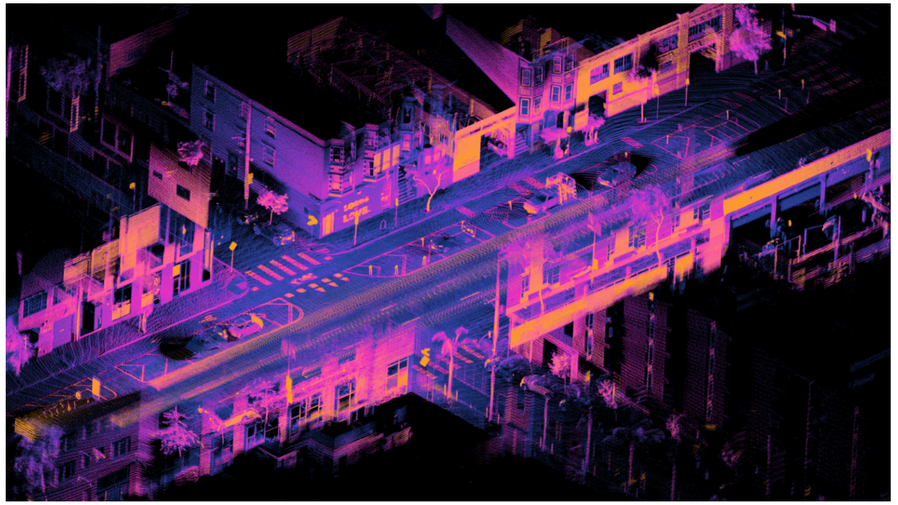

LIDAR SLAM

Machines use maps

Shakey the Robot, 1972

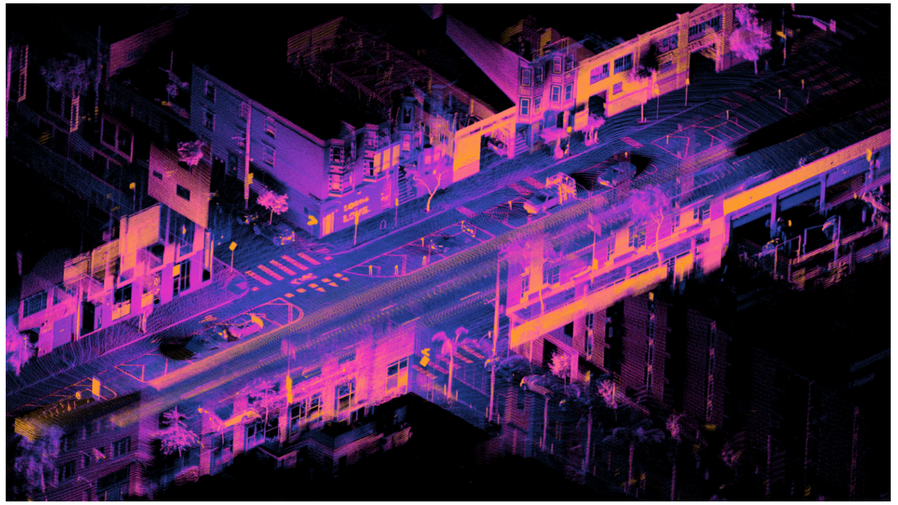

LIDAR SLAM

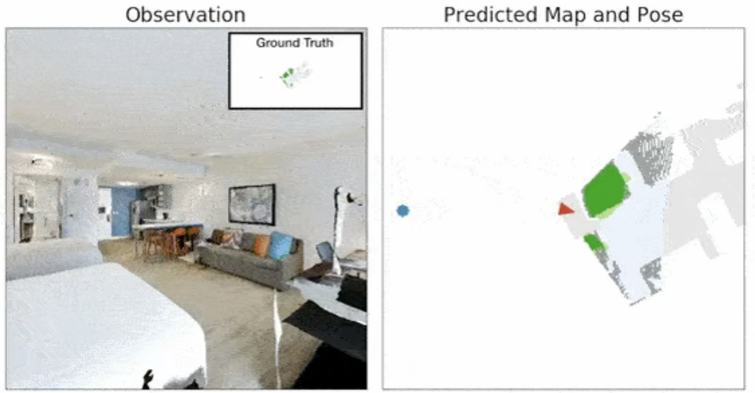

Neural SLAM

Do AI agents learn to build maps in the course of learning to navigate?

Do AI agents learn to build maps in the course of learning to navigate?

- Would shed light on the internal workings of black box AI navigation agents

Do AI agents learn to build maps in the course of learning to navigate?

- Would shed light on the internal workings of black box AI navigation agents
  - Recent results have shown high performance with ‘map-free’ navigators

DD-PPO, Wijmans et al, 2020
EmbClip: Khandelwal et al, 2022
Gato: Reed et al, 2022

Do AI agents learn to build maps in the course of learning to navigate?

- Would shed light on the internal workings of black box AI navigation agents
  - Recent results have shown high performance with ‘map-free’ navigators
- In a manner similar to convergent evolution, it would imply that maps are a natural solution to navigation

Convergent Evolution

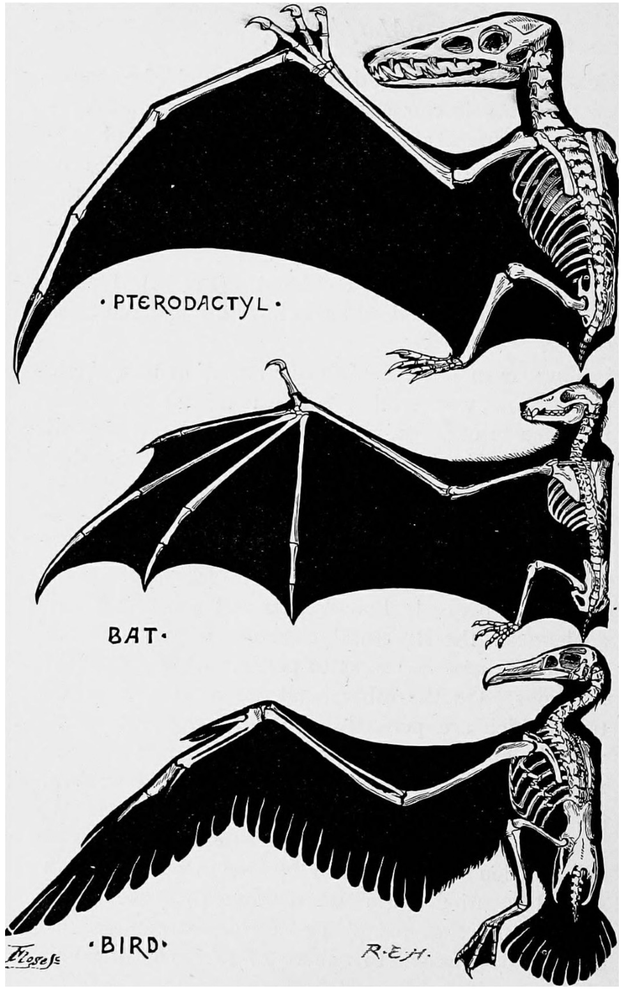

PointGoal Navigation

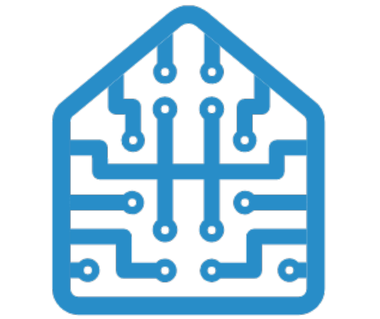

Habitat

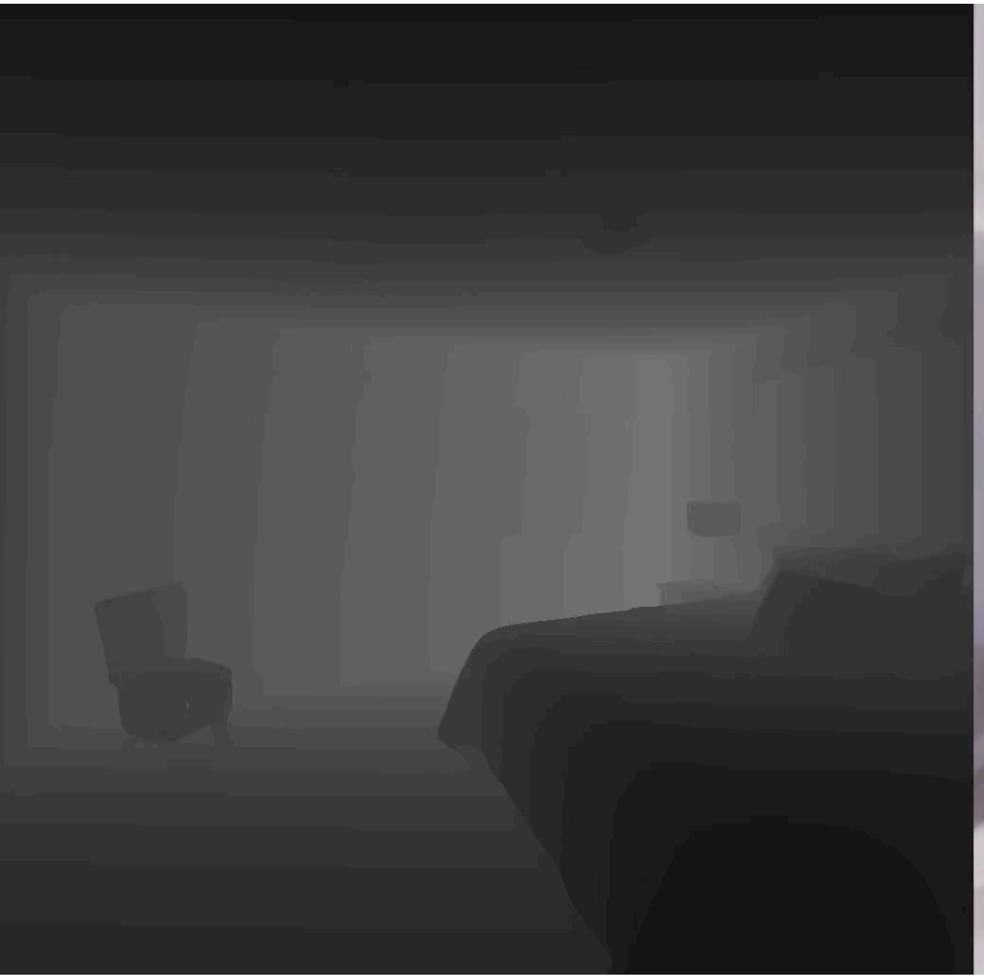

Depth

RGB and GPS+Compass

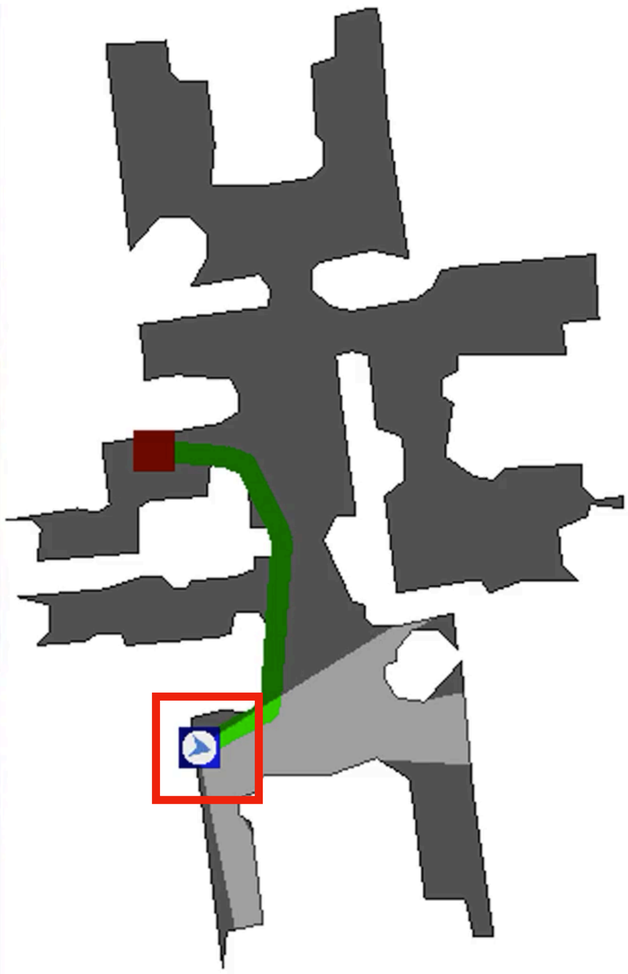

Top Down Map

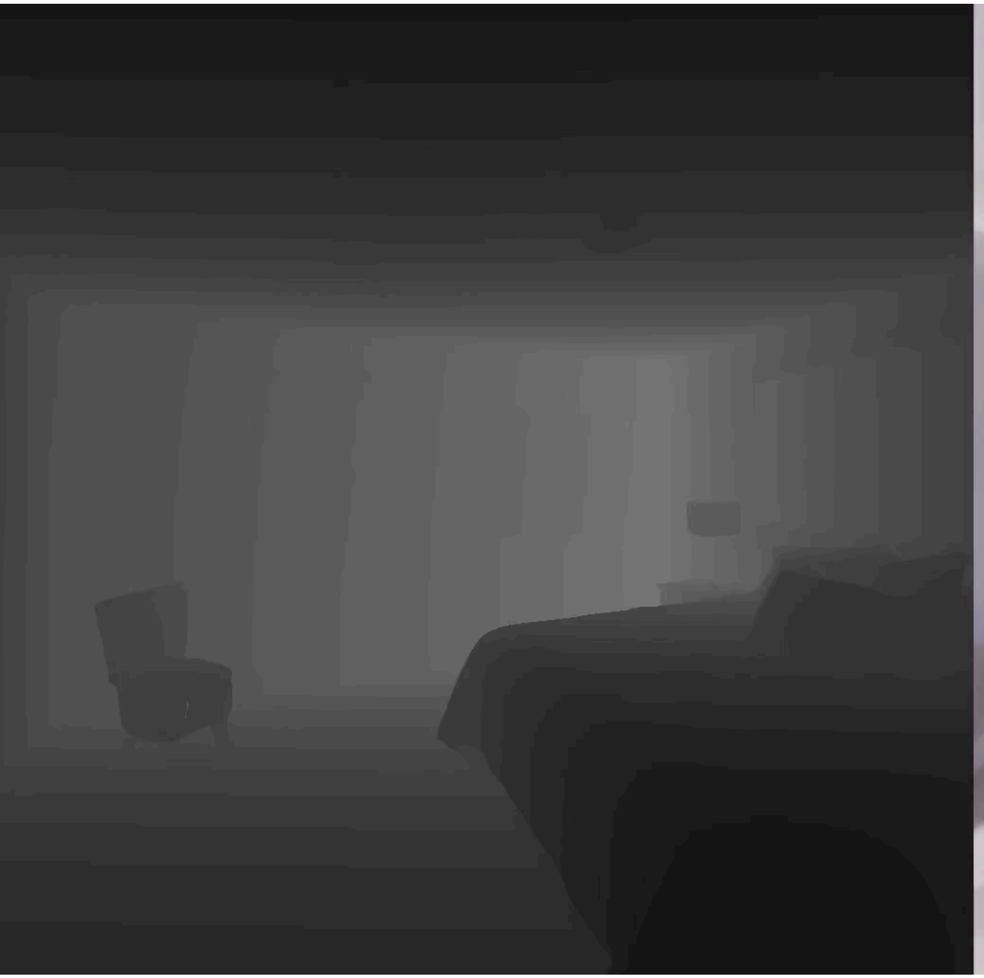

Depth

RGB and GPS+Compass

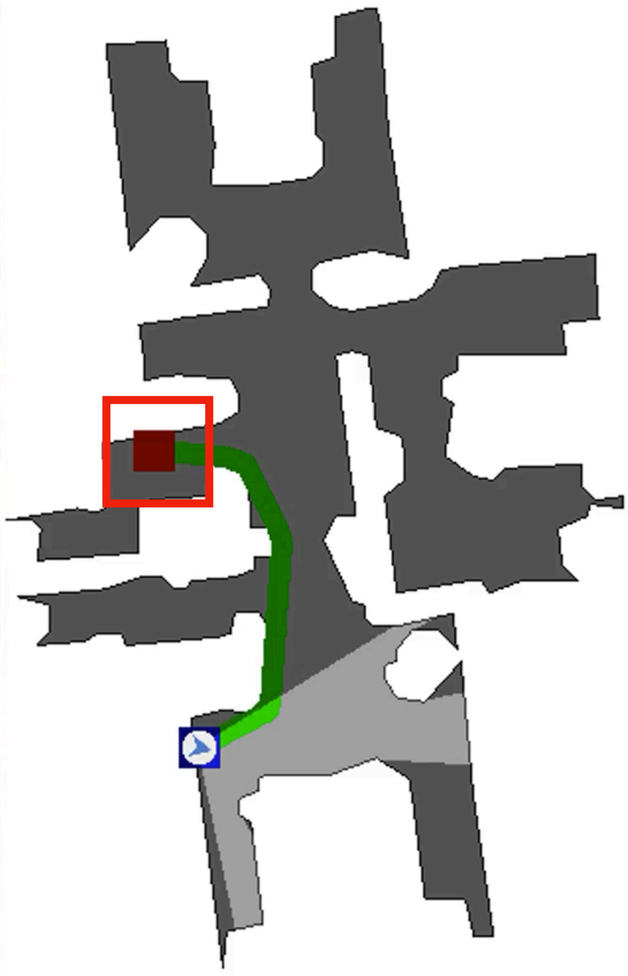

Top Down Map

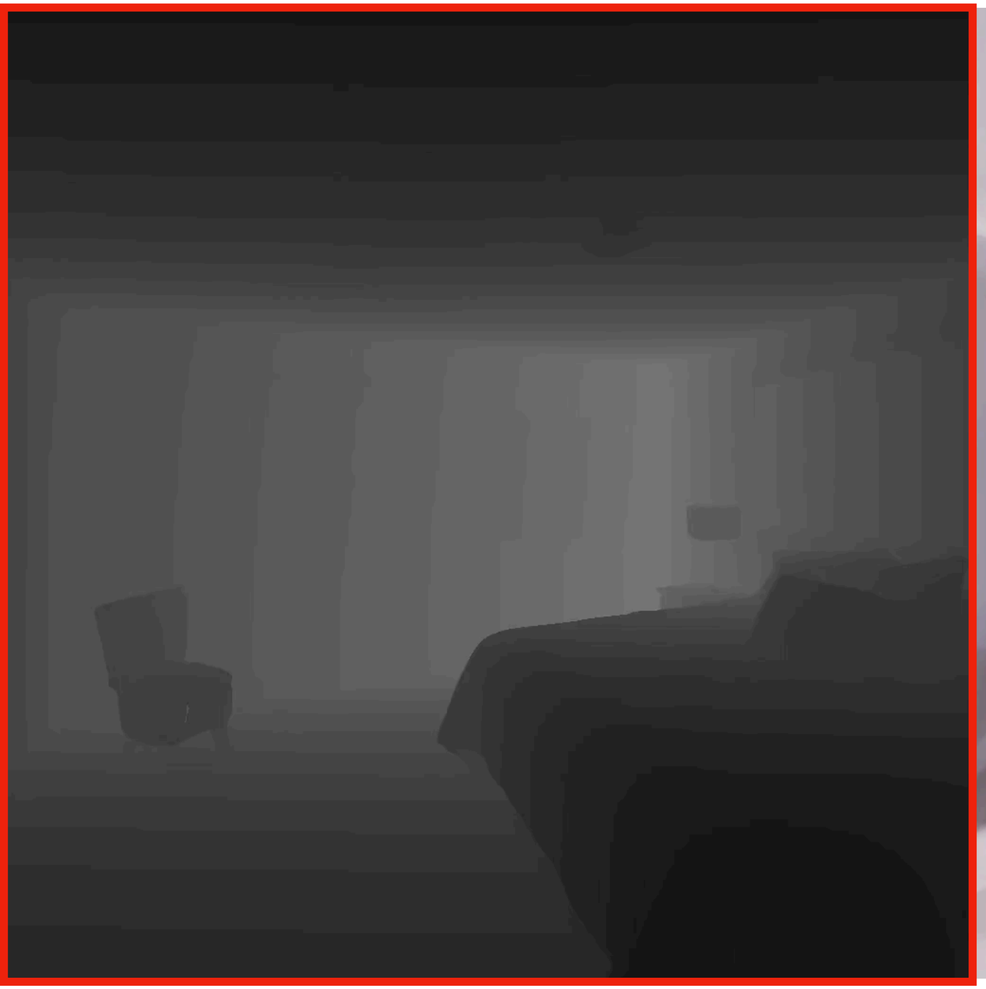

Depth

RGB and GPS+Compass

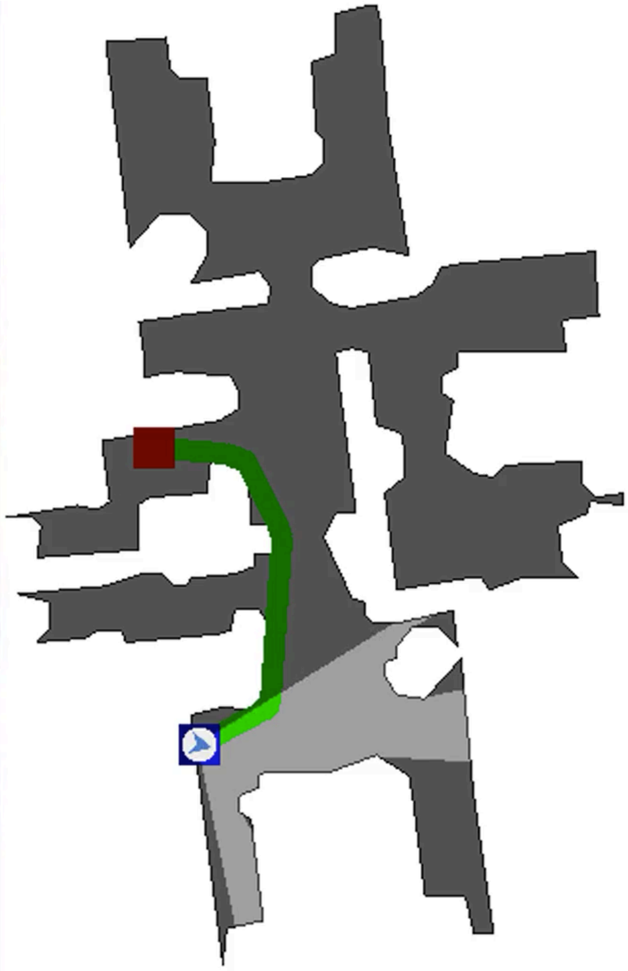

Top Down Map

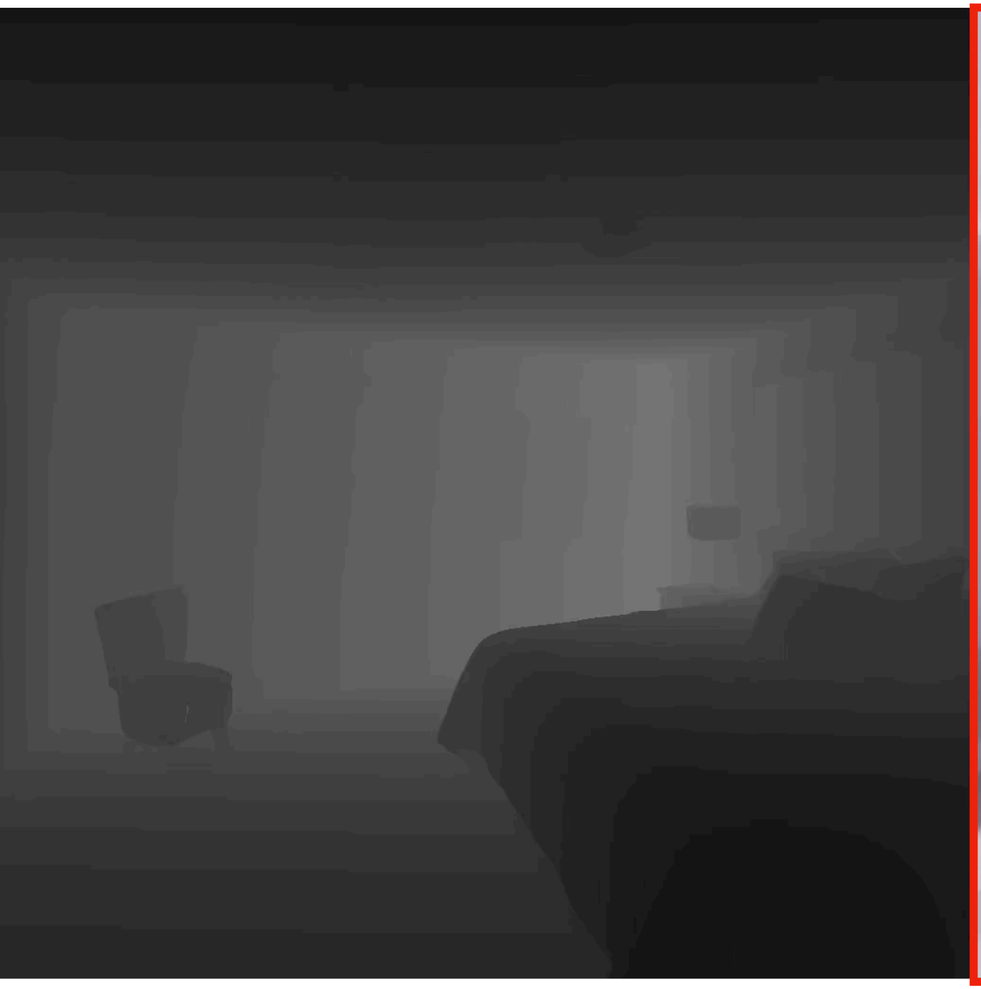

Depth

RGB and GPS+Compass

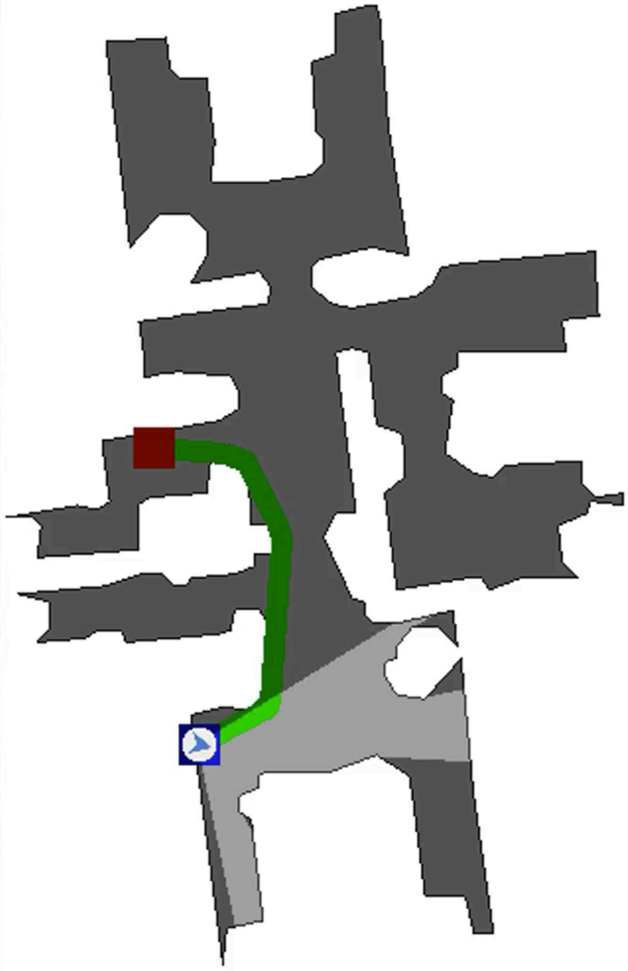

Top Down Map

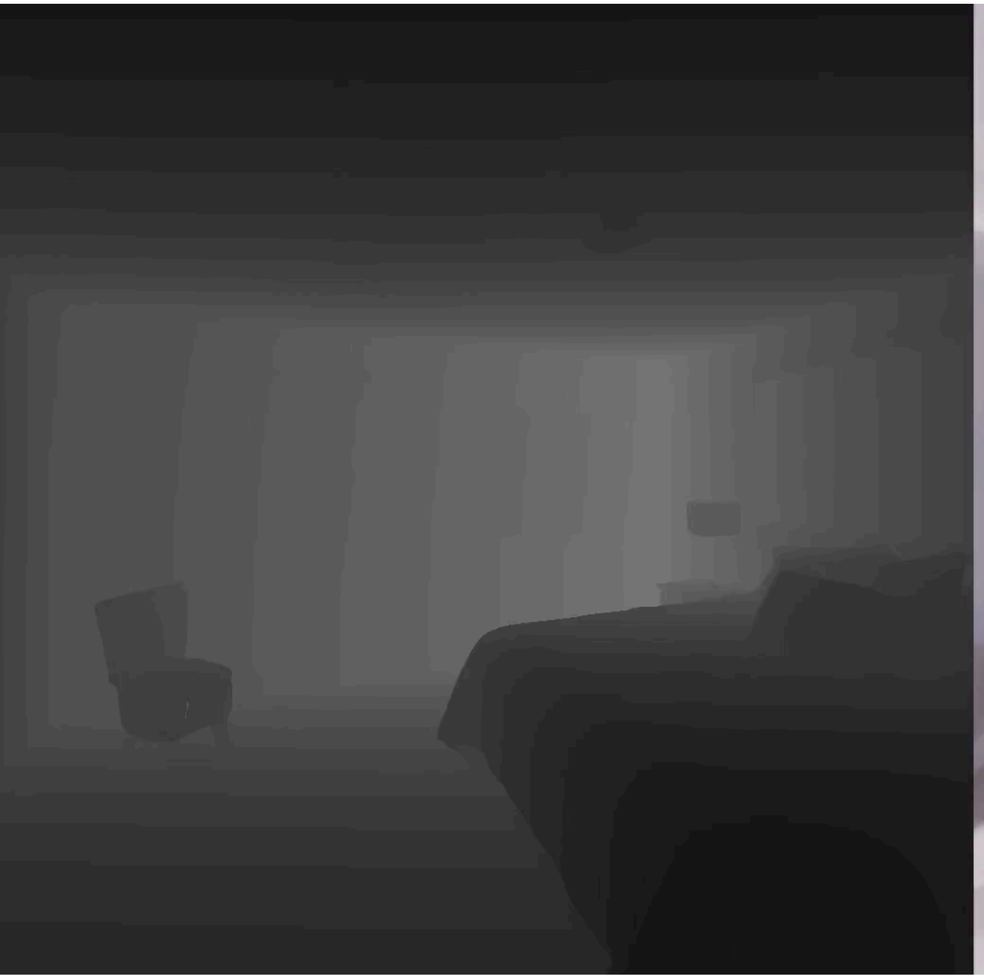

Depth

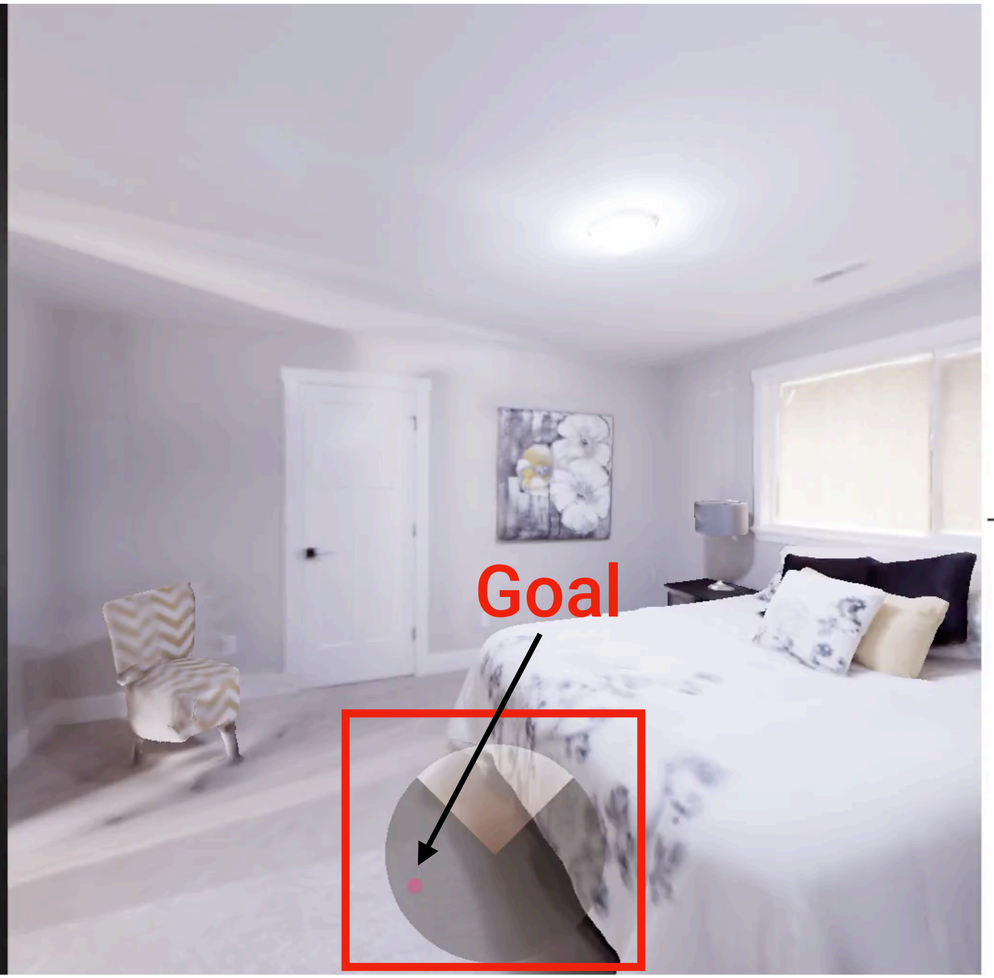

RGB and GPS+Compass

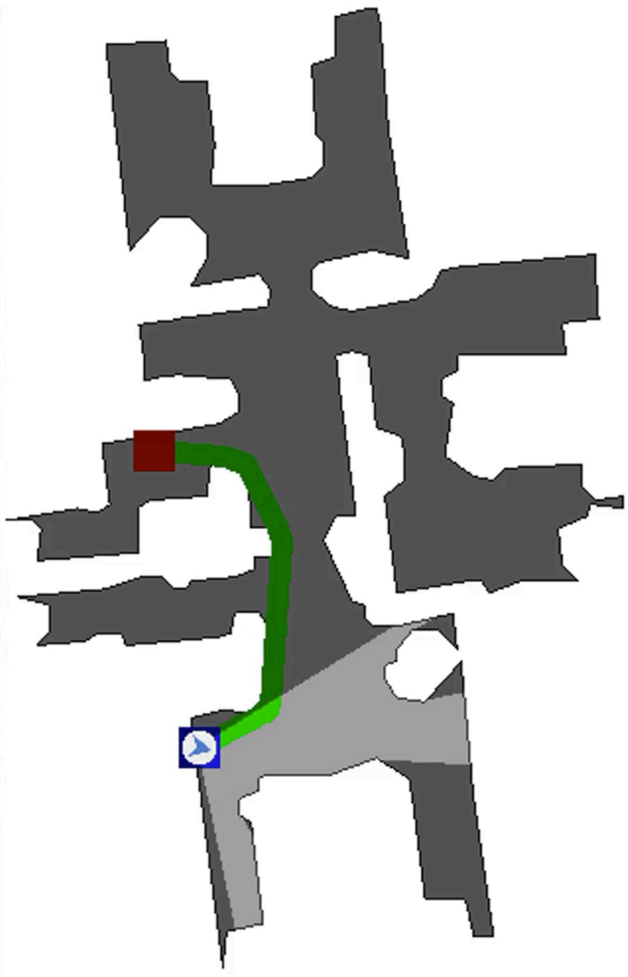

Top Down Map

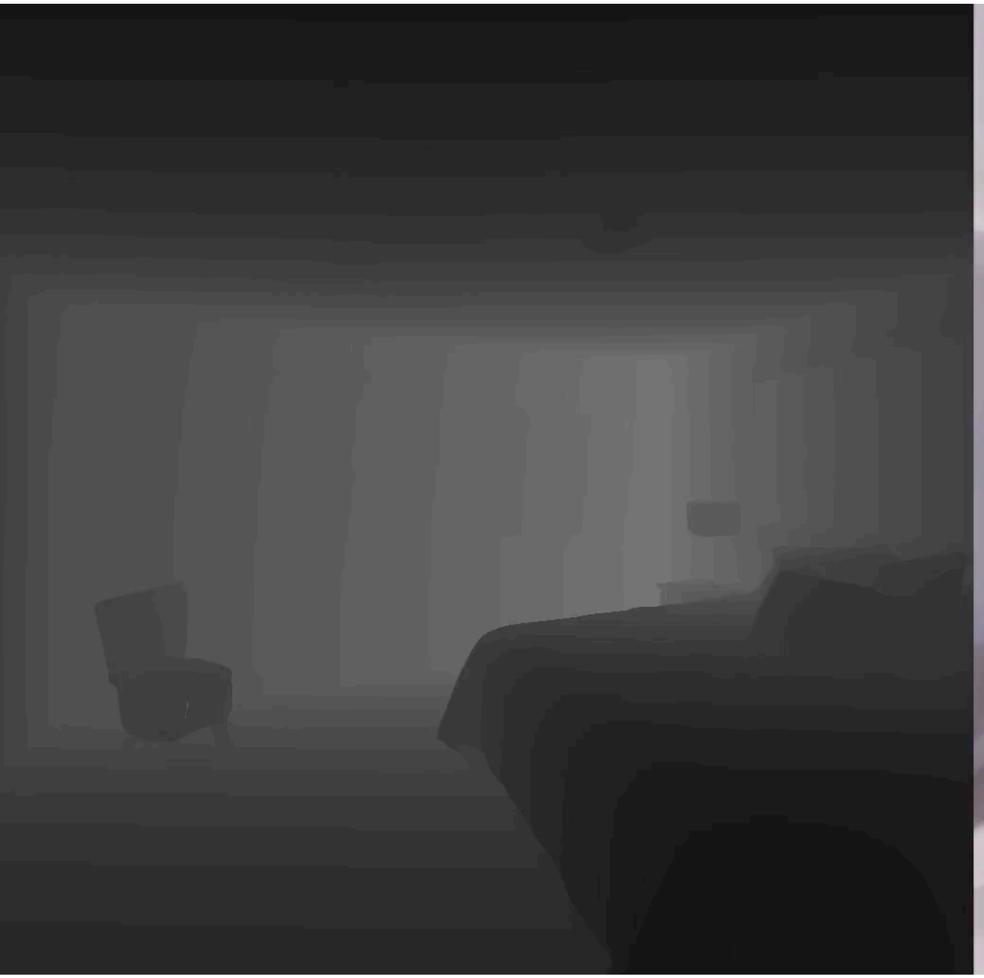

Depth

RGB and GPS+Compass

Top Down Map

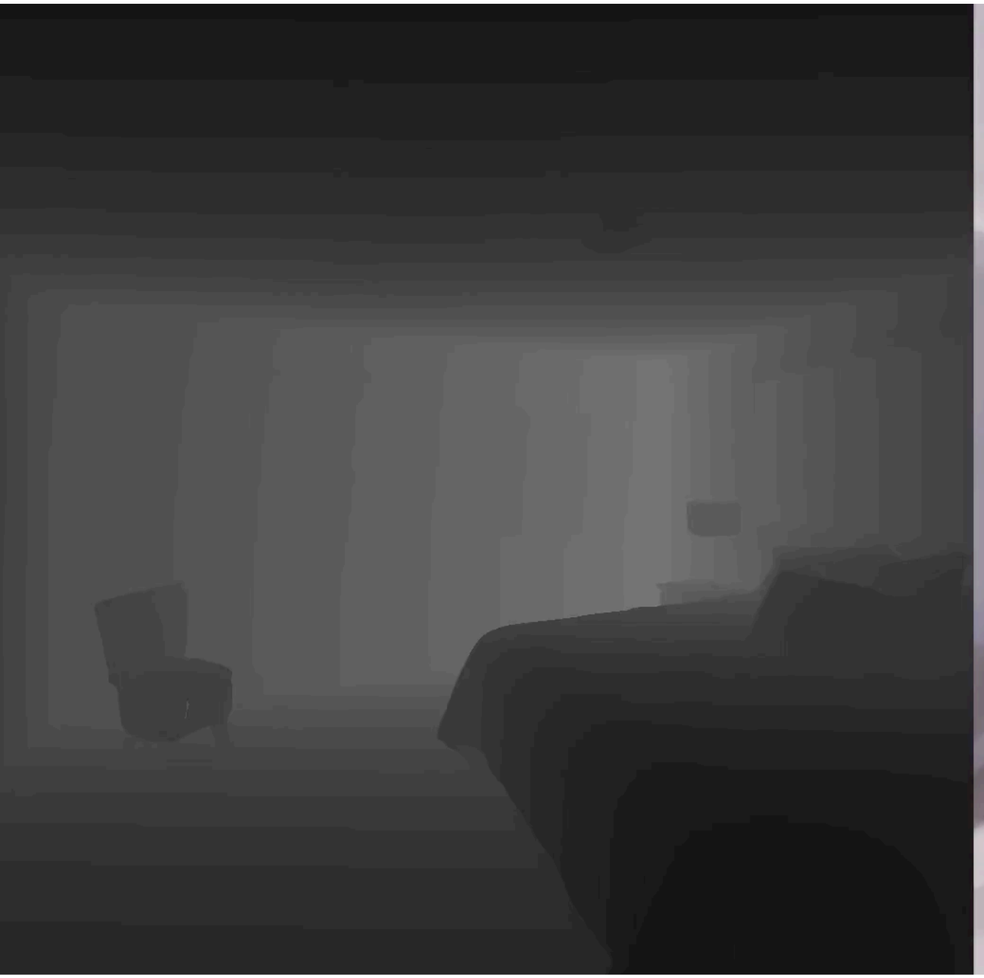

Depth

RGB and GPS+Compass

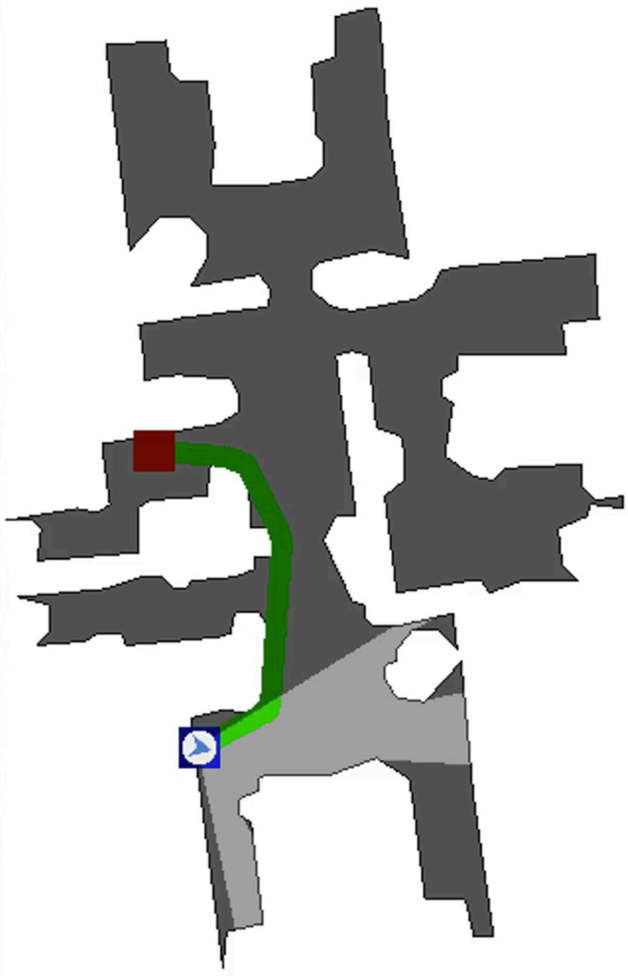

Top Down Map

Similar to Honey Bee Waggle Dance

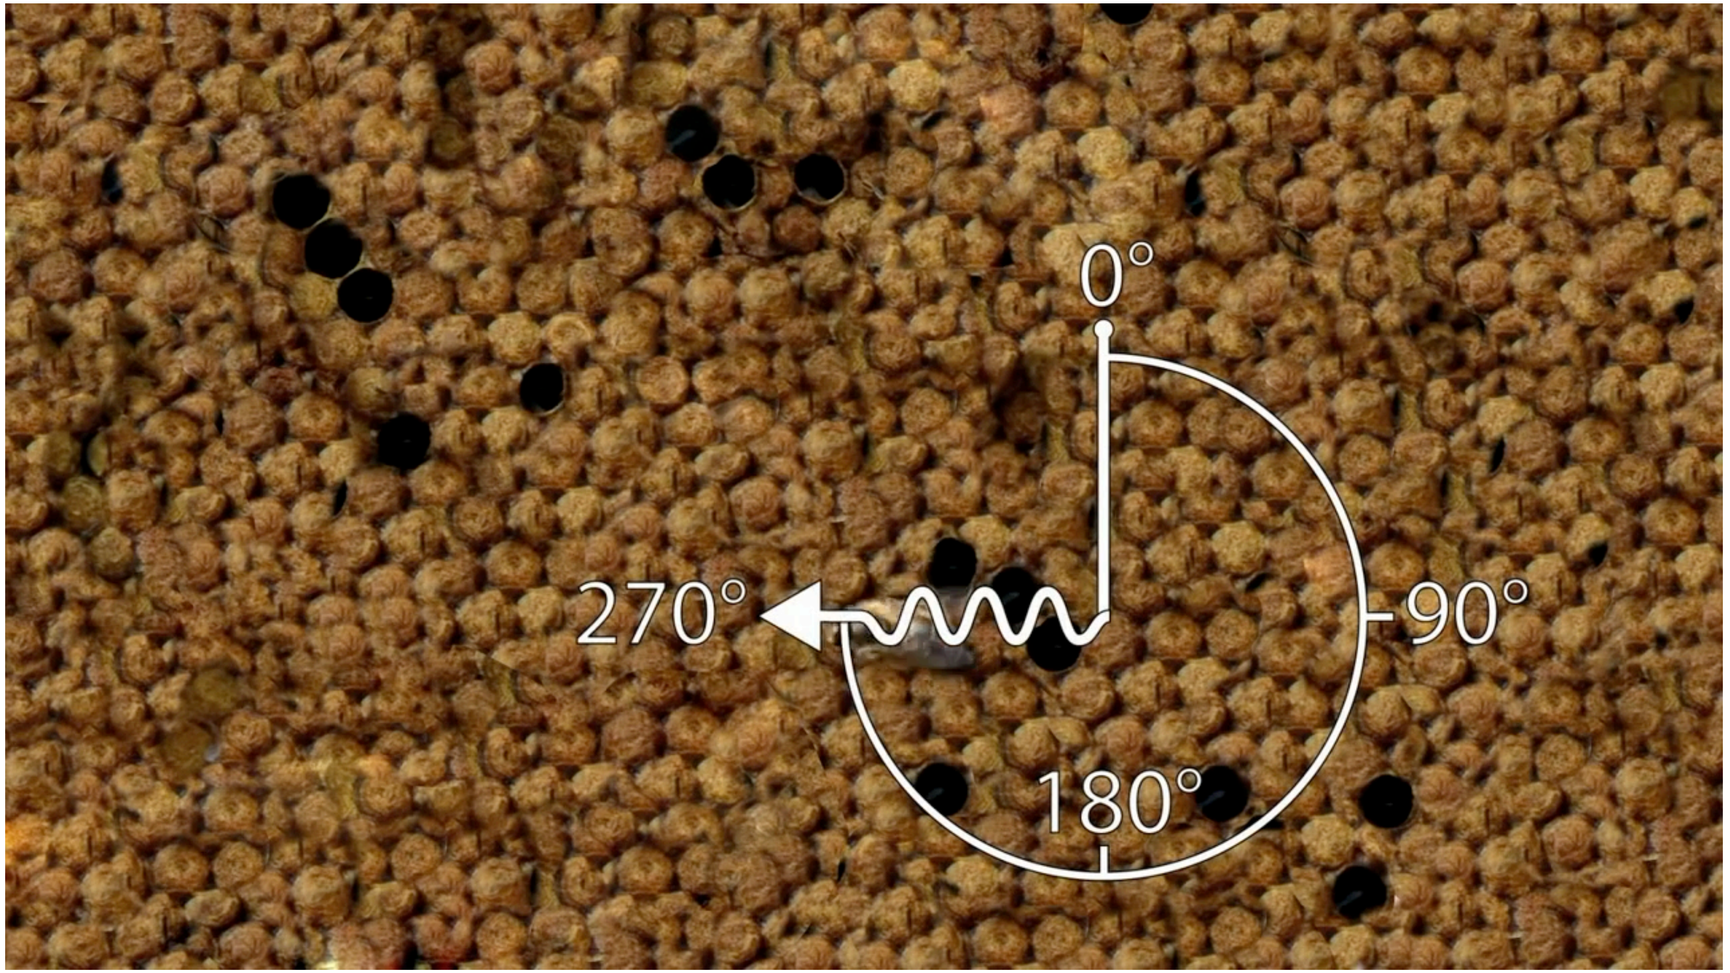

Video Credit: The Bee Group @ VT (https://www.freelyflyingbees.com/beegroupvt.html)

Causality

Causality

- How can we be sure that the agent *uses* its map?

Causality

- How can we be sure that the agent *uses* its map?
- Similar to issues in animal navigation

Causality

- How can we be sure that the agent *uses* its map?
- Similar to issues in animal navigation
  - Maybe landmark detection

Causality

- How can we be sure that the agent *uses* its map?
- Similar to issues in animal navigation
  - Maybe landmark detection
  - Maybe smell

Causality

- How can we be sure that the agent *uses* its map?
- Similar to issues in animal navigation
  - Maybe landmark detection
  - Maybe smell
  - Maybe sun direction

Causality

- How can we be sure that the agent *uses* its map?
- Similar to issues in animal navigation
  - Maybe landmark detection
  - Maybe smell
  - Maybe sun direction
  - Maybe the earth’s magnetic field

Depth

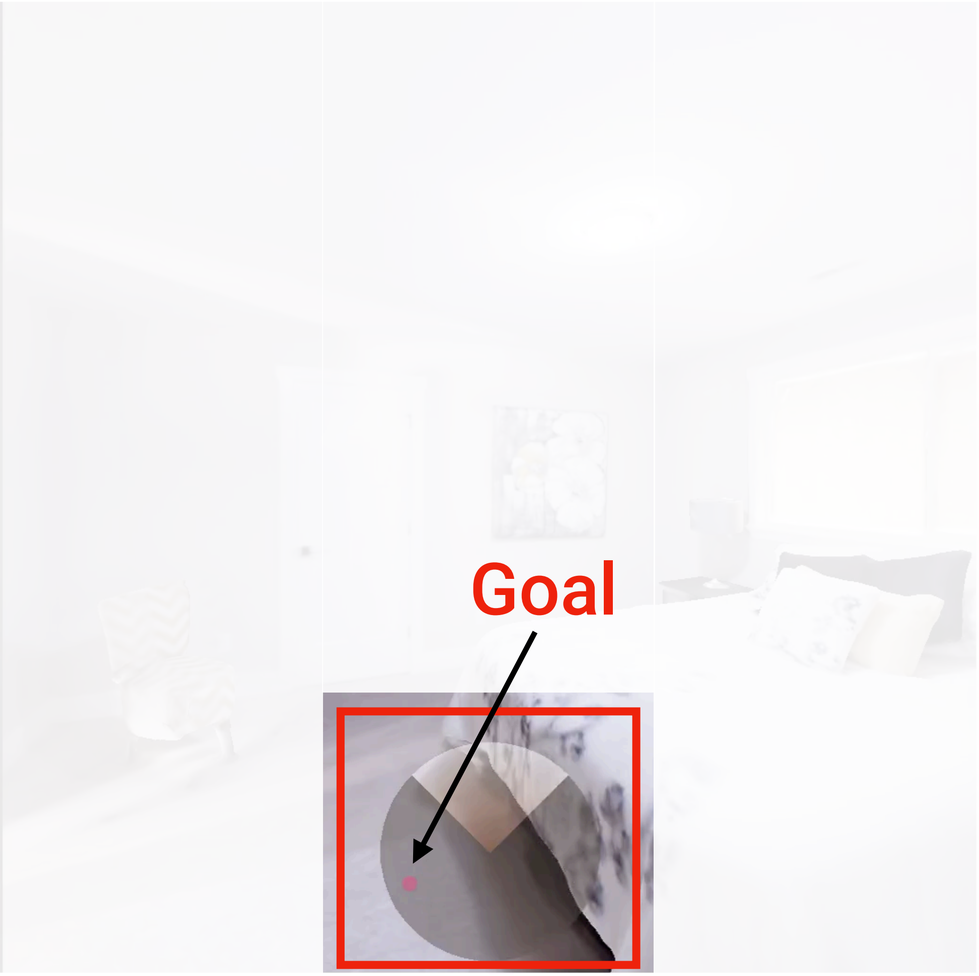

RGB and GPS+Compass

Top Down Map

Agent Model Design

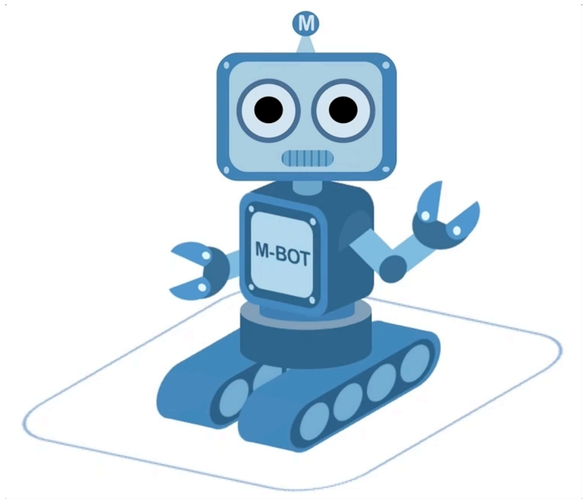

Agent Model Design

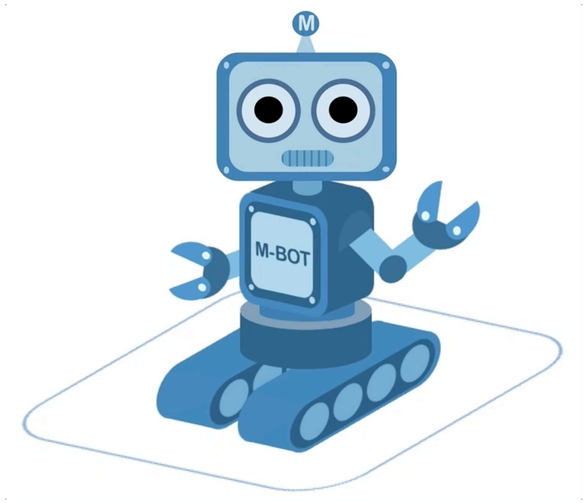

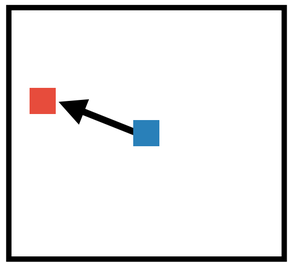

Agent Model Design

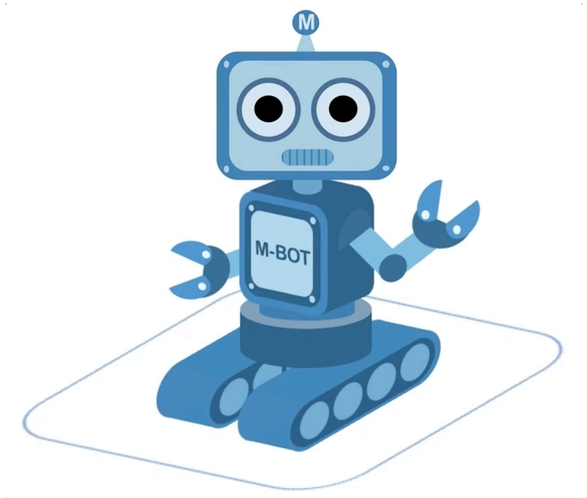

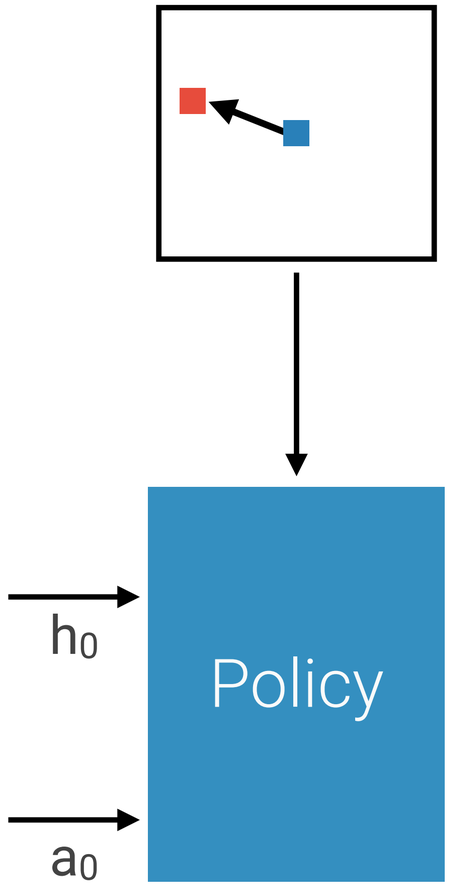

Agent Model Design

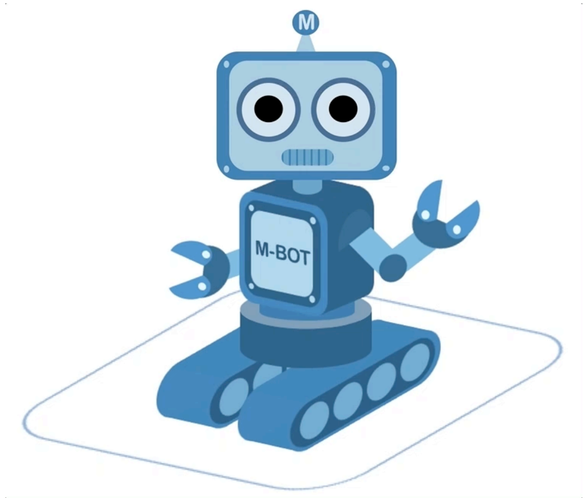

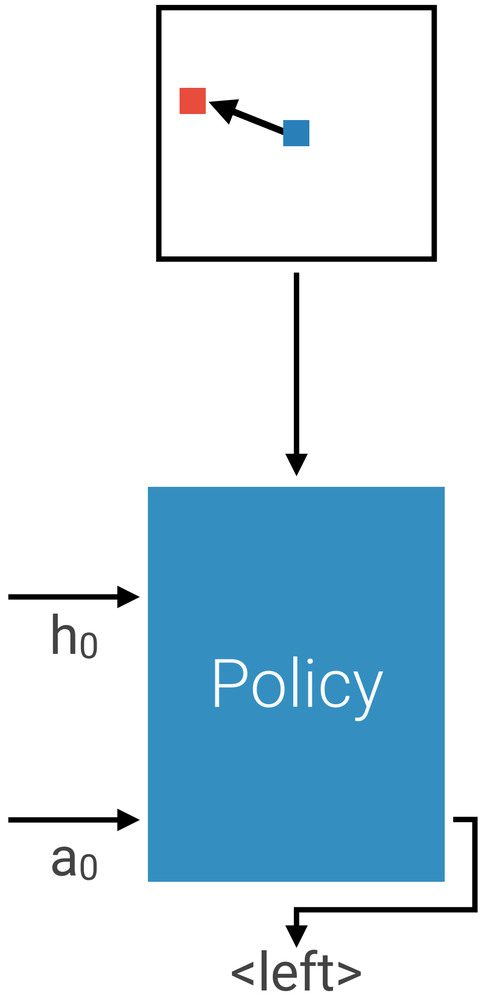

Agent Model Design

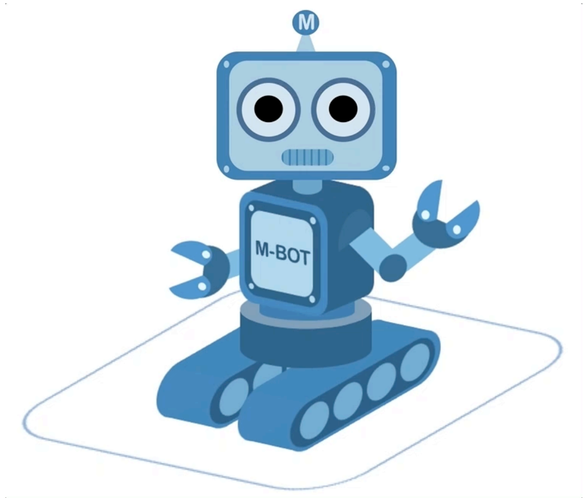

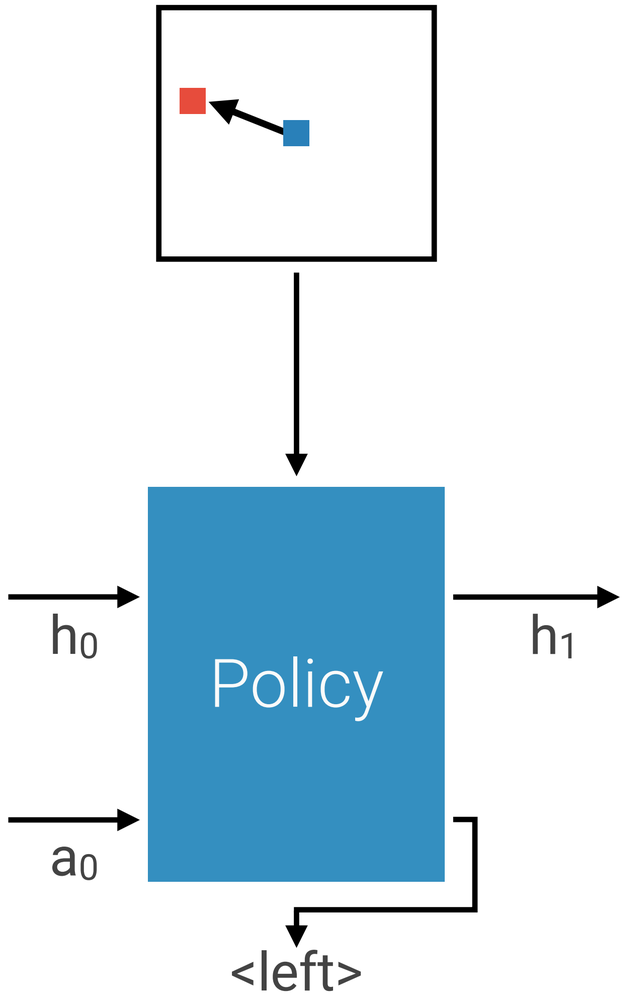

Agent Model Design

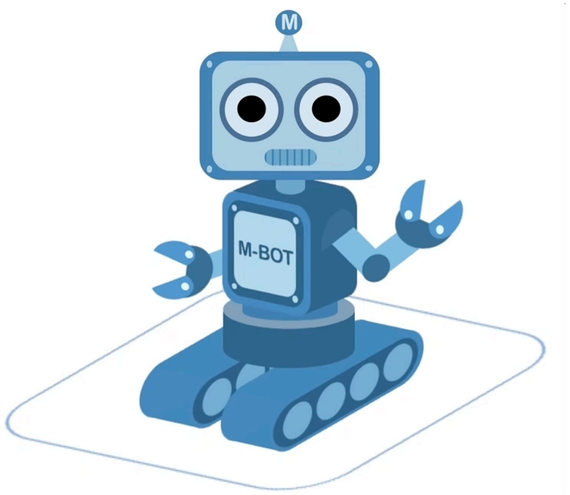

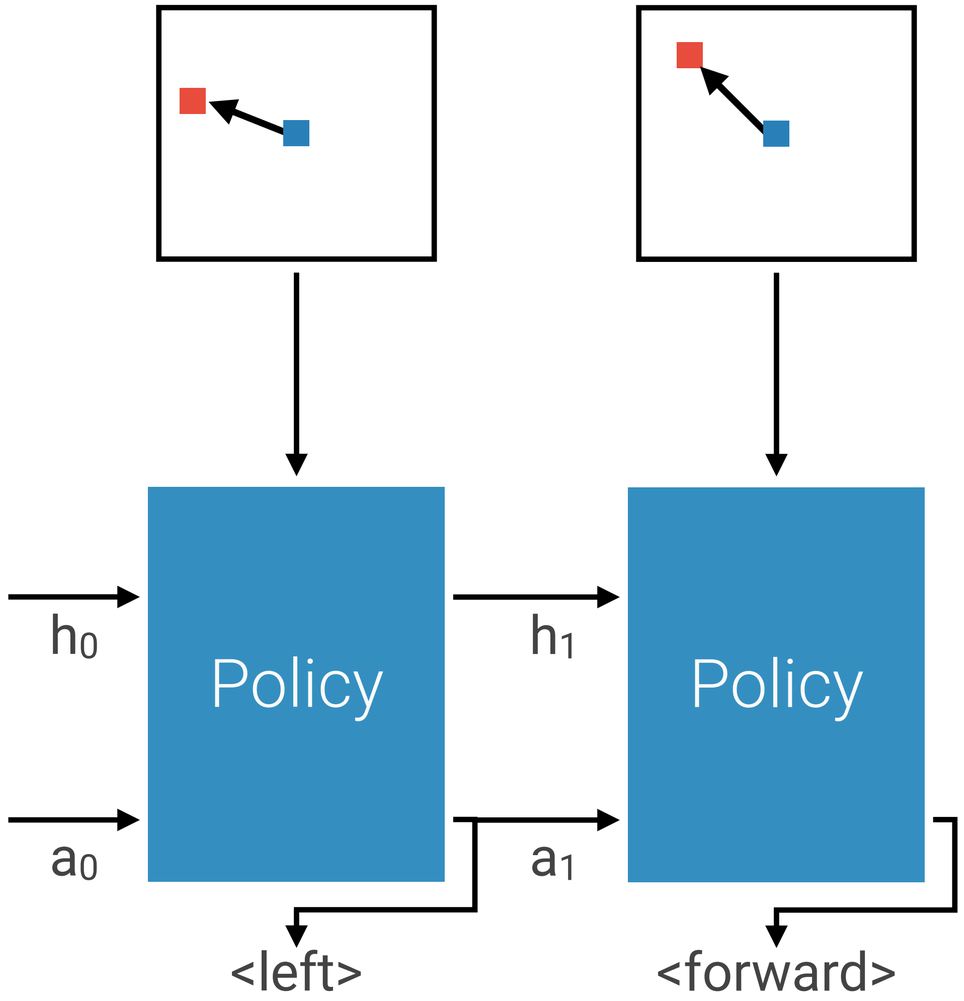

Agent Model Design

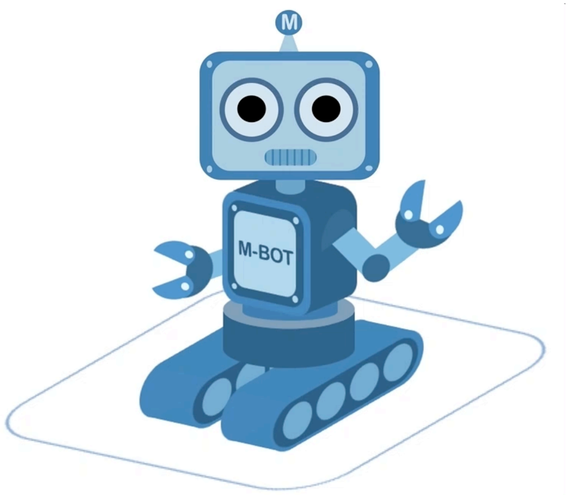

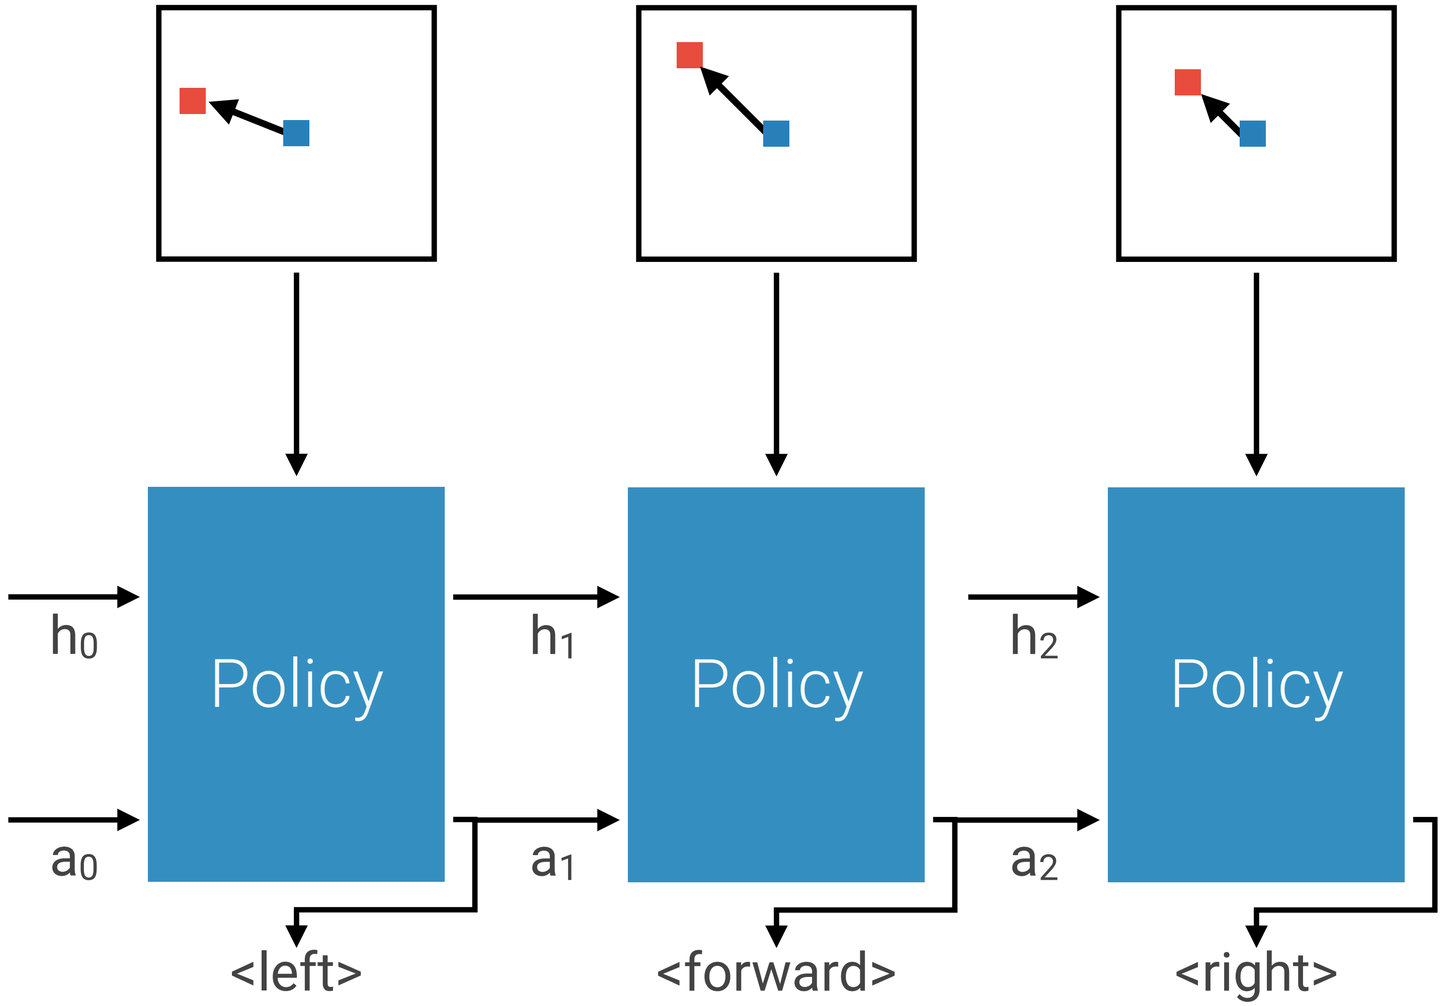

Agent Model Design

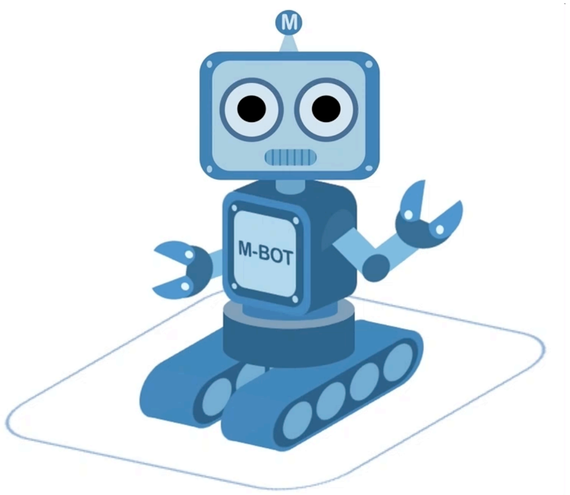

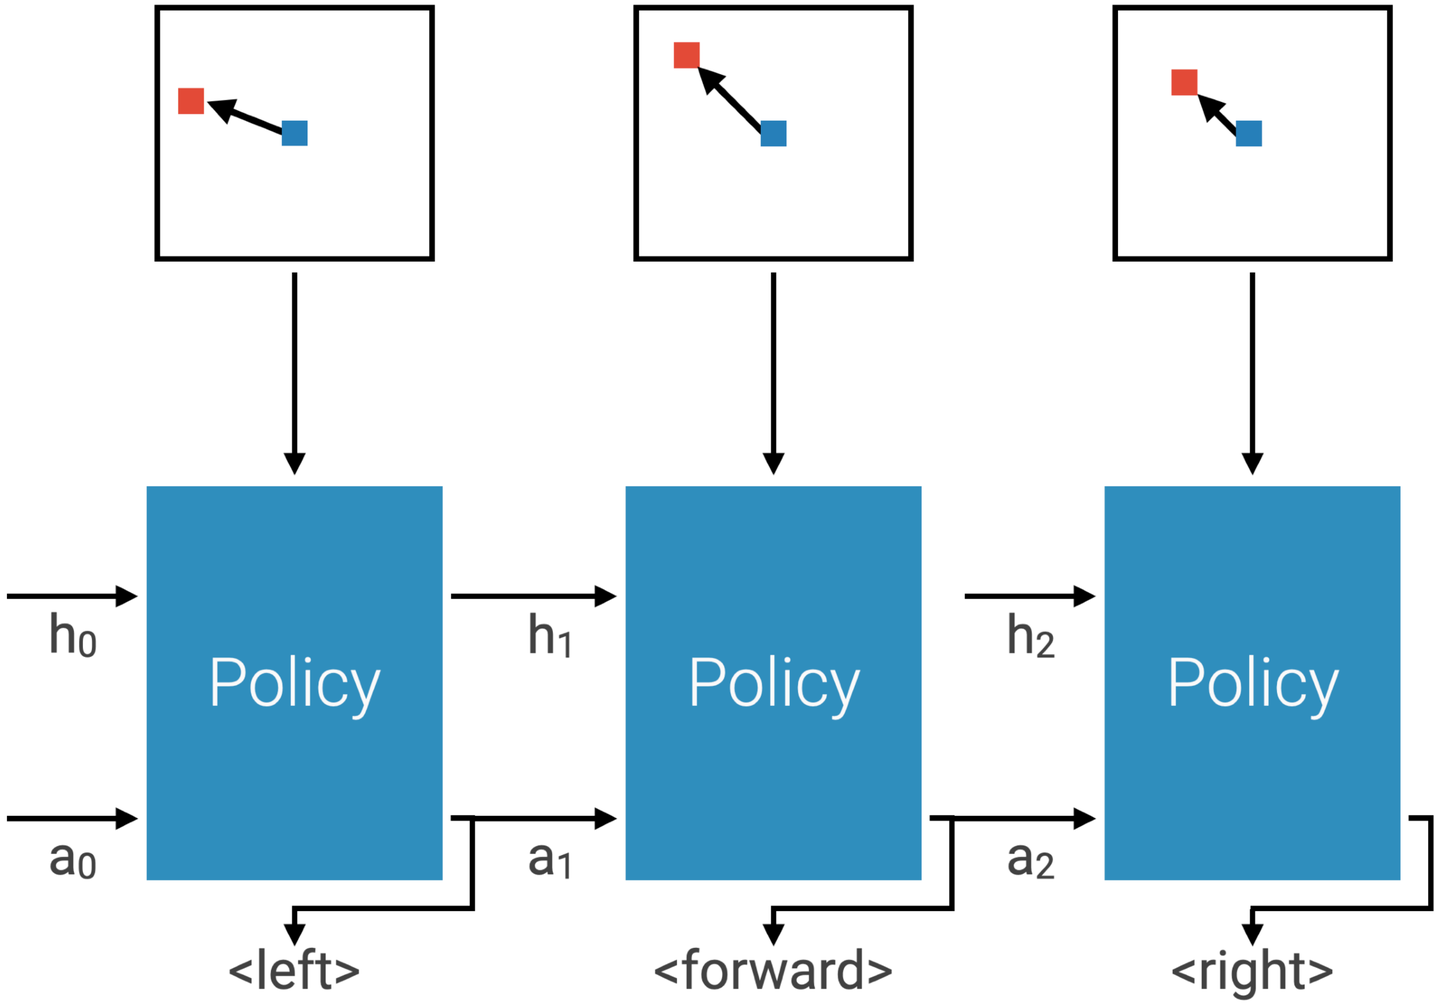

Training

- Training with generic on-policy RL (Proximal Policy Optimization)

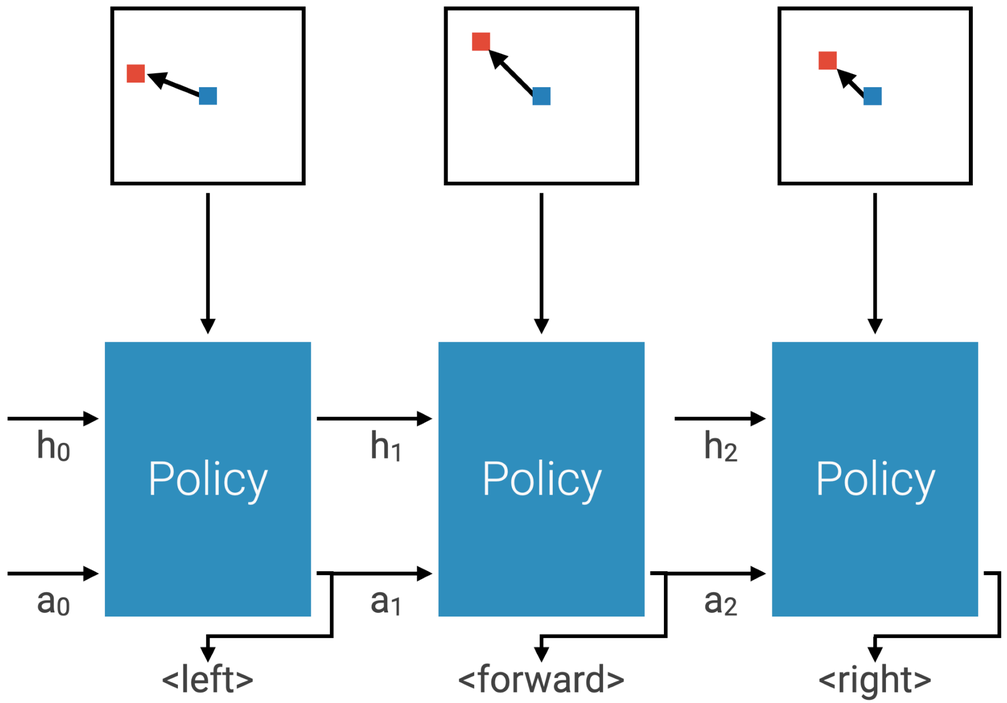

Training

- Training with generic on-policy RL (Proximal Policy Optimization)
- Both architecture and training regime are generic

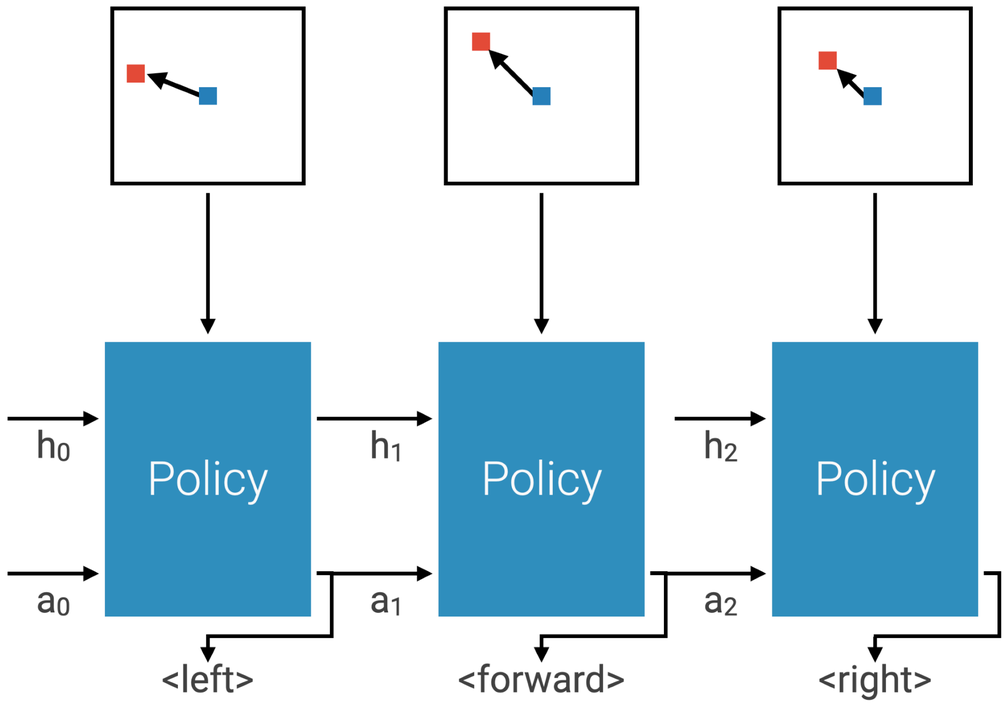

Effective navigation with only egomotion sensing

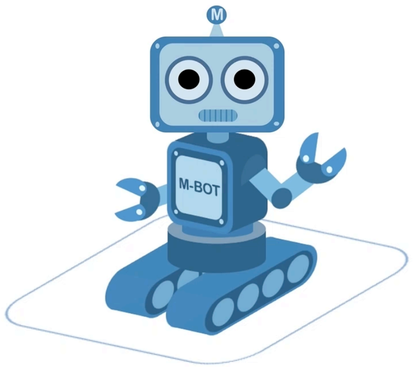

Effective navigation with only egomotion sensing

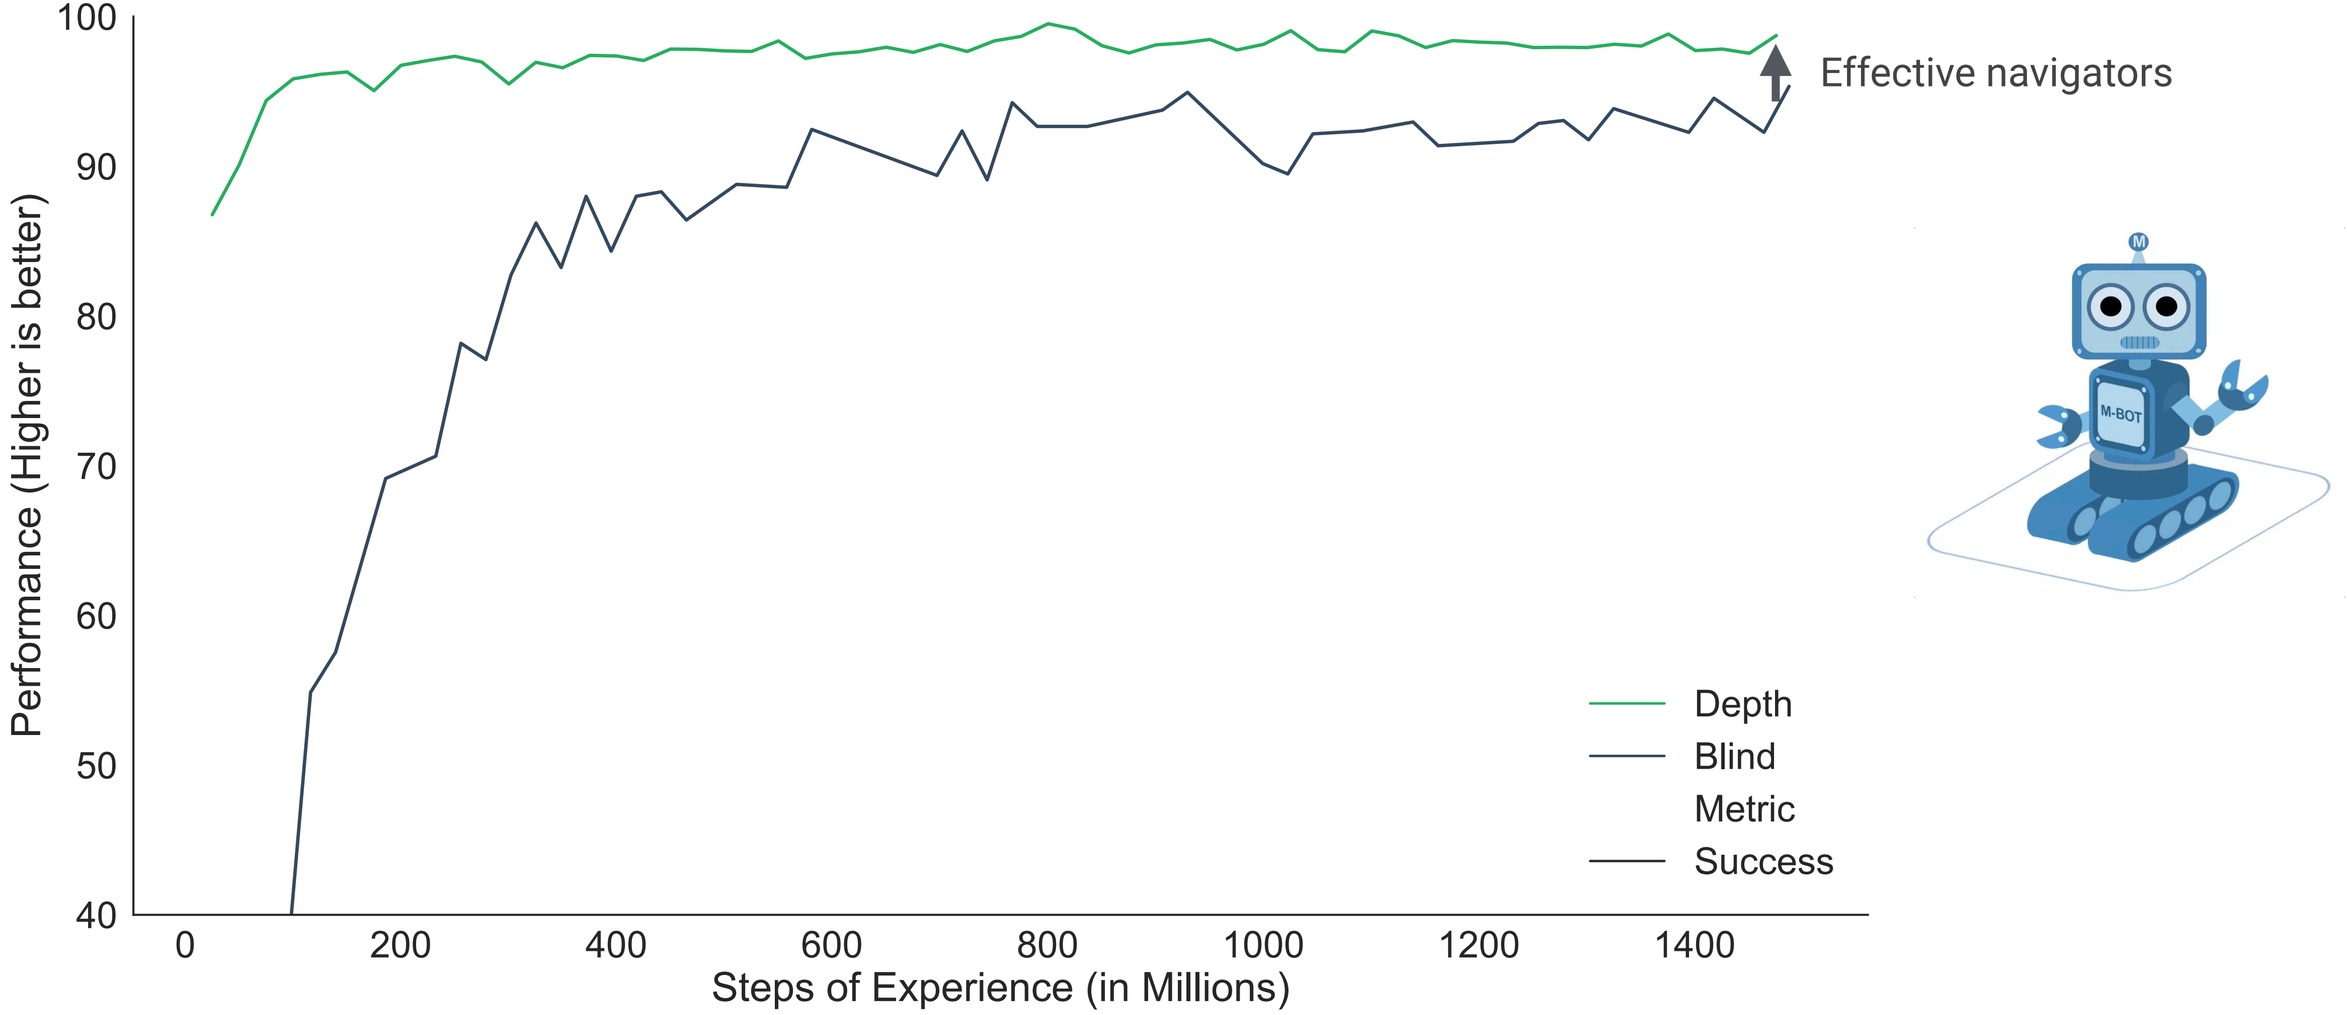

Effective navigation with only egomotion sensing

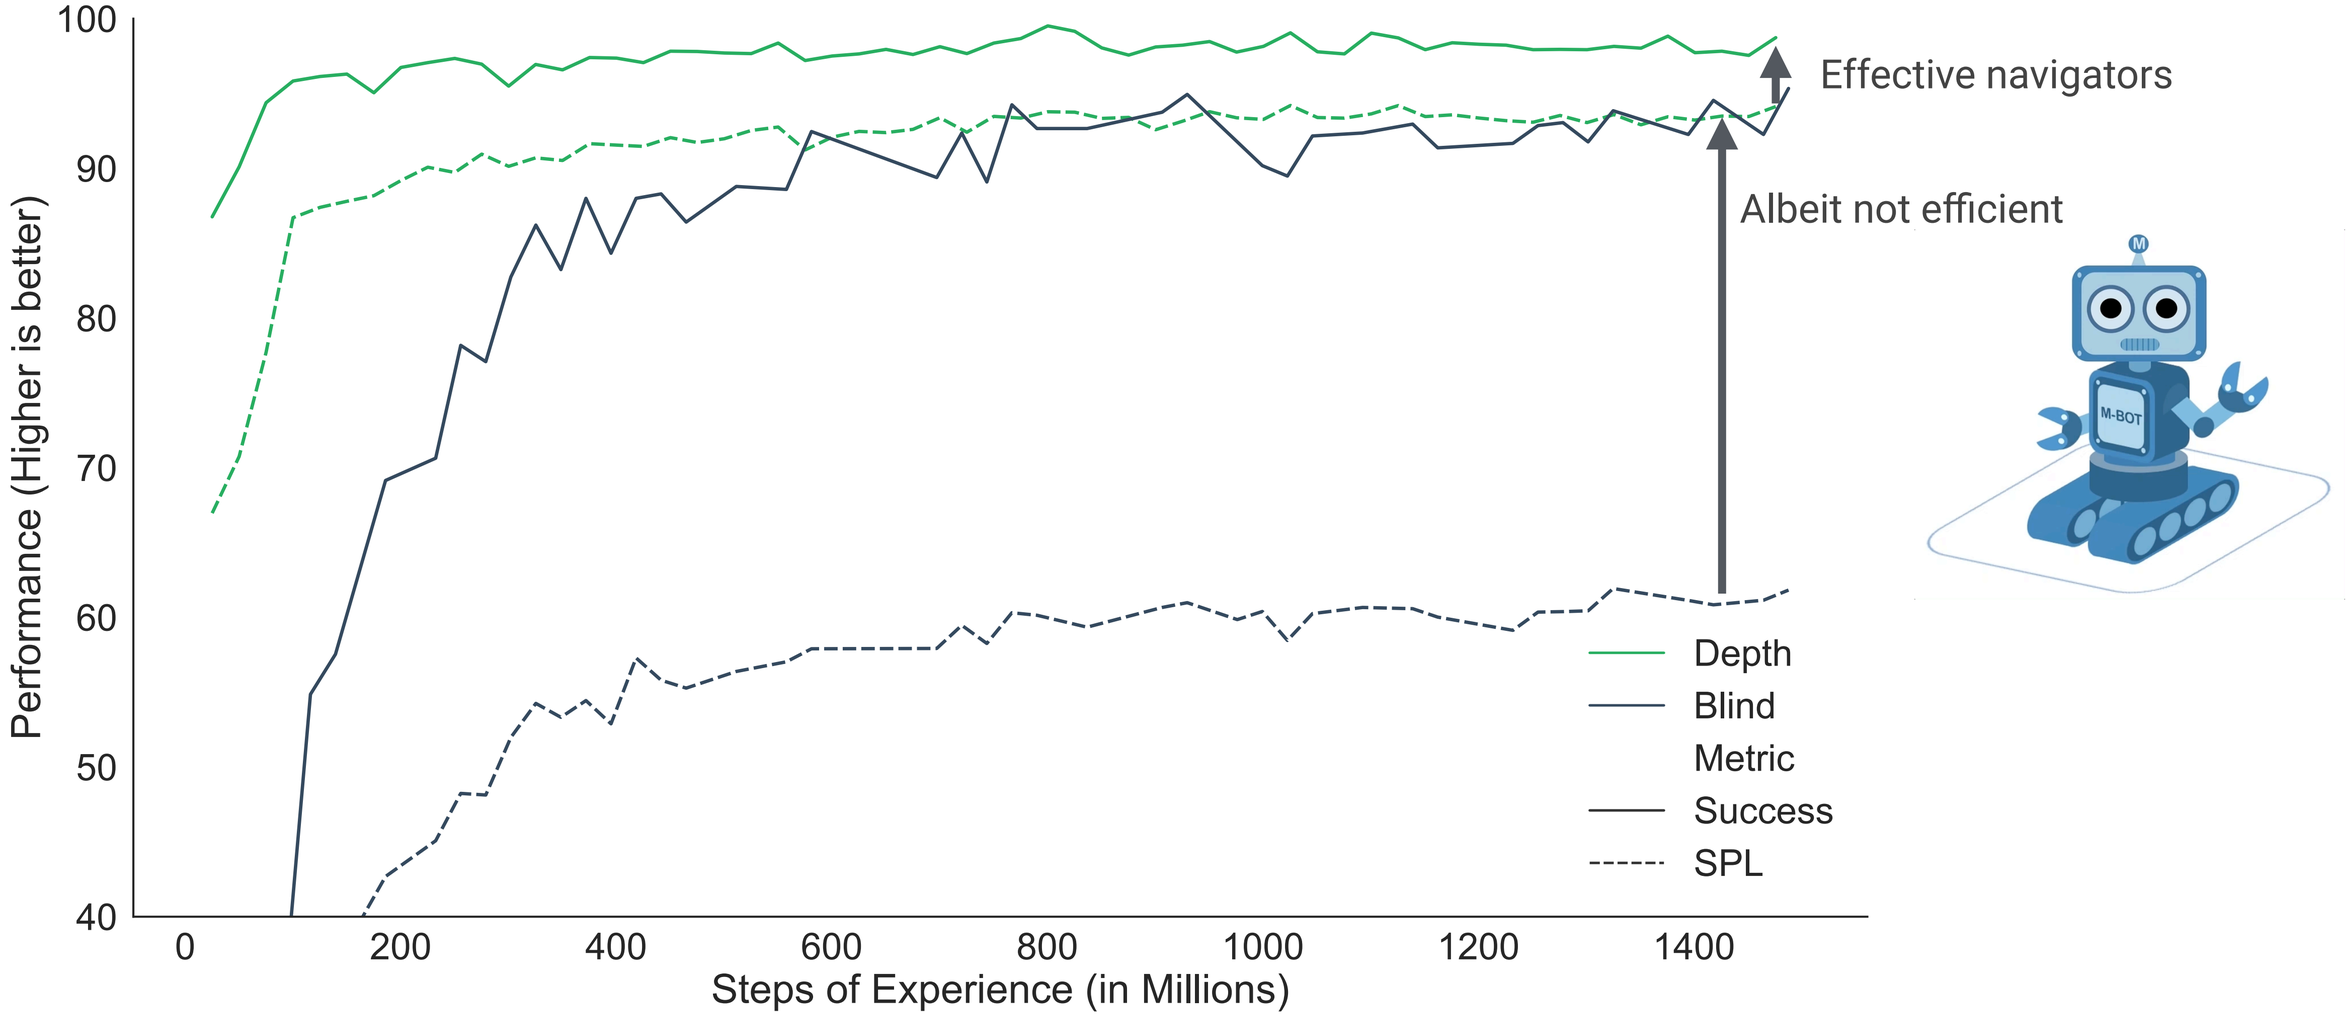

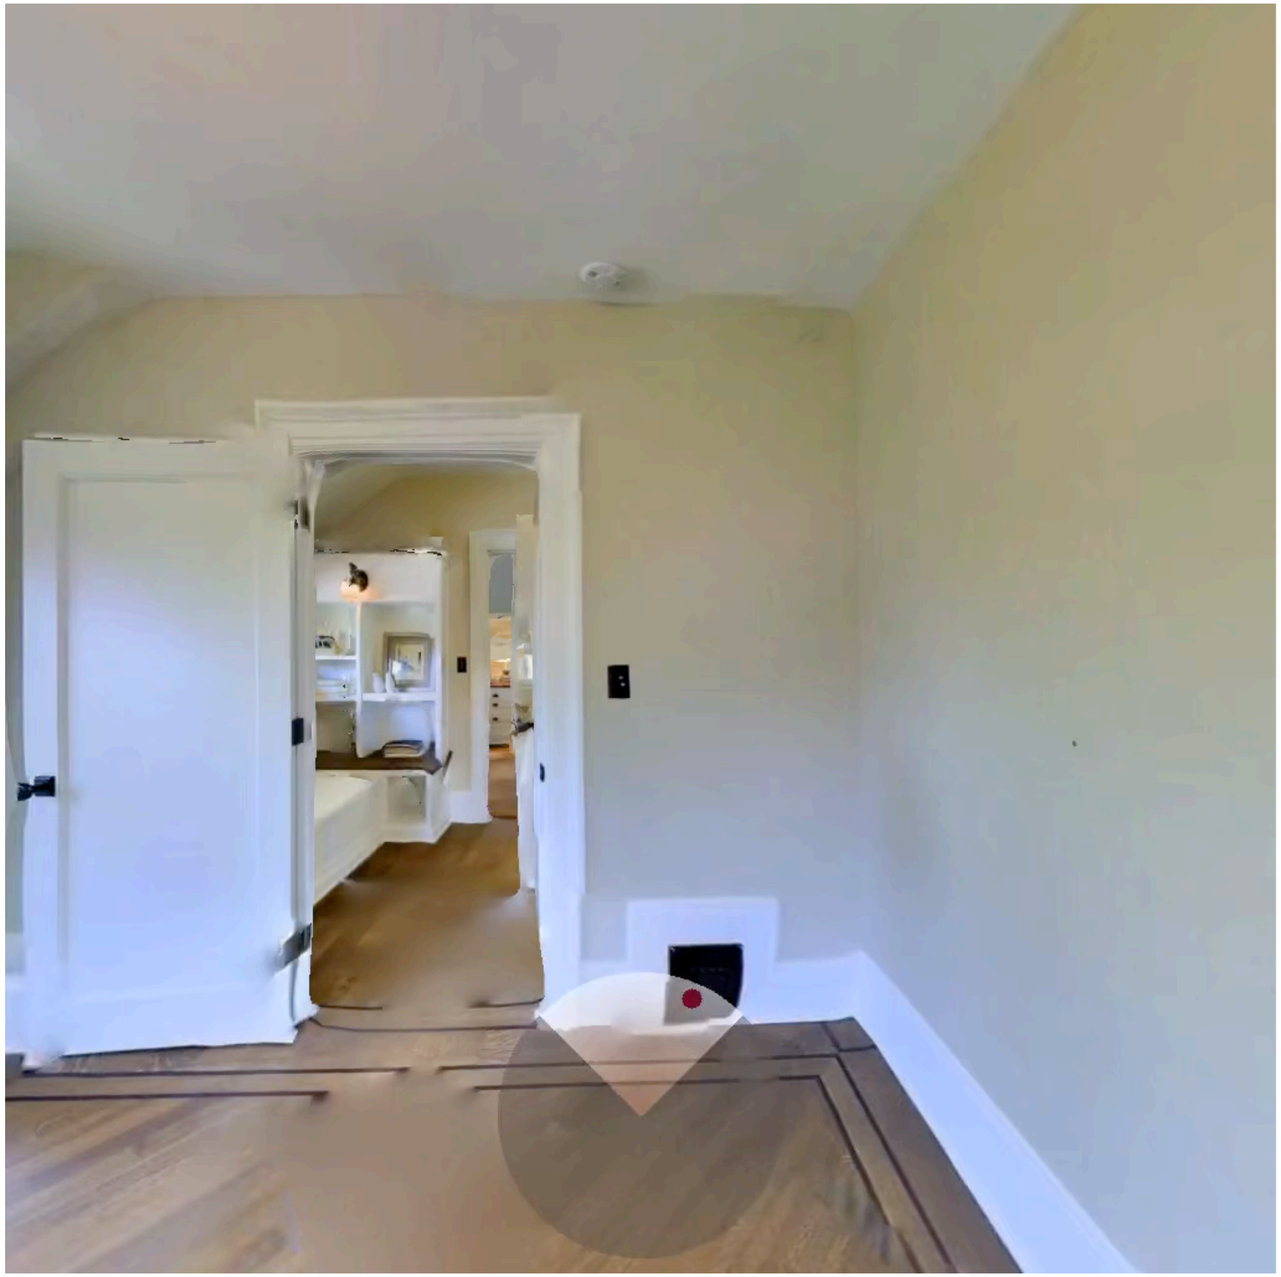

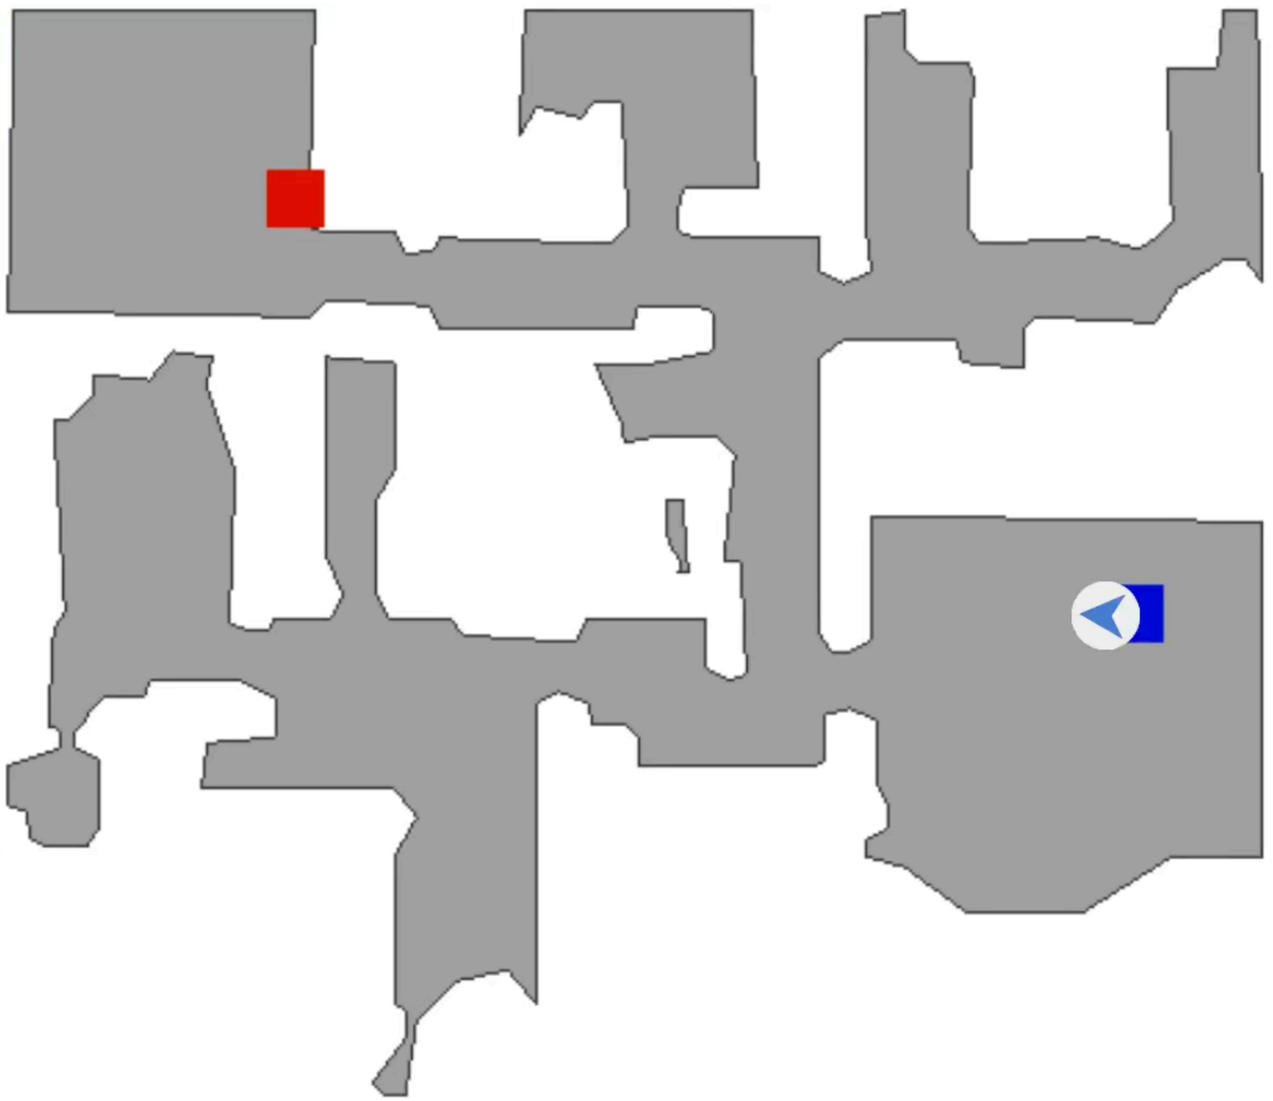

Wall following behavior

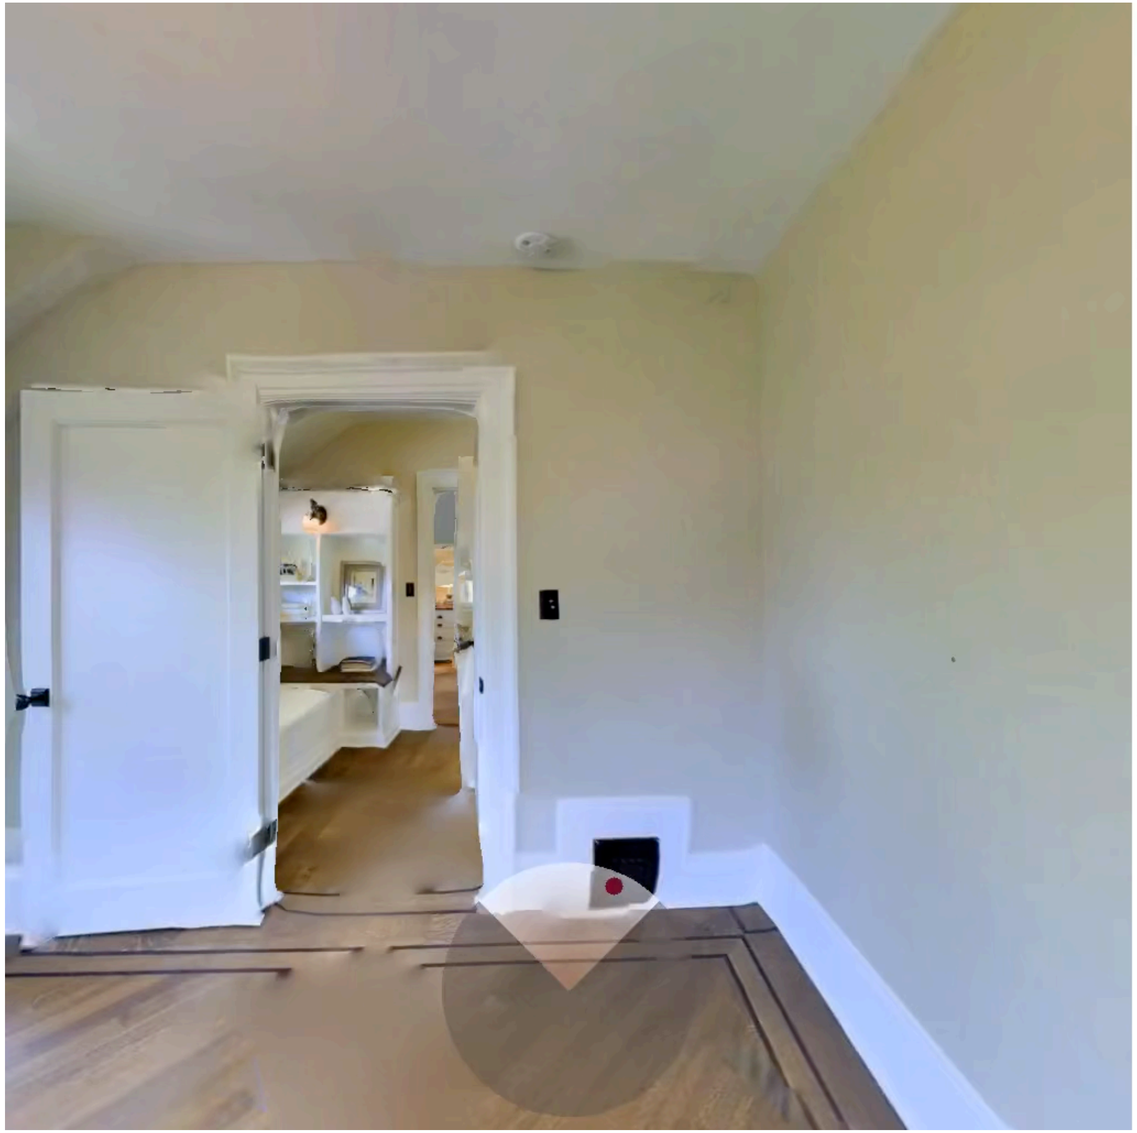

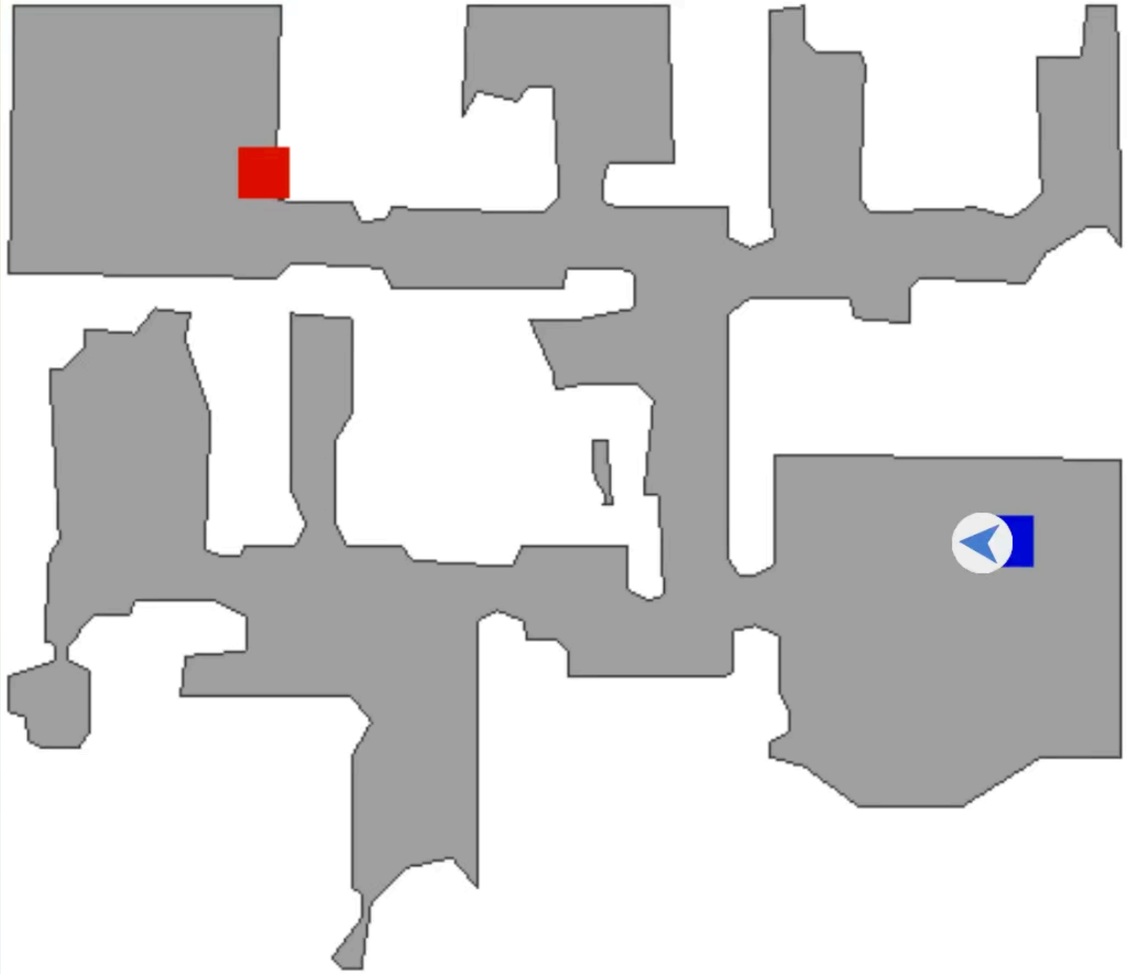

Wall following behavior

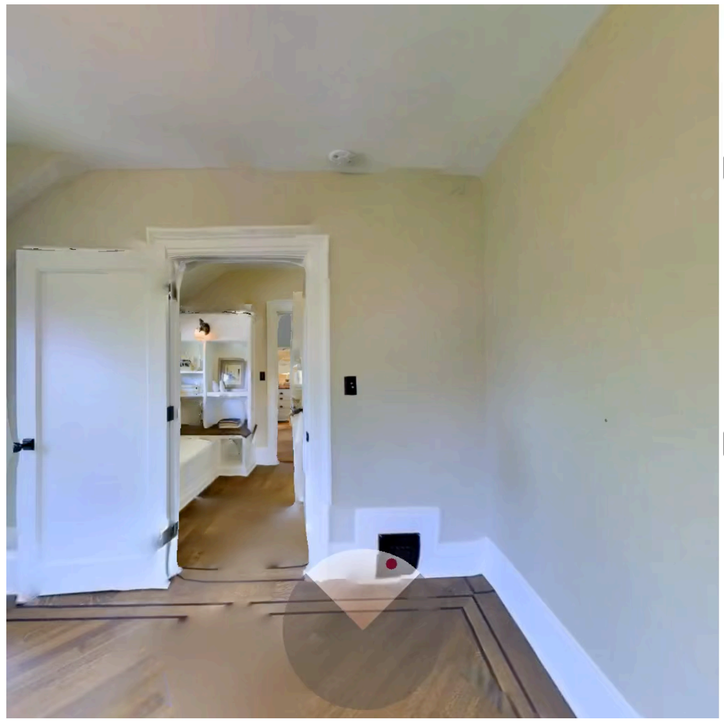

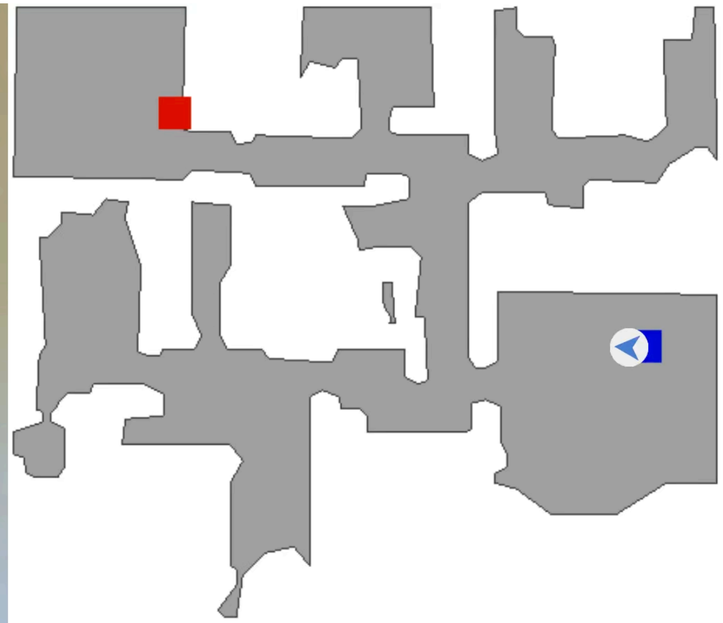

- Environments are unknown

Wall following behavior

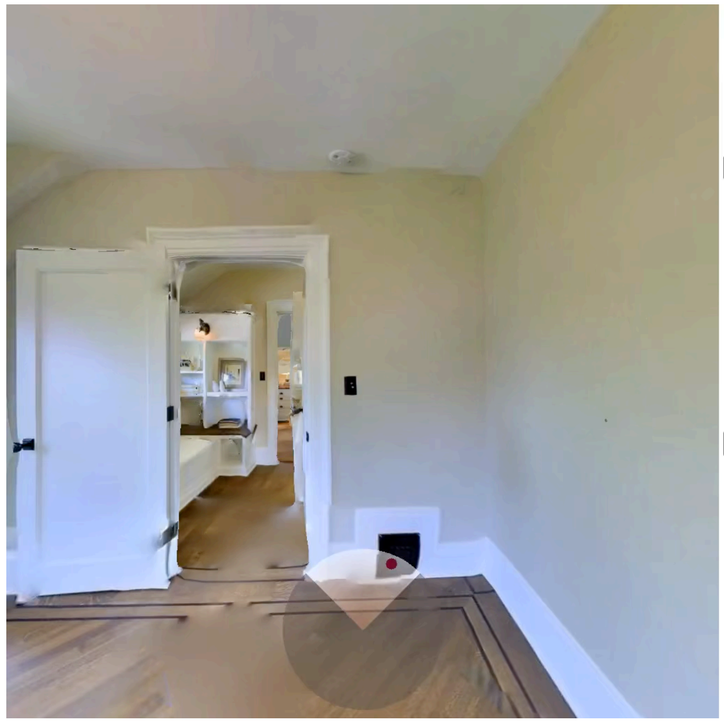

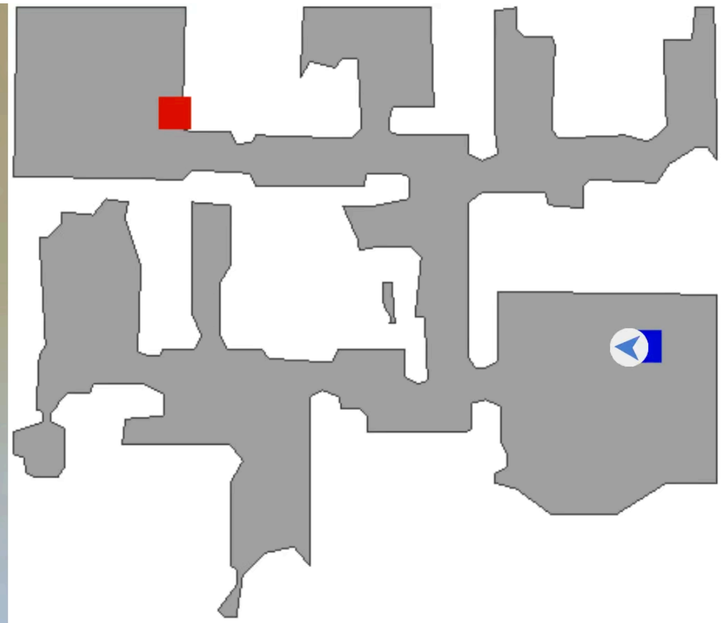

- Environments are unknown
- Coordinate system is episodic

Wall following behavior

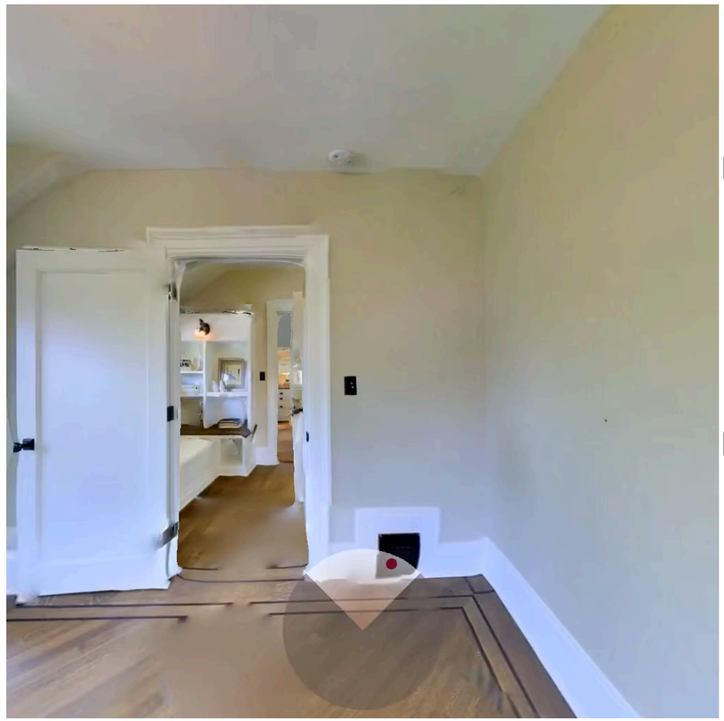

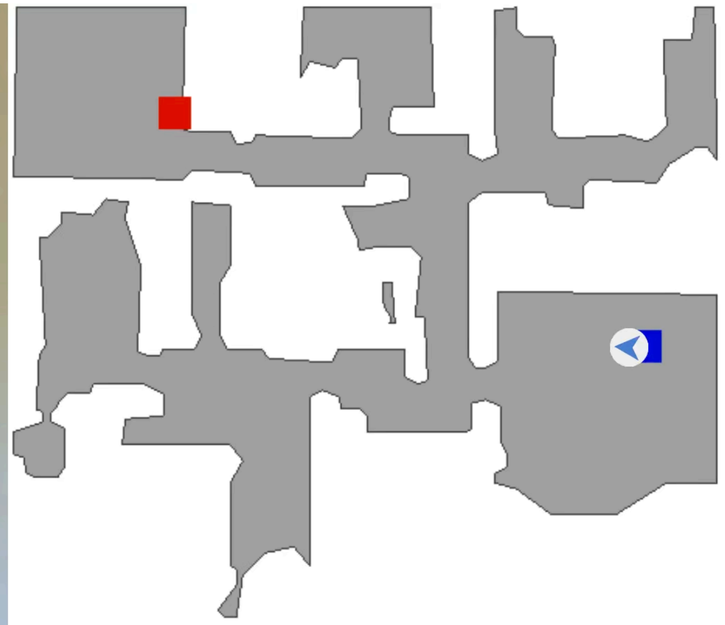

- Environments are unknown
- Coordinate system is episodic
- Learns a policy for navigation in *unknown* environments

Known environment, global coordinates

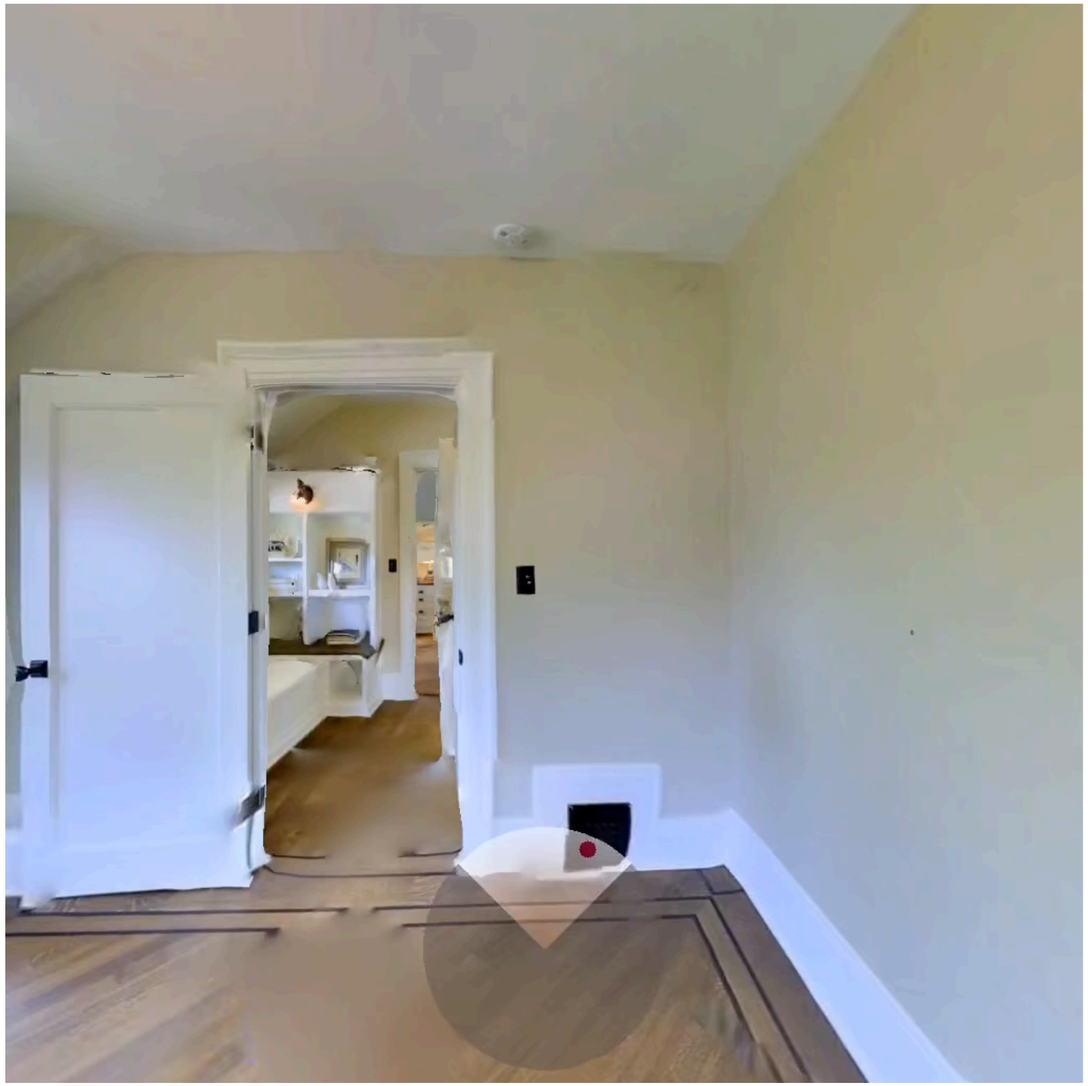

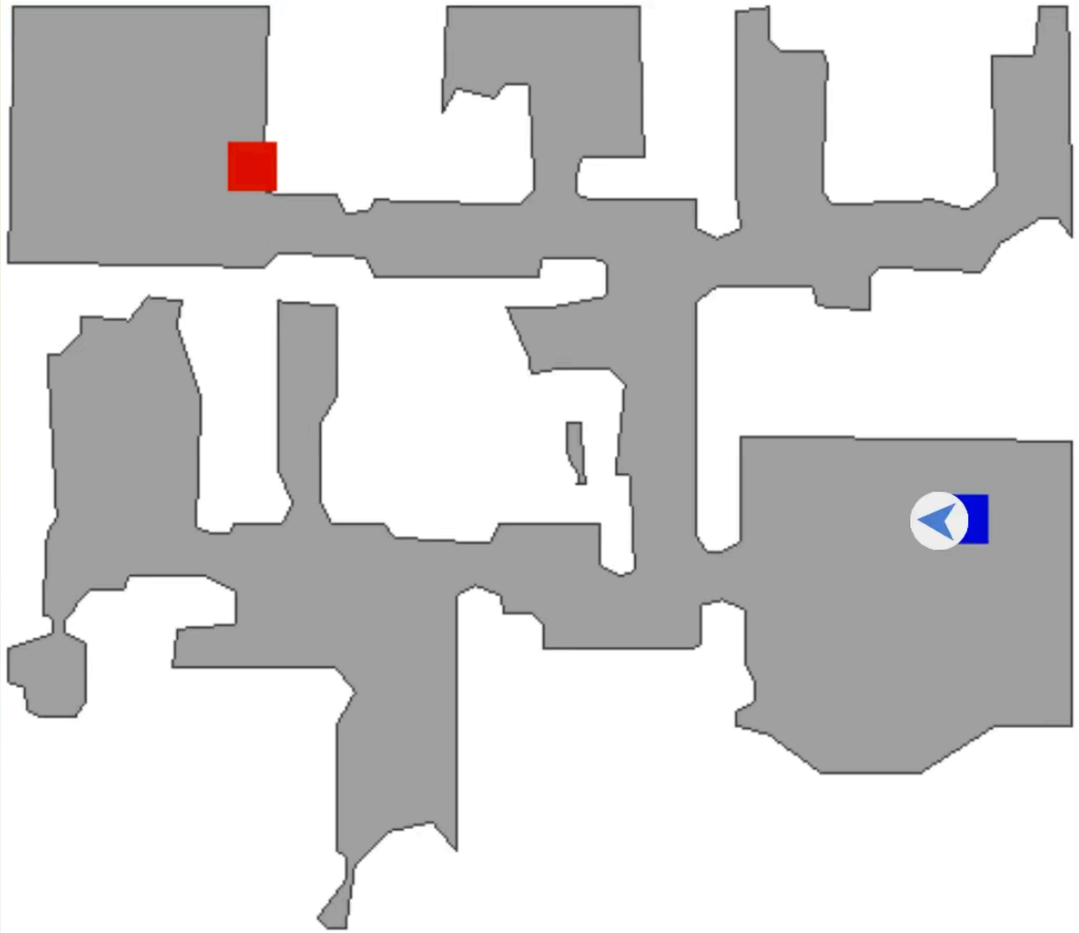

“Bug” algorithms

Algorithmica (1987) 2: 403-430

Algorithmica

Path-Planning Strategies for a Point Mobile Automaton Moving Amidst Unknown Obstacles of Arbitrary Shape[1]

Vladimir J. Lumelsky[2] and Alexander A. Stepanov[3]

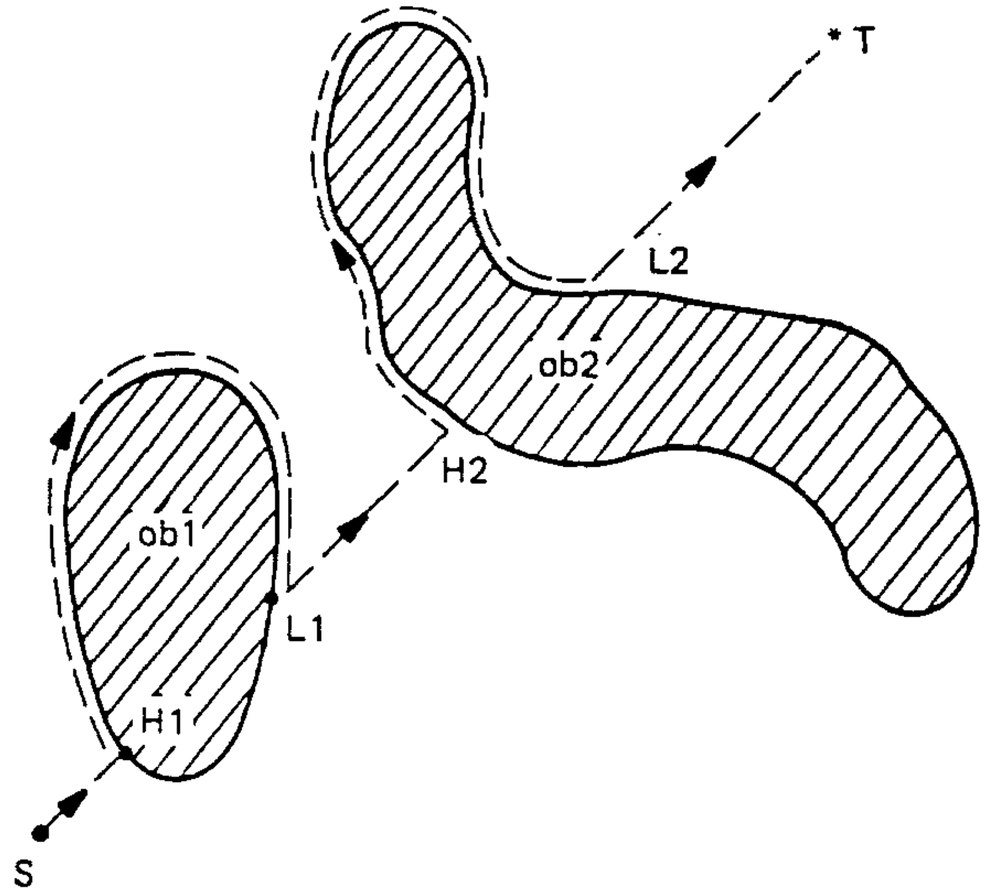

Fig. 4. Automaton’s path (dotted line) under Algorithm Bug2.

Memory is key to performance

Memory is key to performance

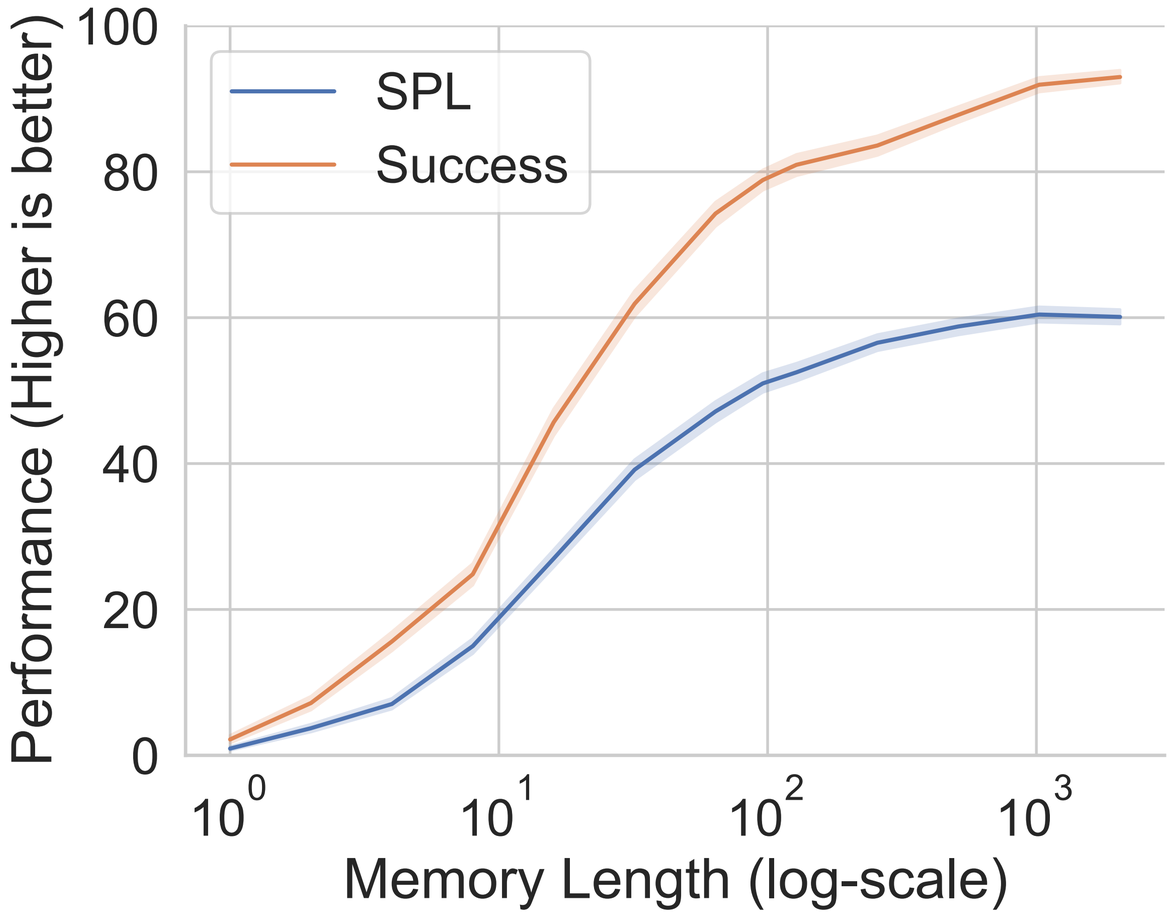

Memory is used to detect collisions

Memory is used to detect collisions

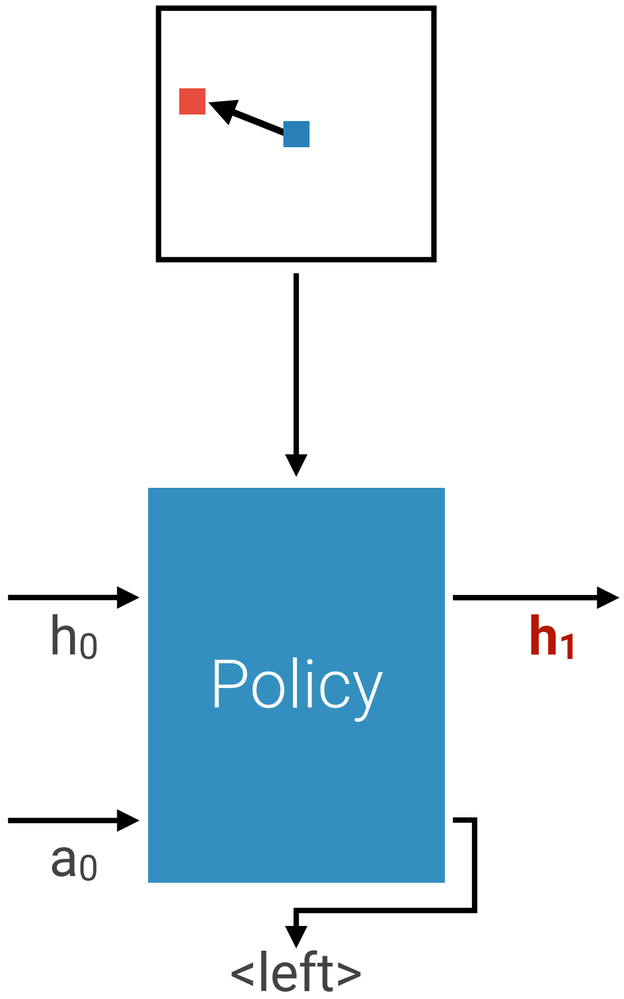

Memory is used to detect collisions

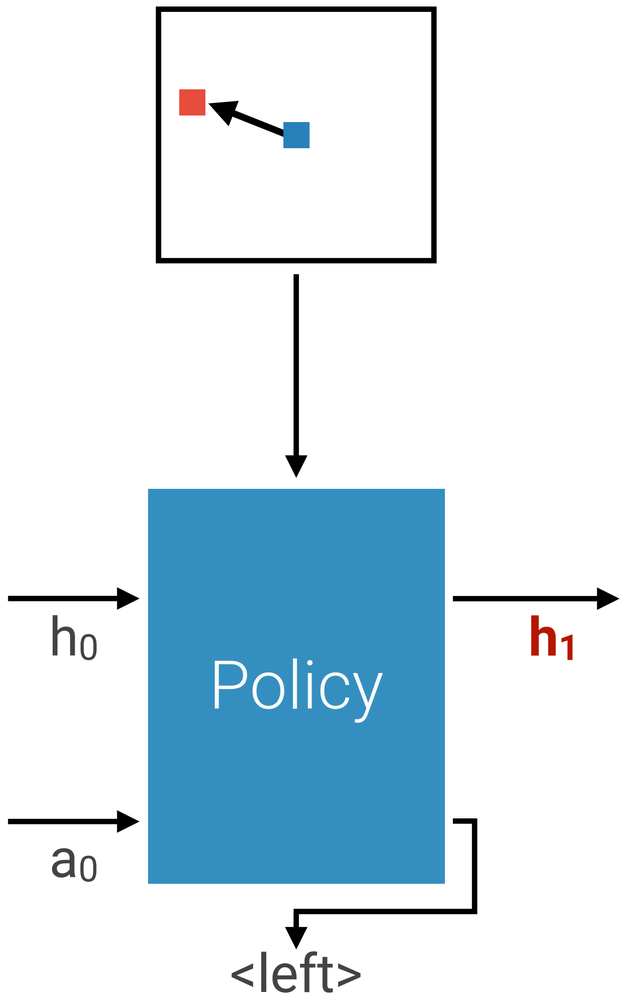

Collision or not

Memory is used to detect collisions

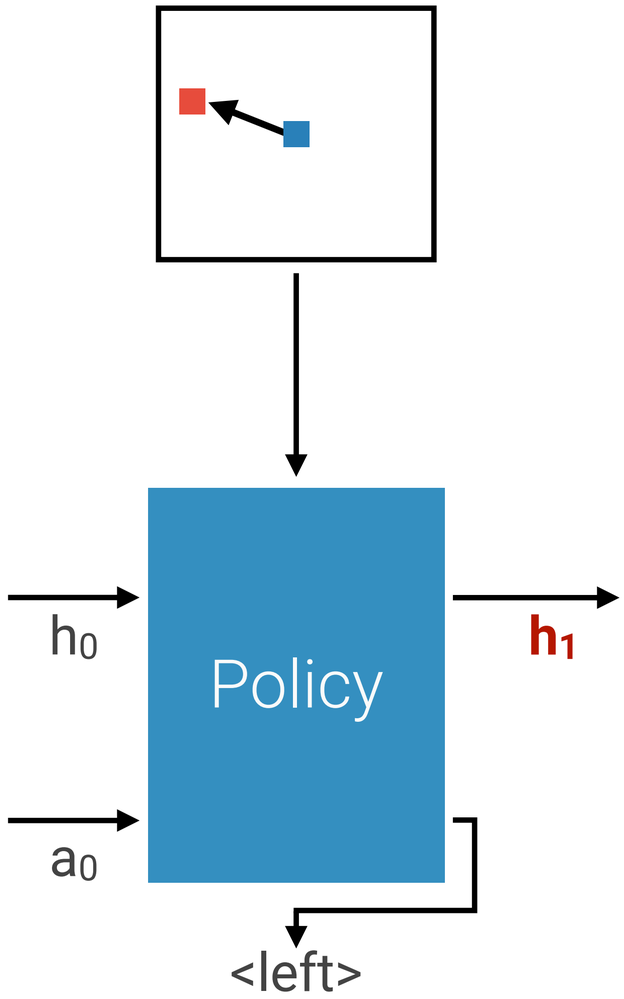

Only for explanation

Collision or not

Memory is used to detect collisions

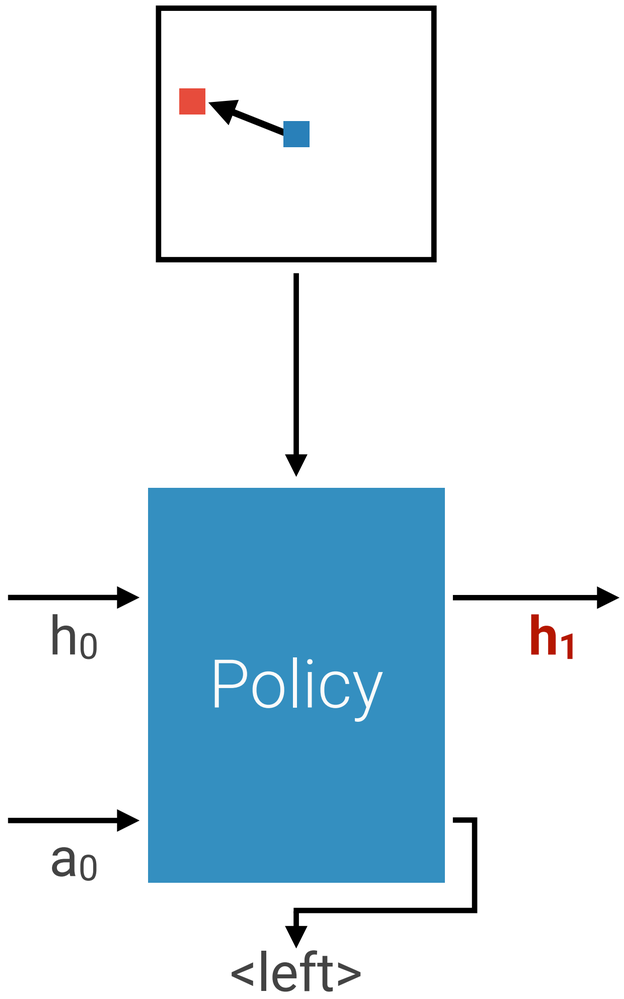

Only for explanation

No gradients into agent

Collision or not

Memory is used to detect collisions

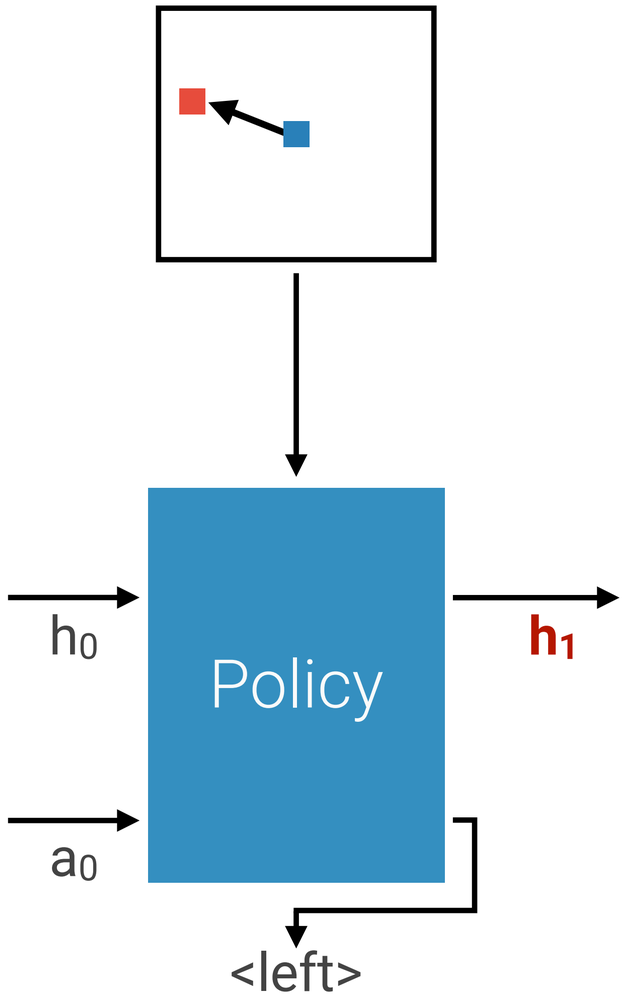

Only for explanation

No gradients into agent

Collision or not

98% accuracy on held-out data

Memory is used to detect collisions

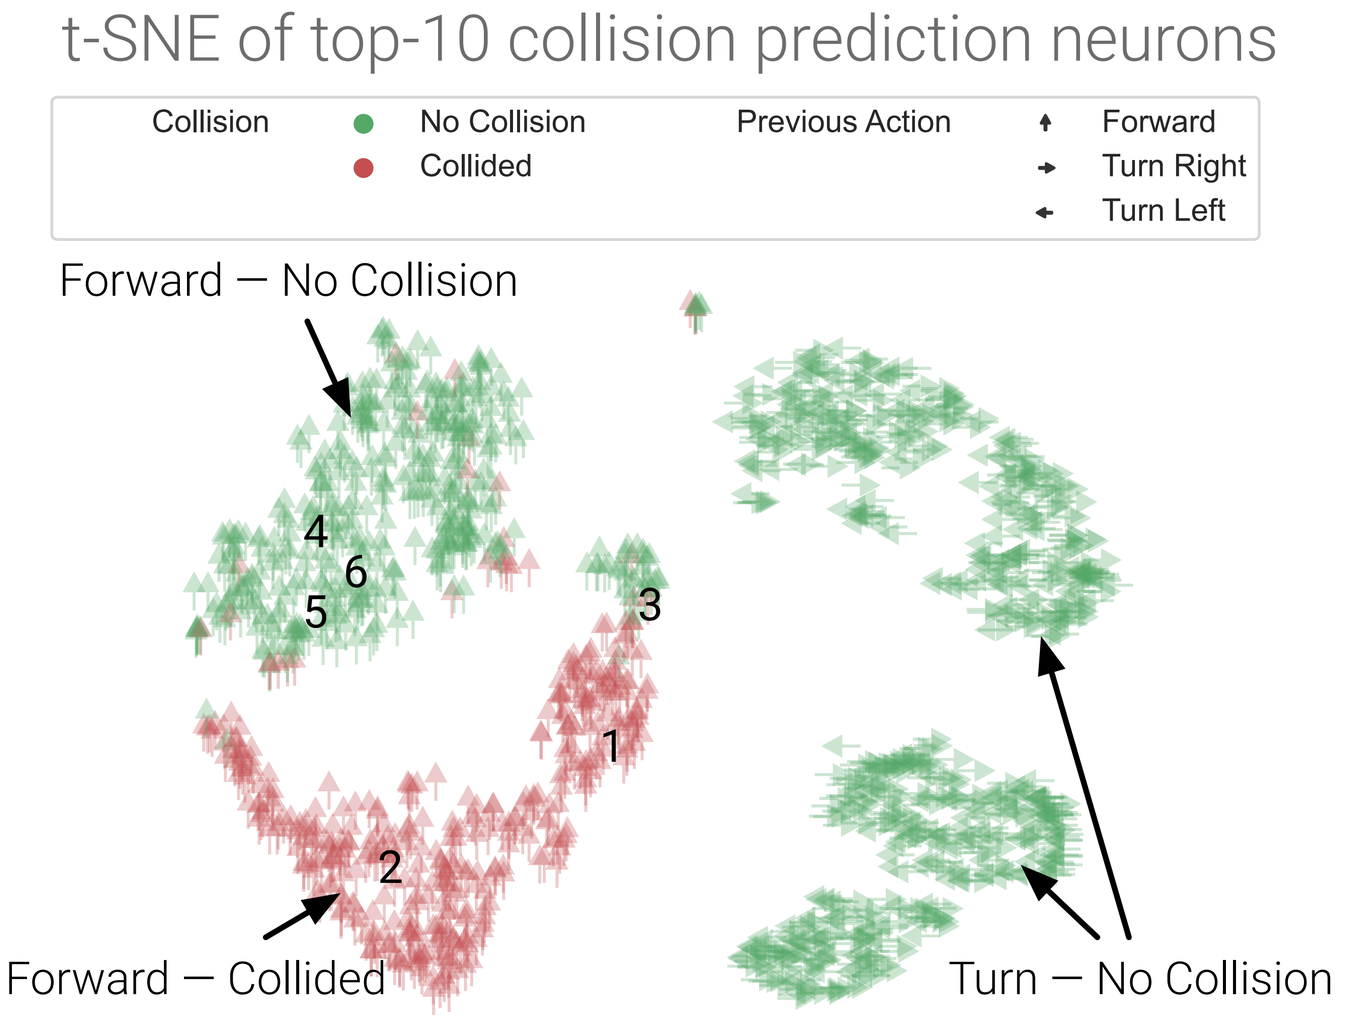

Memory is used for mapping

Memory is used for mapping

- AI rendition of Menzel (1973)’s chimpanzee experiment


T
S


T
S


T
S

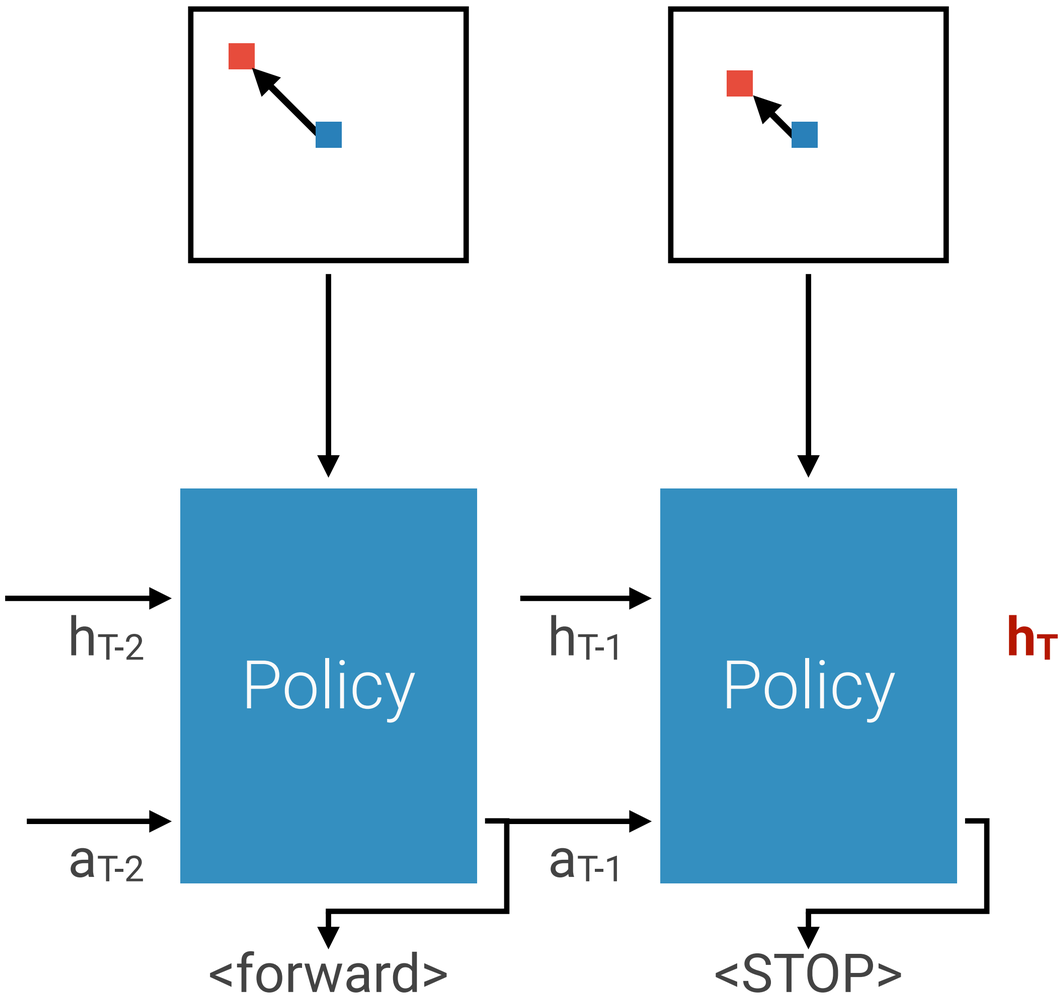
Policy
Policy
h_{T-2}
h_{T-1}
h_T
a_{T-2}
a_{T-1}
<forward>
<STOP>

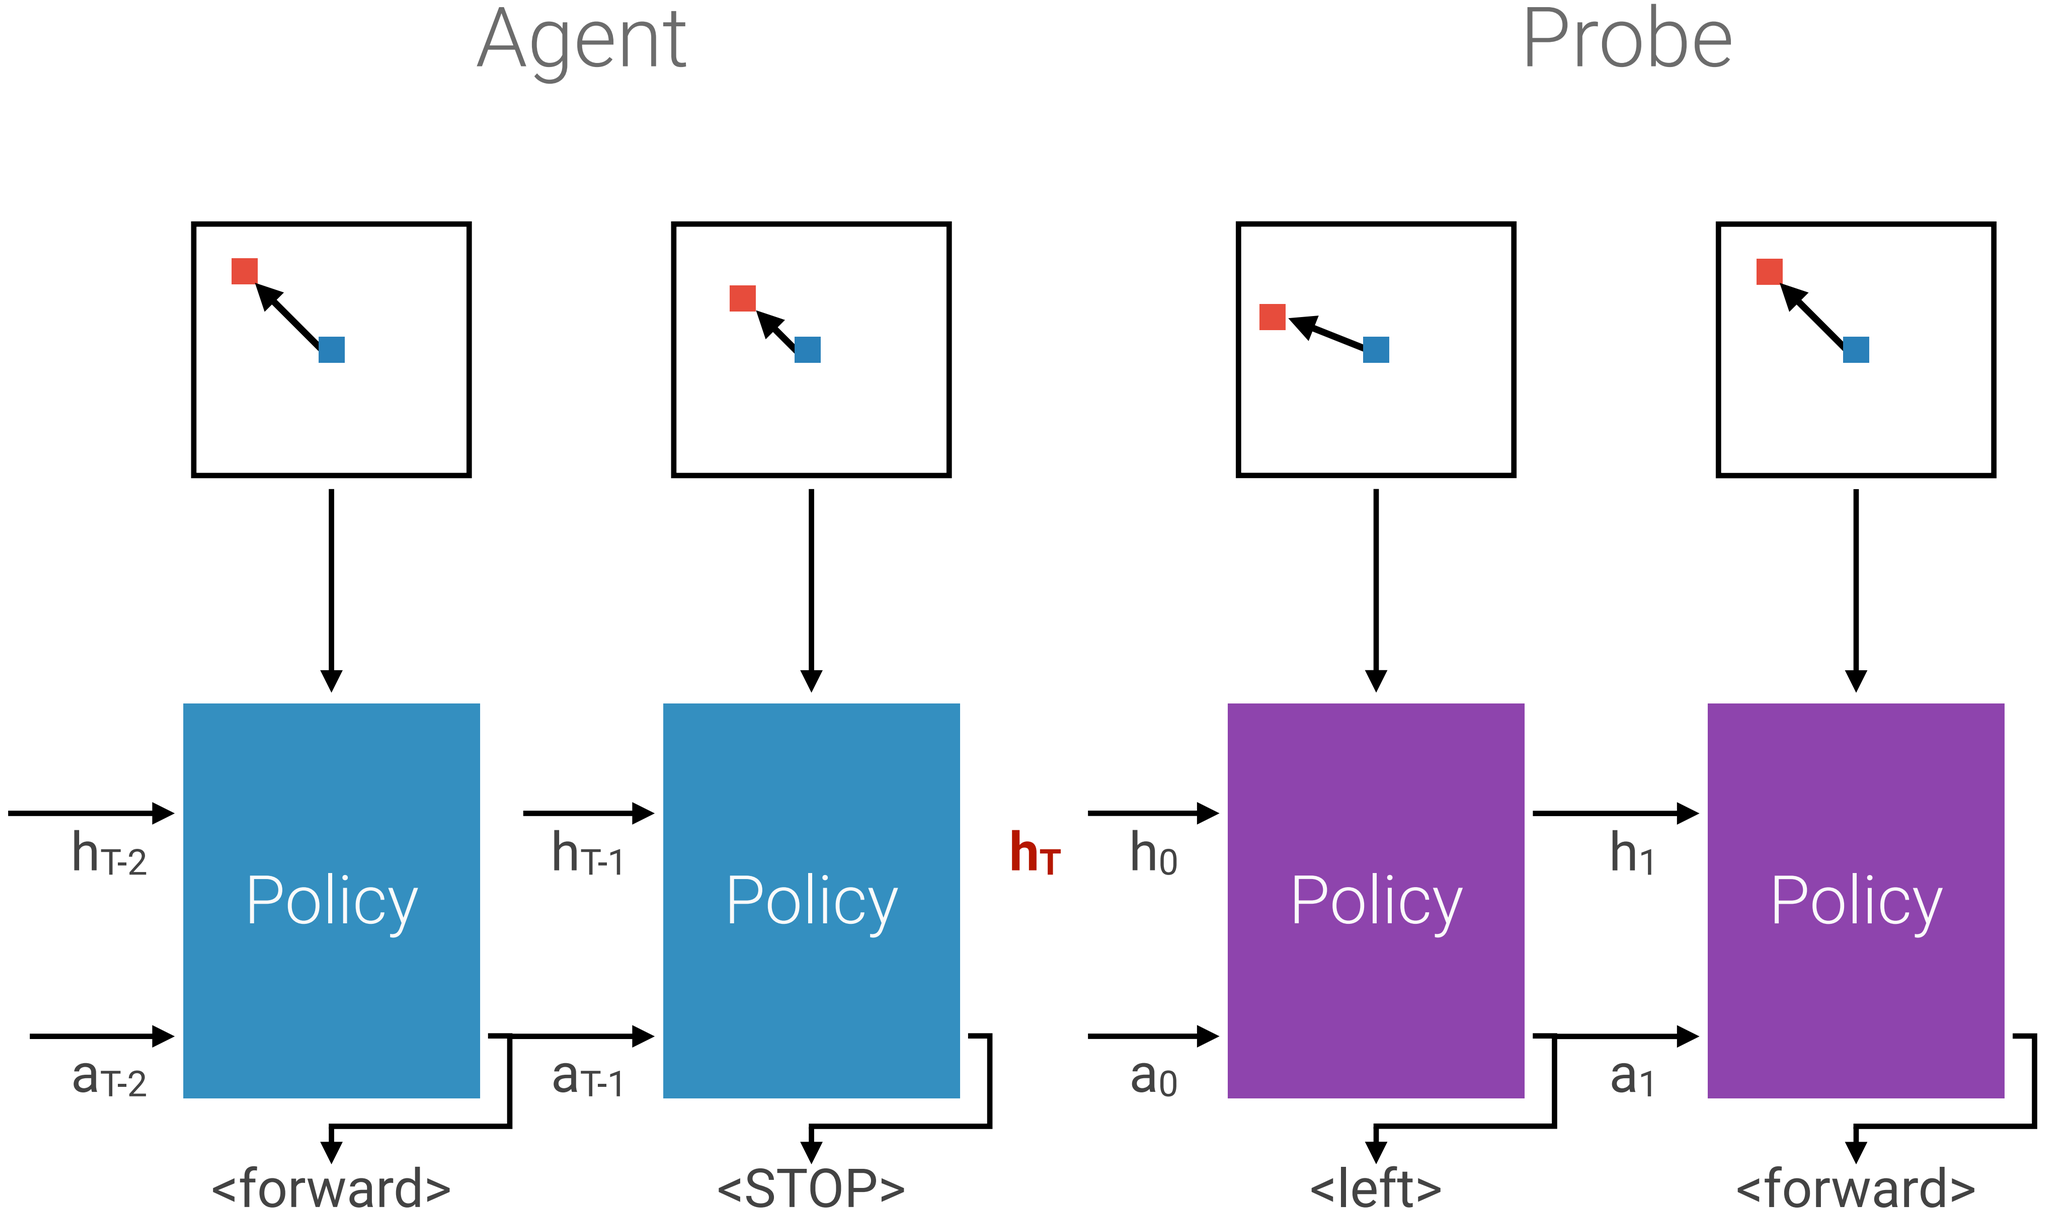
Agent
Probe
Policy
Policy
Policy
Policy
h_{T-2}
h_{T-1}
h_T
h_0
h_1
a_{T-2}
a_{T-1}
a_0
a_1
<forward>
<STOP>
<left>
<forward>


T
S

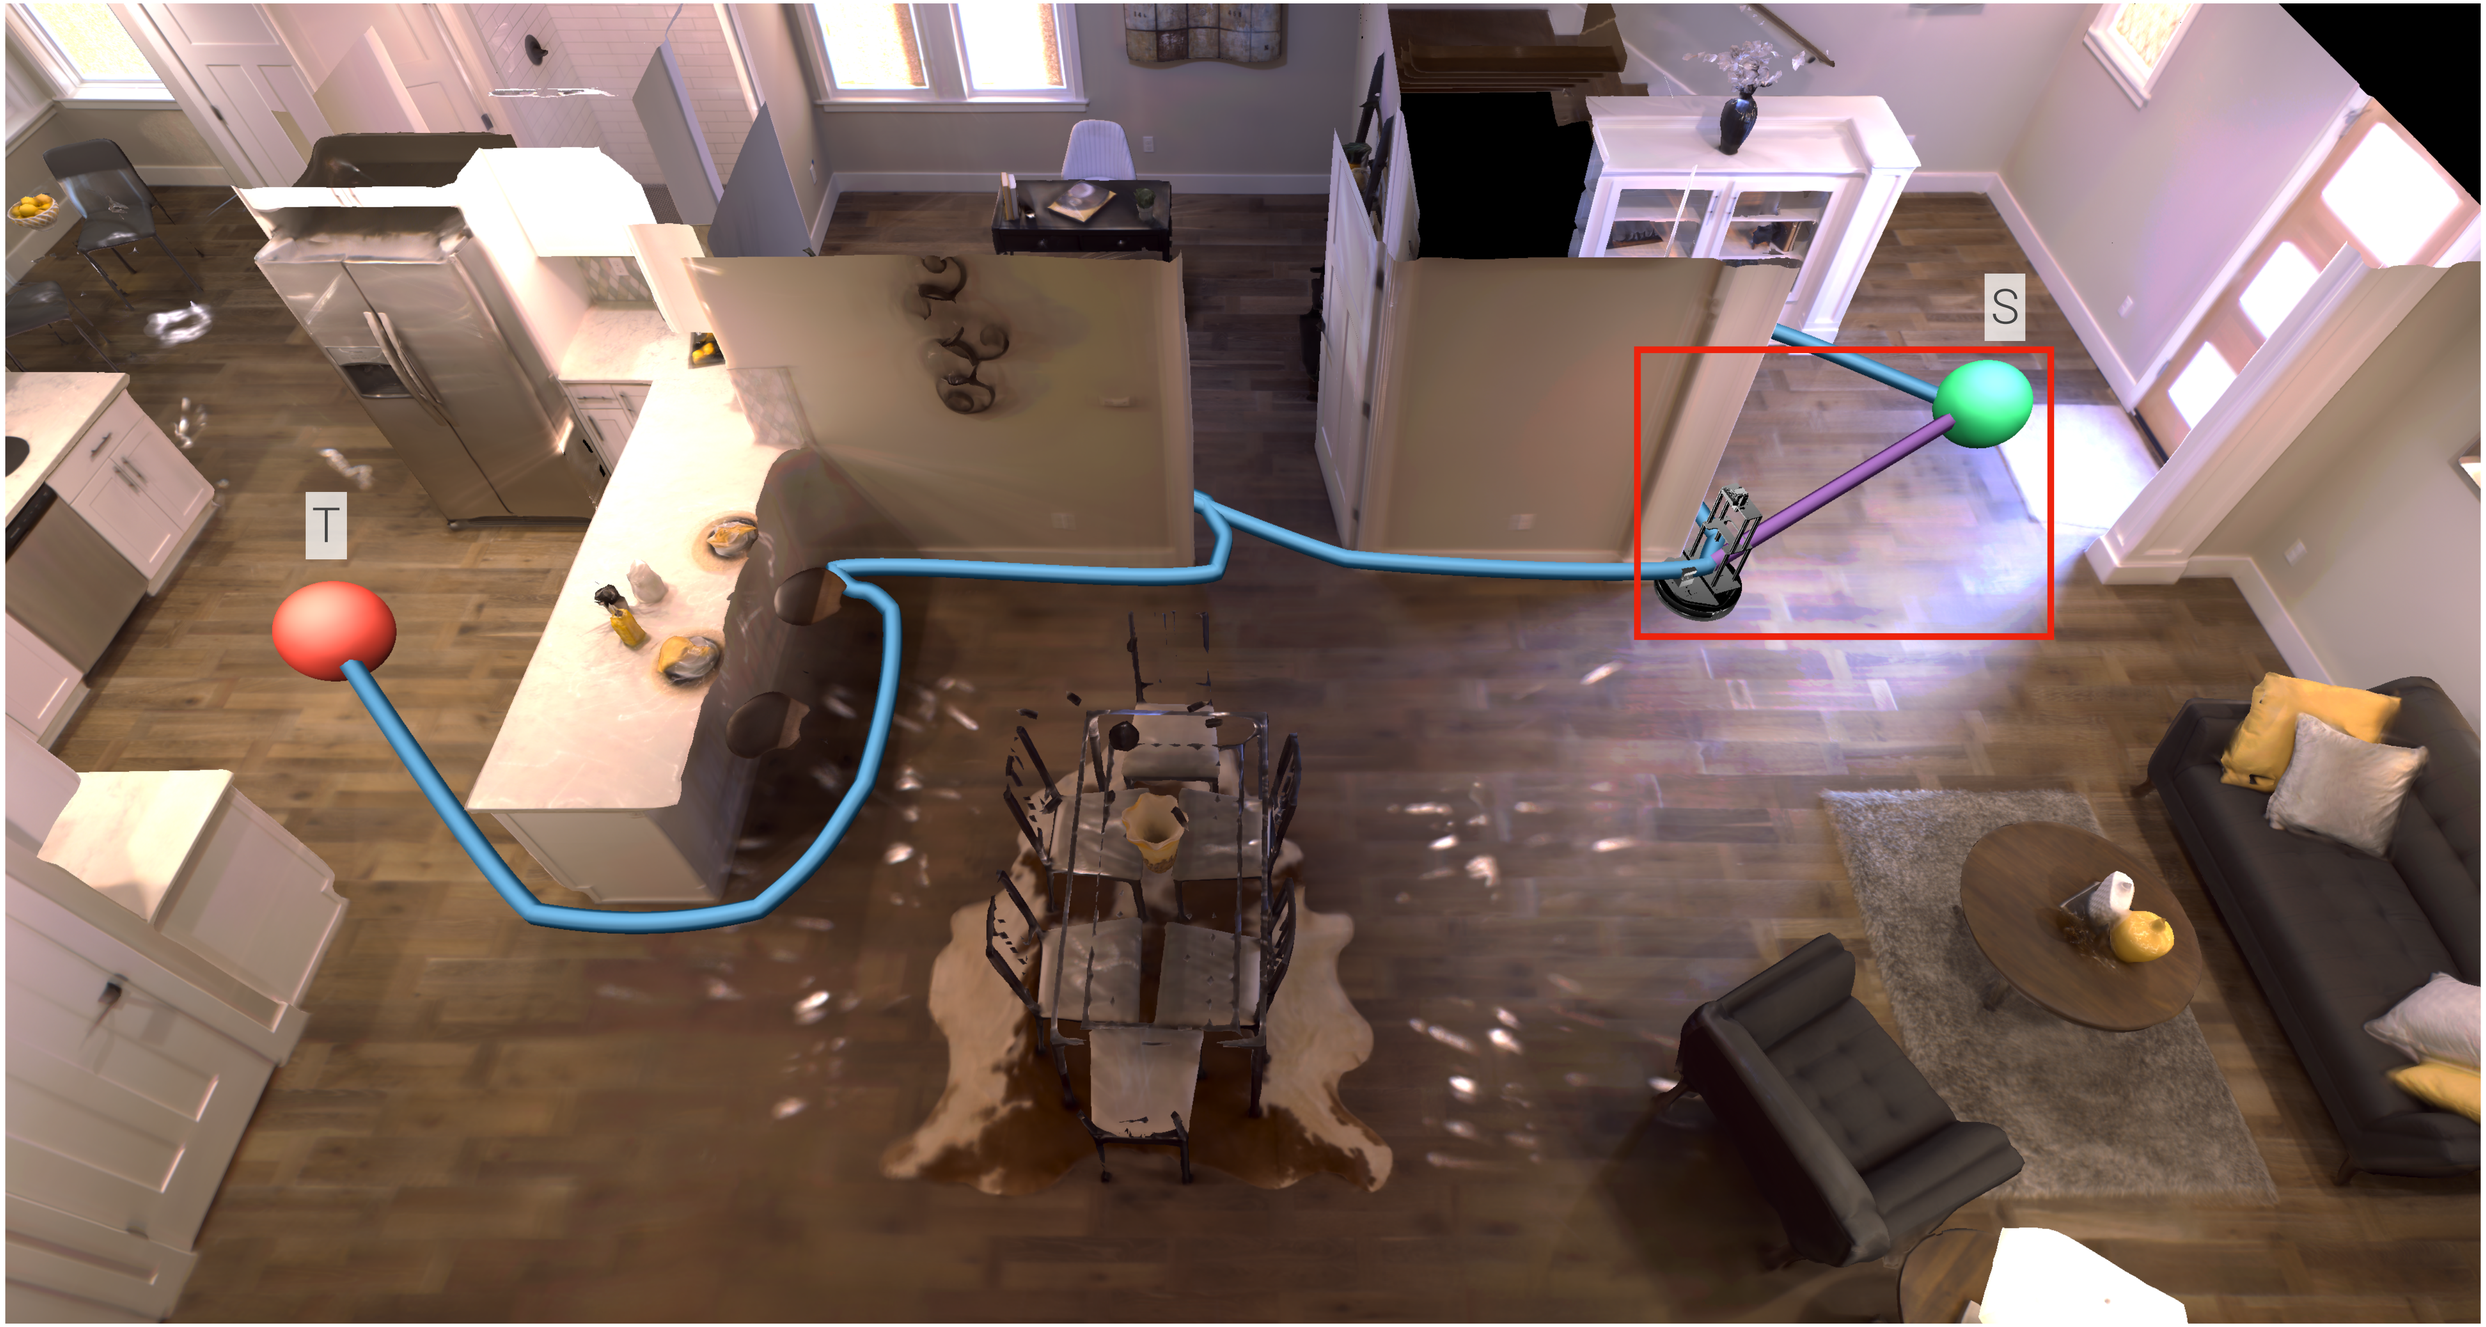
T
S

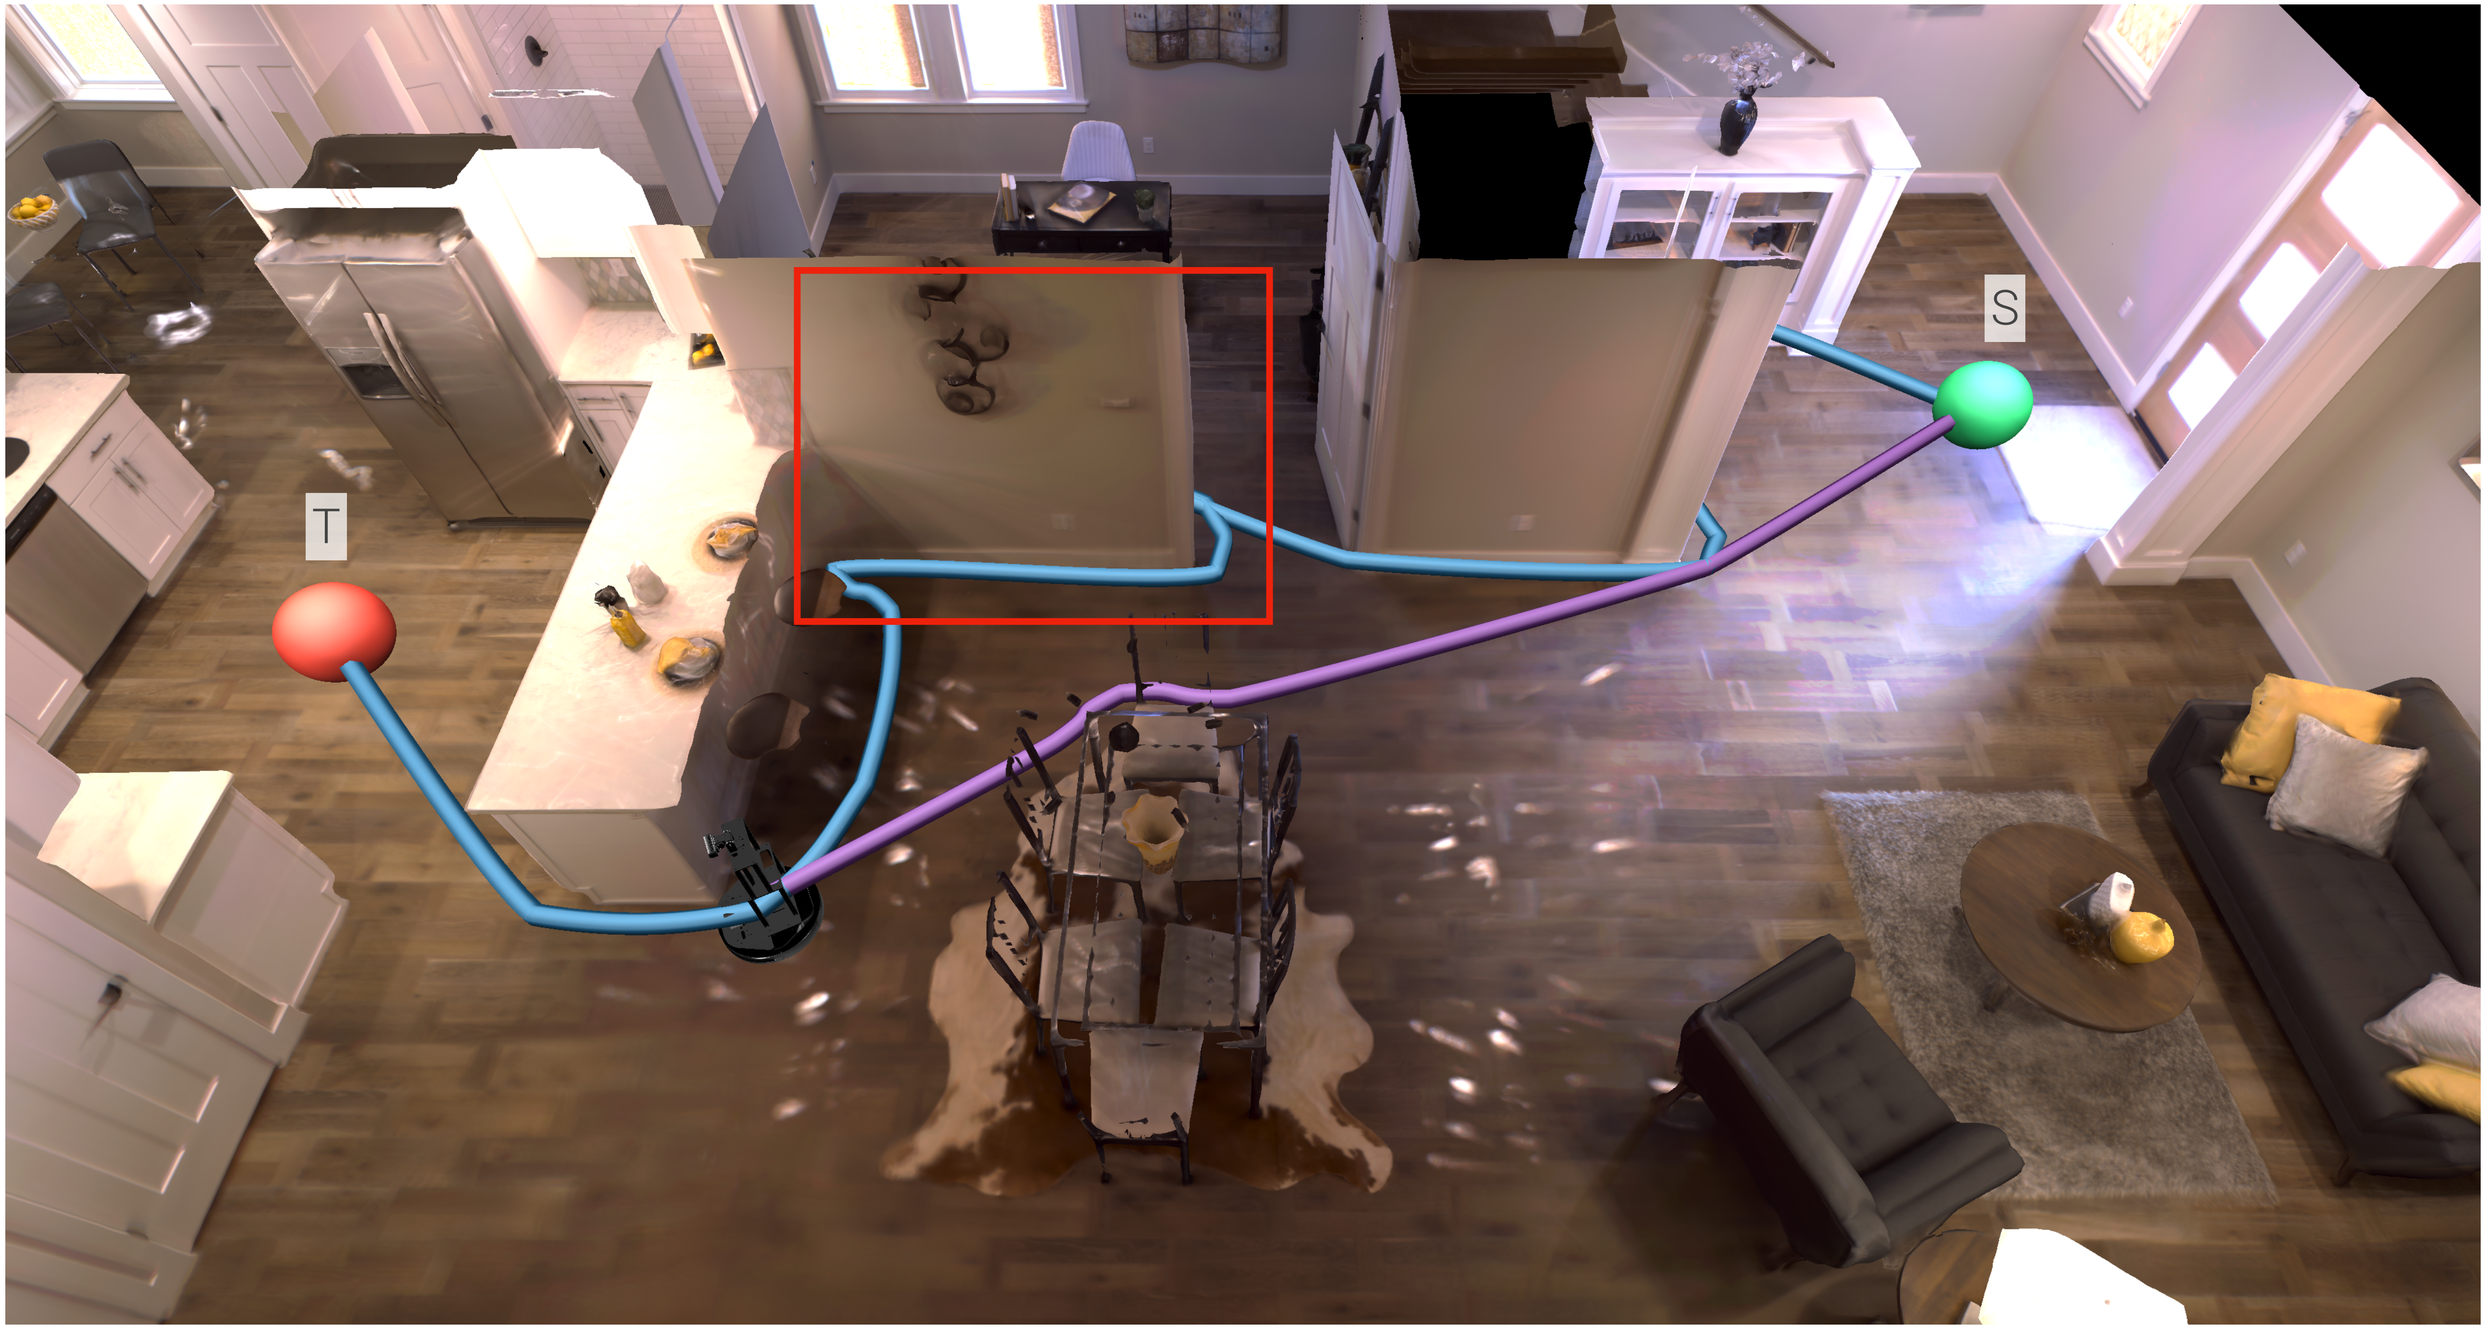
T
S


T
S

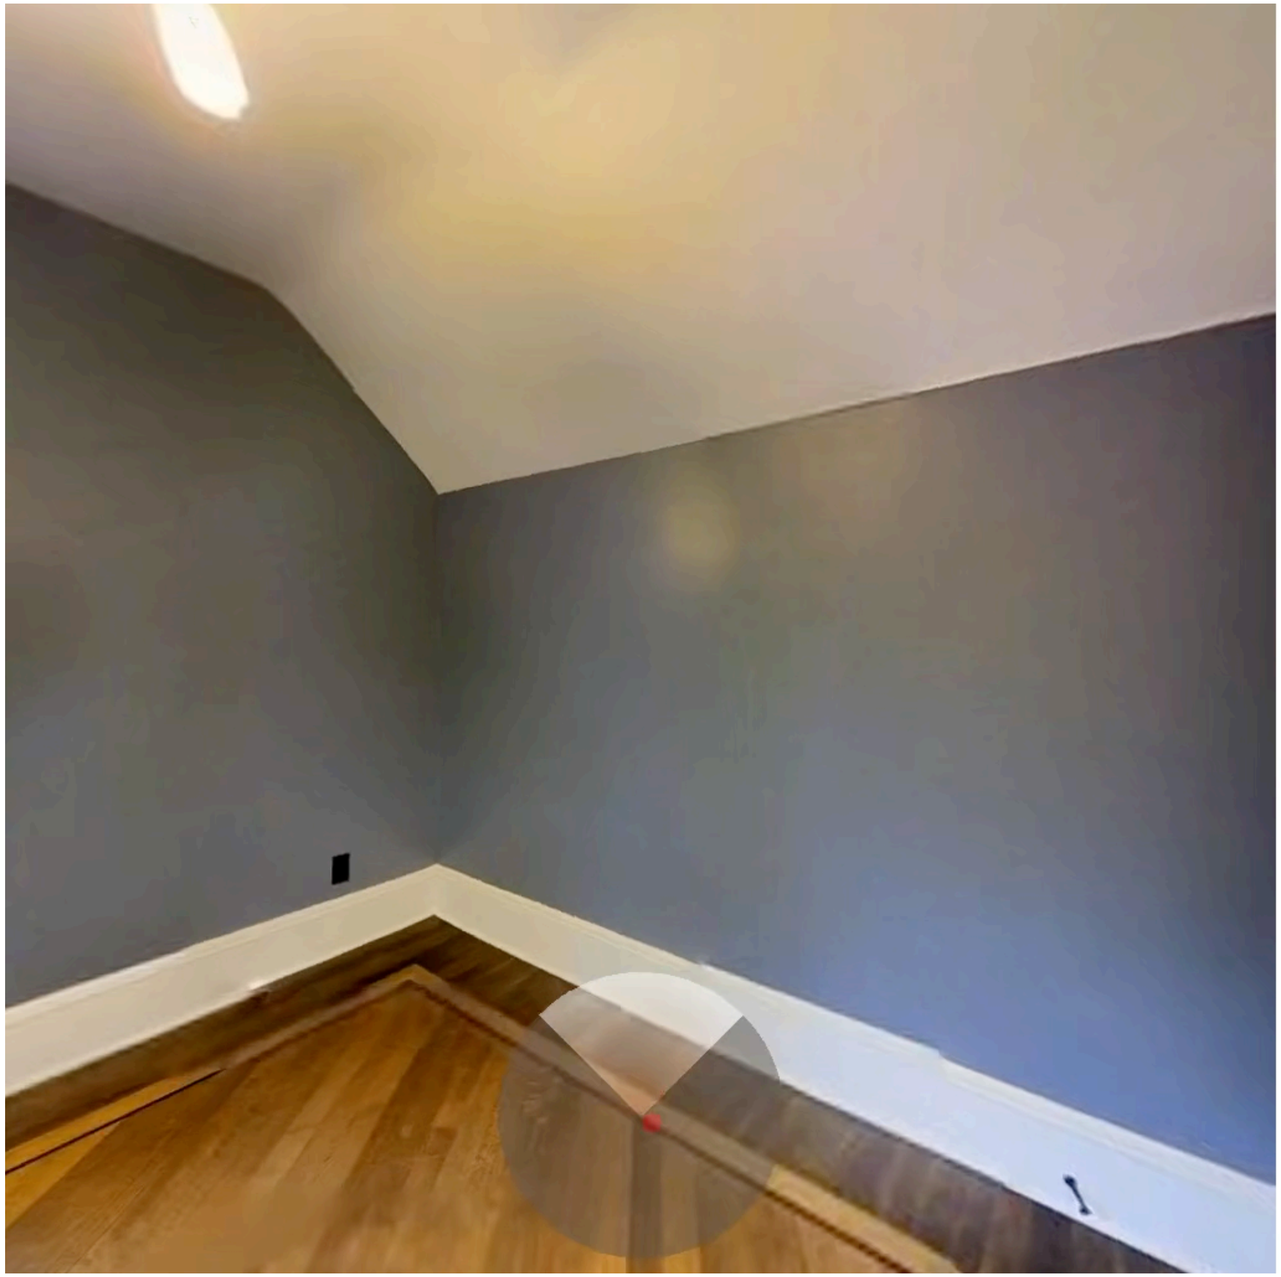

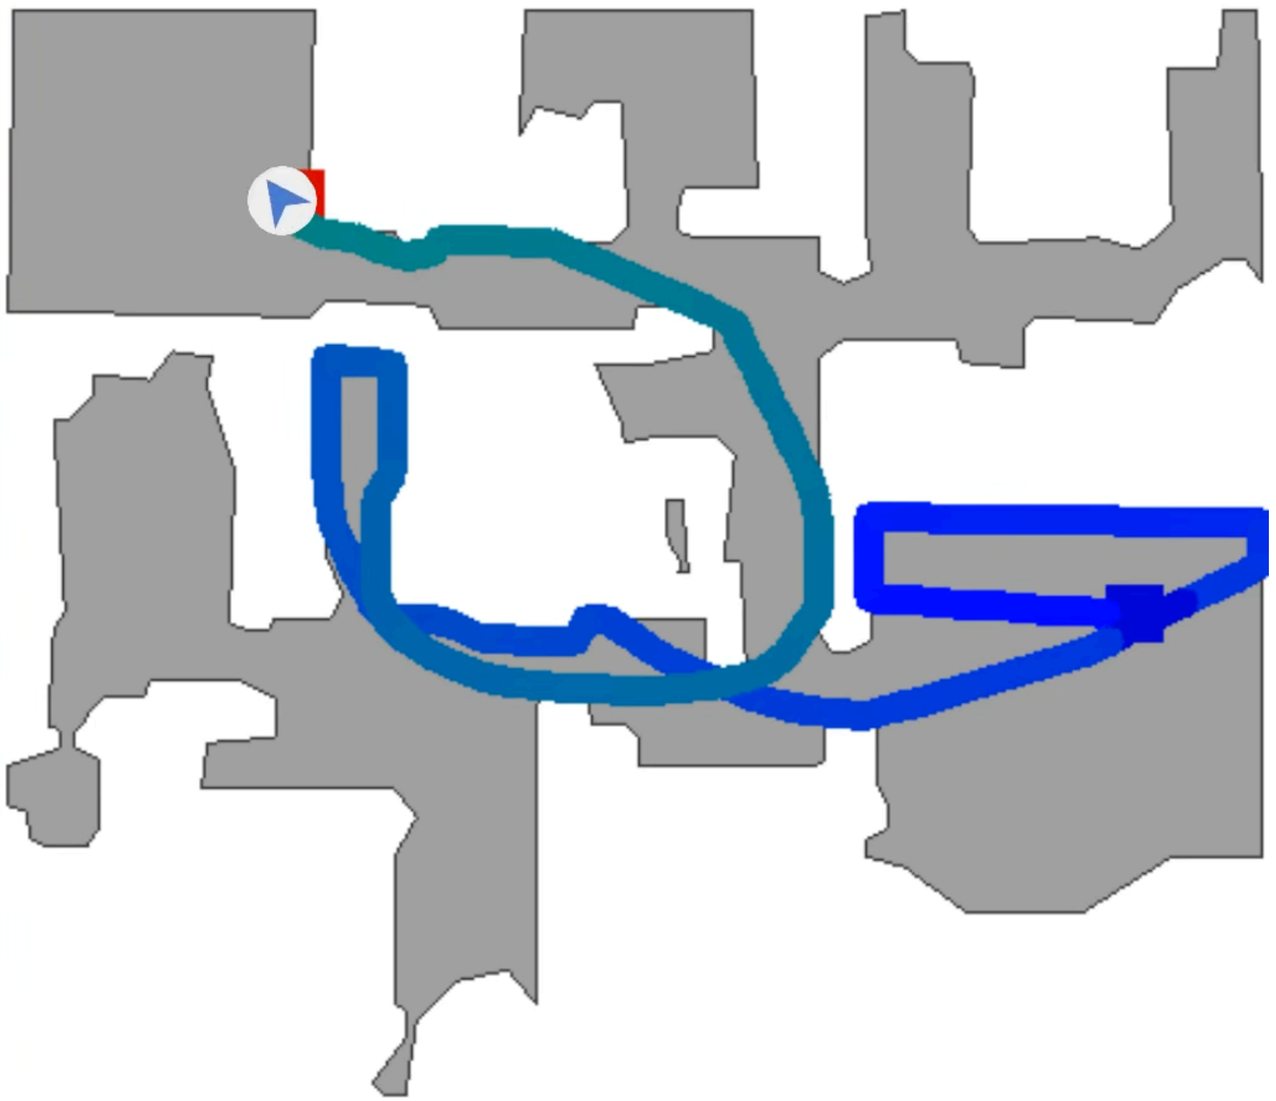

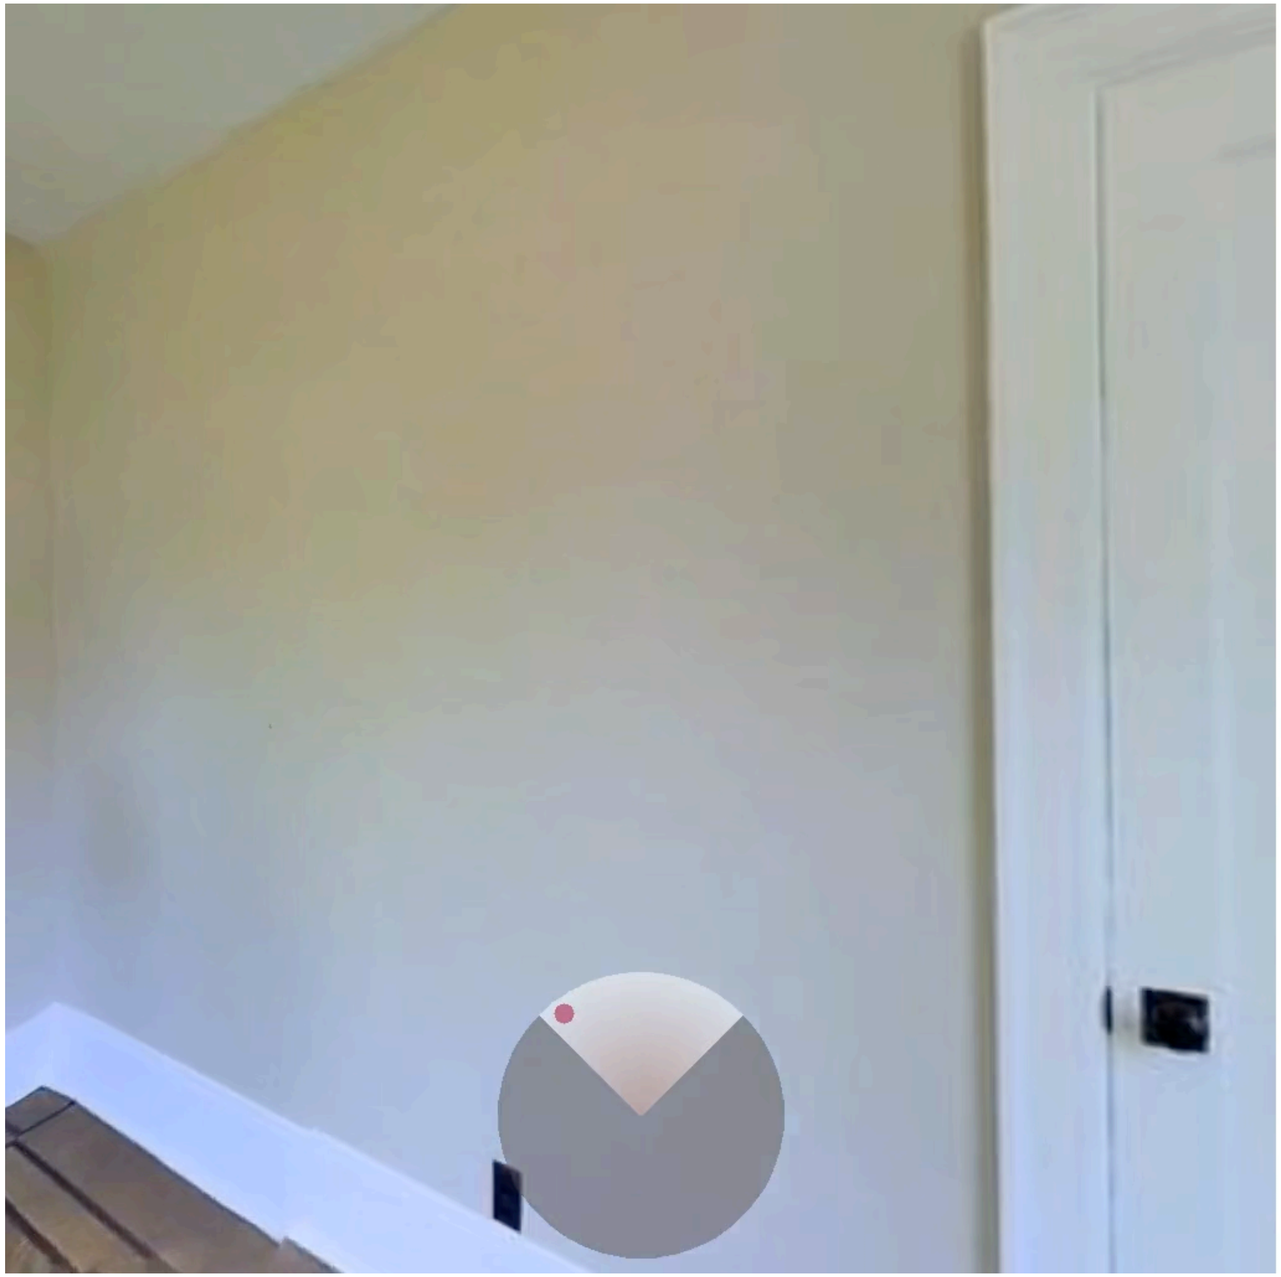

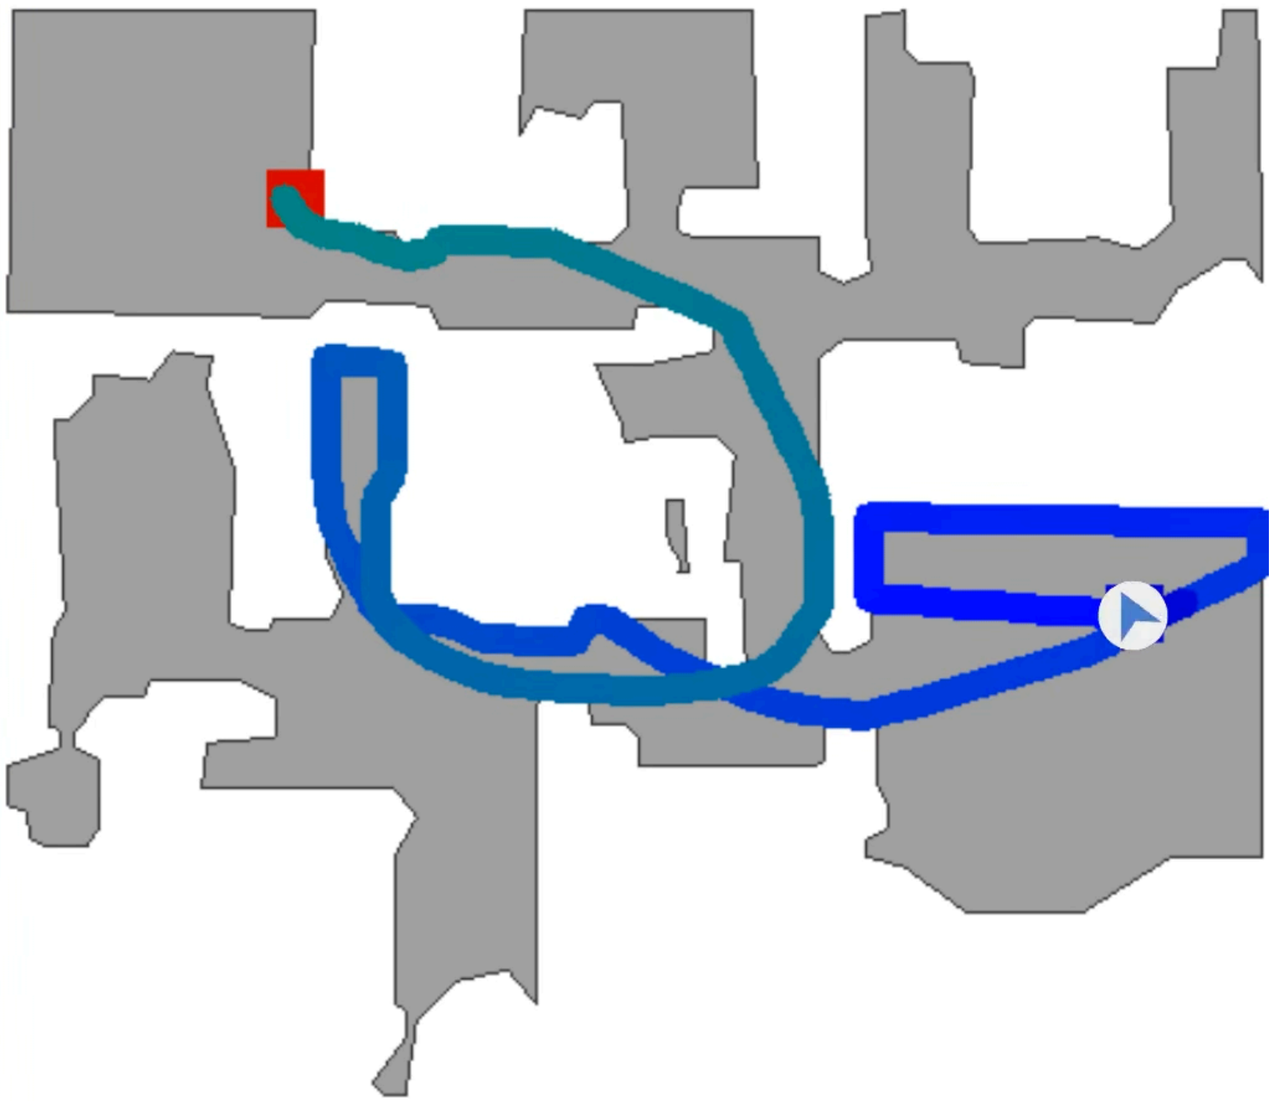

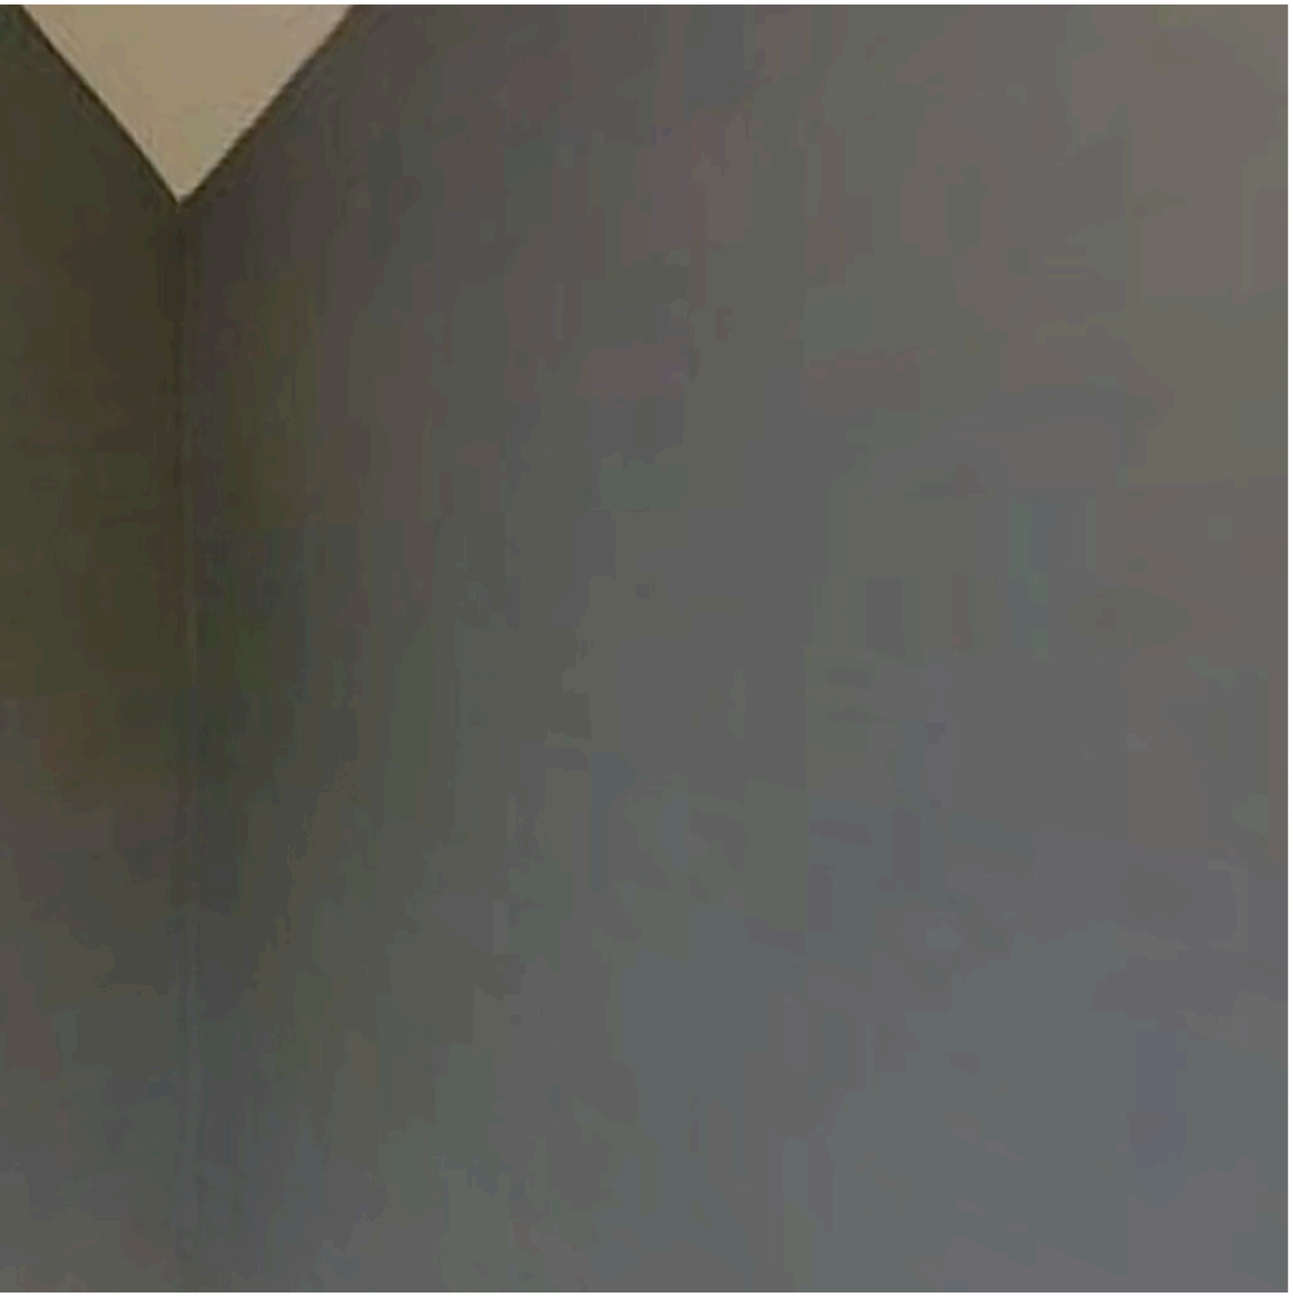

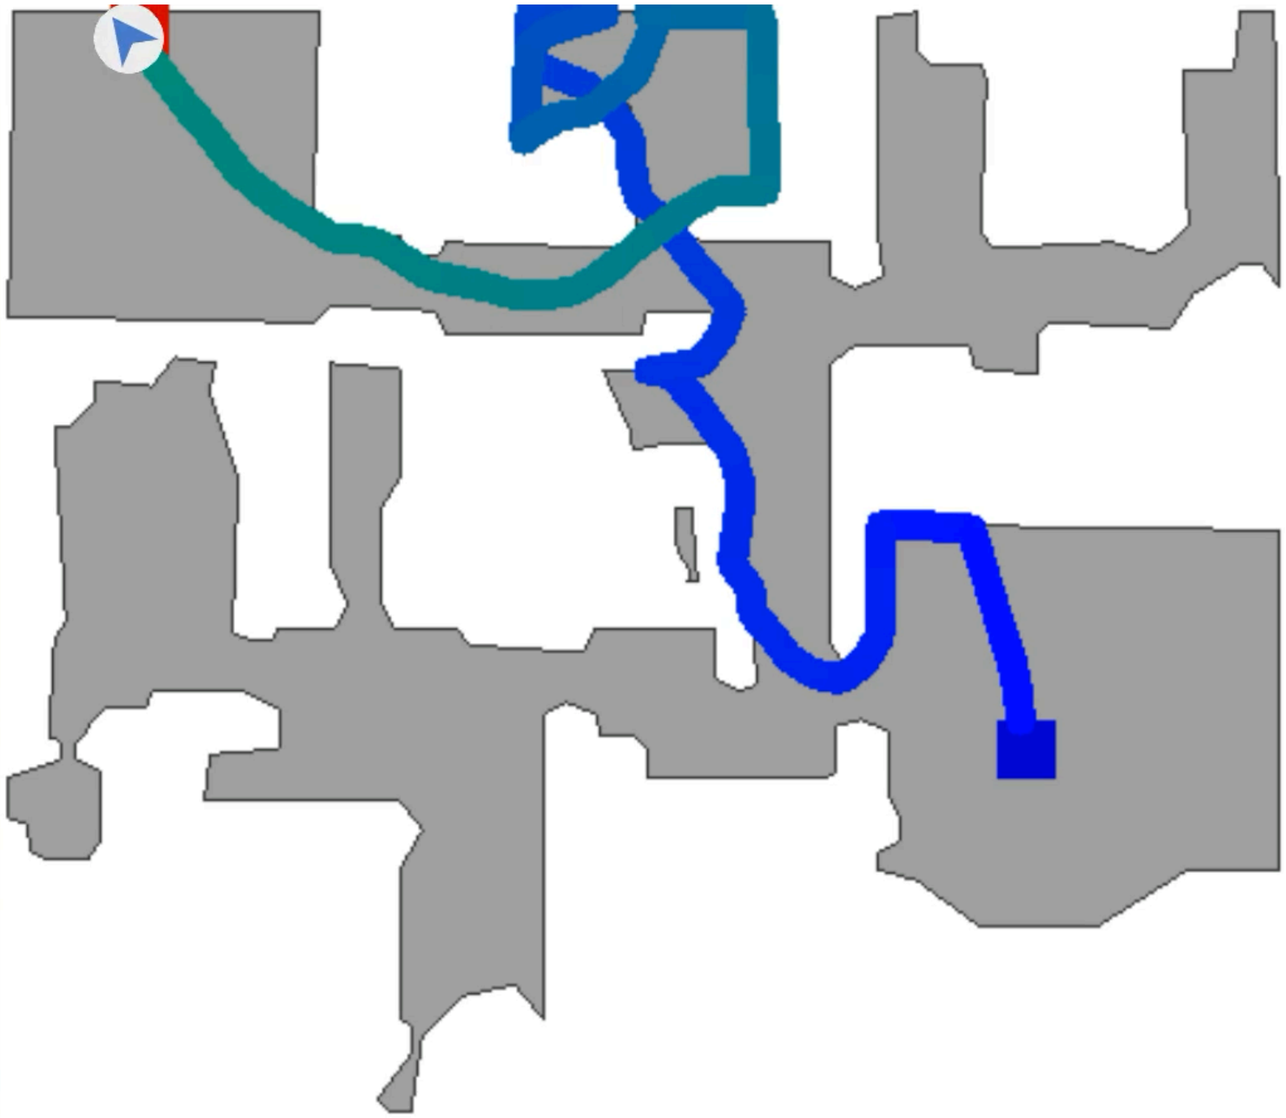

Memory is used for mapping

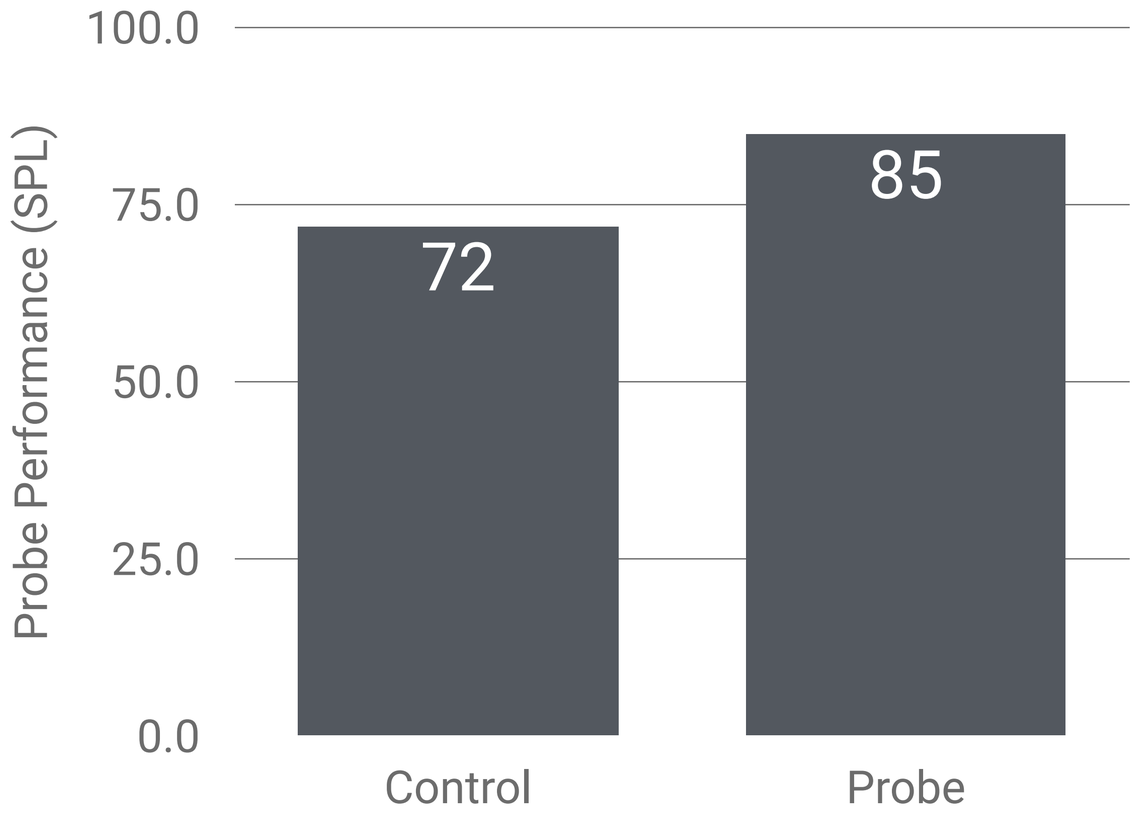

Memory is used for mapping

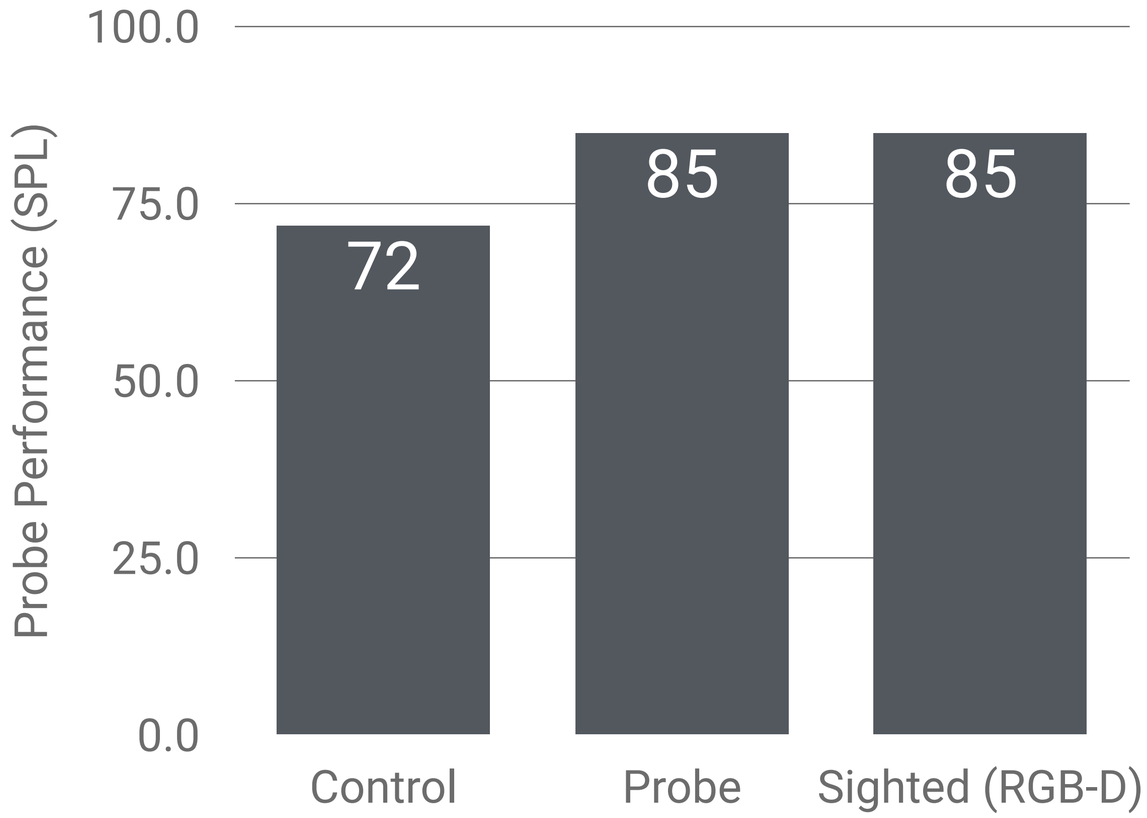

Sighted (RGB): *Ramakrishnan et al, 2021*

Decoding a map

Decoding a map

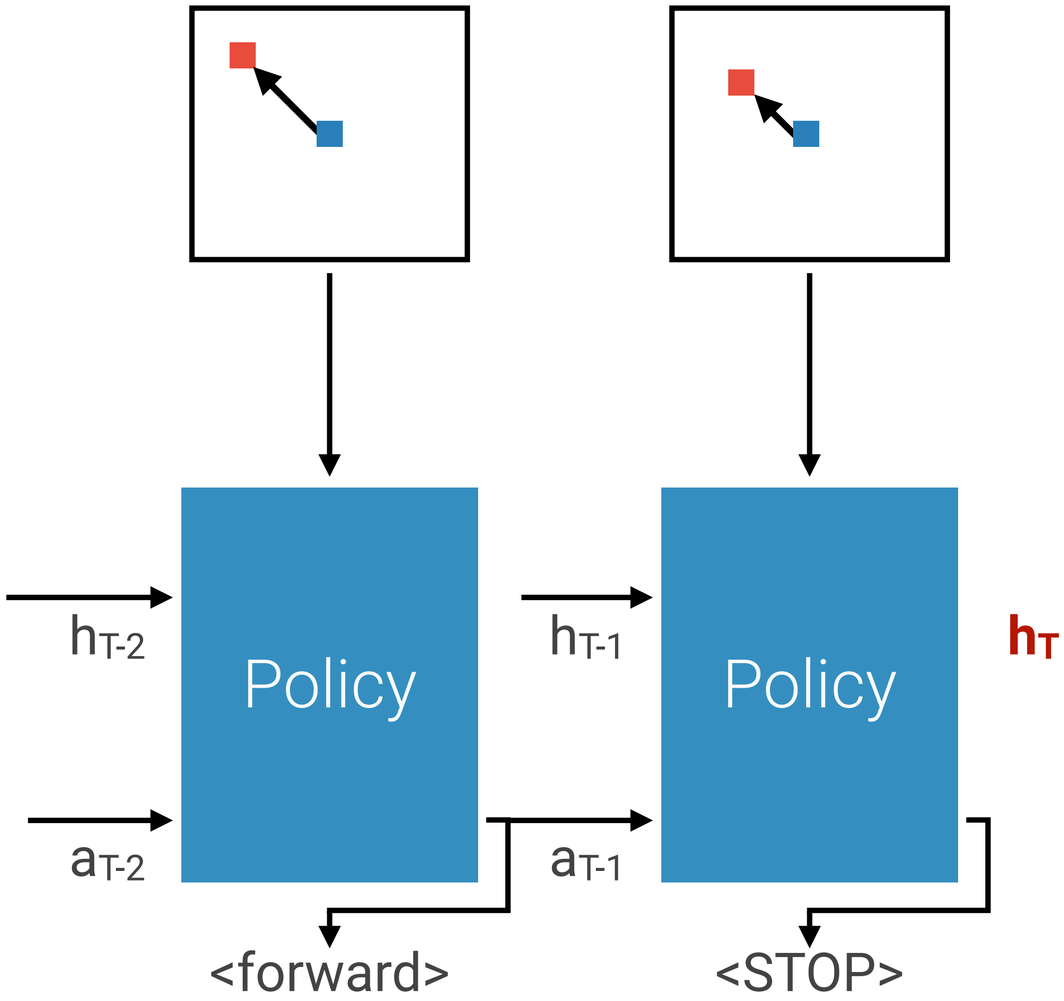

Decoding a map

h_T

Decoding a map

h_T

Decoding a map

h_T

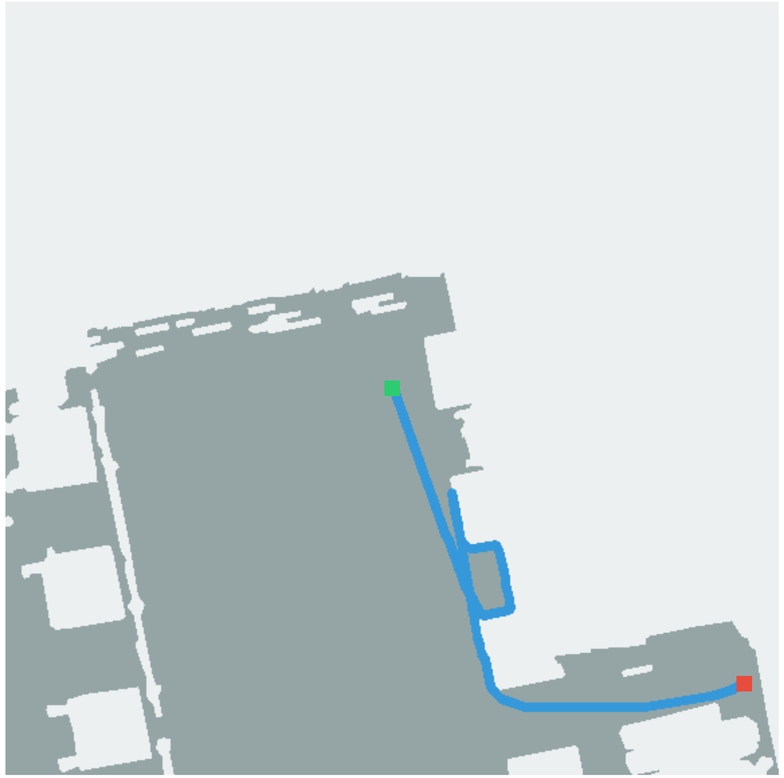

Non-navigable Navigable

Decoding a map

h_T

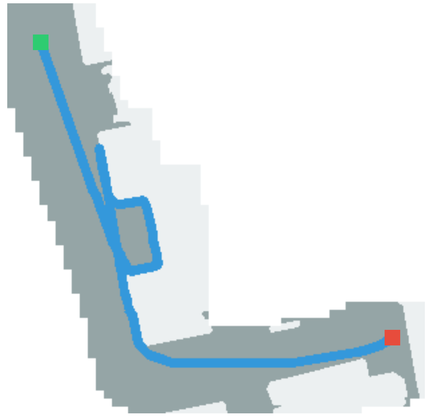

Non-navigable Navigable

Decoding a map

h_T

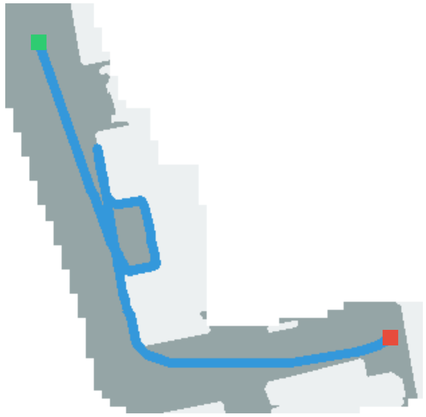

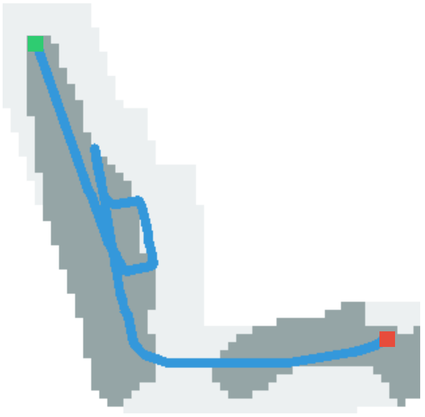

Non-navigable Navigable

Decoding a map

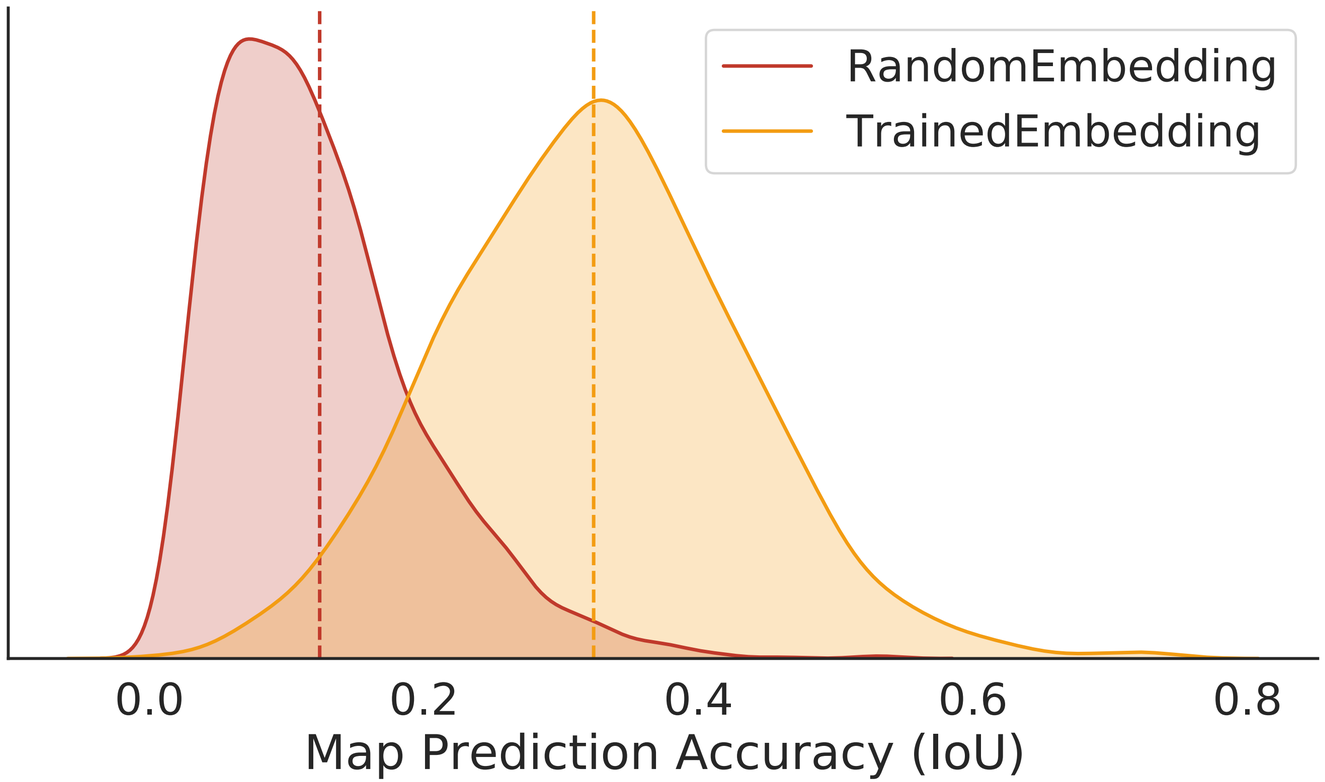

Memory is task dependent

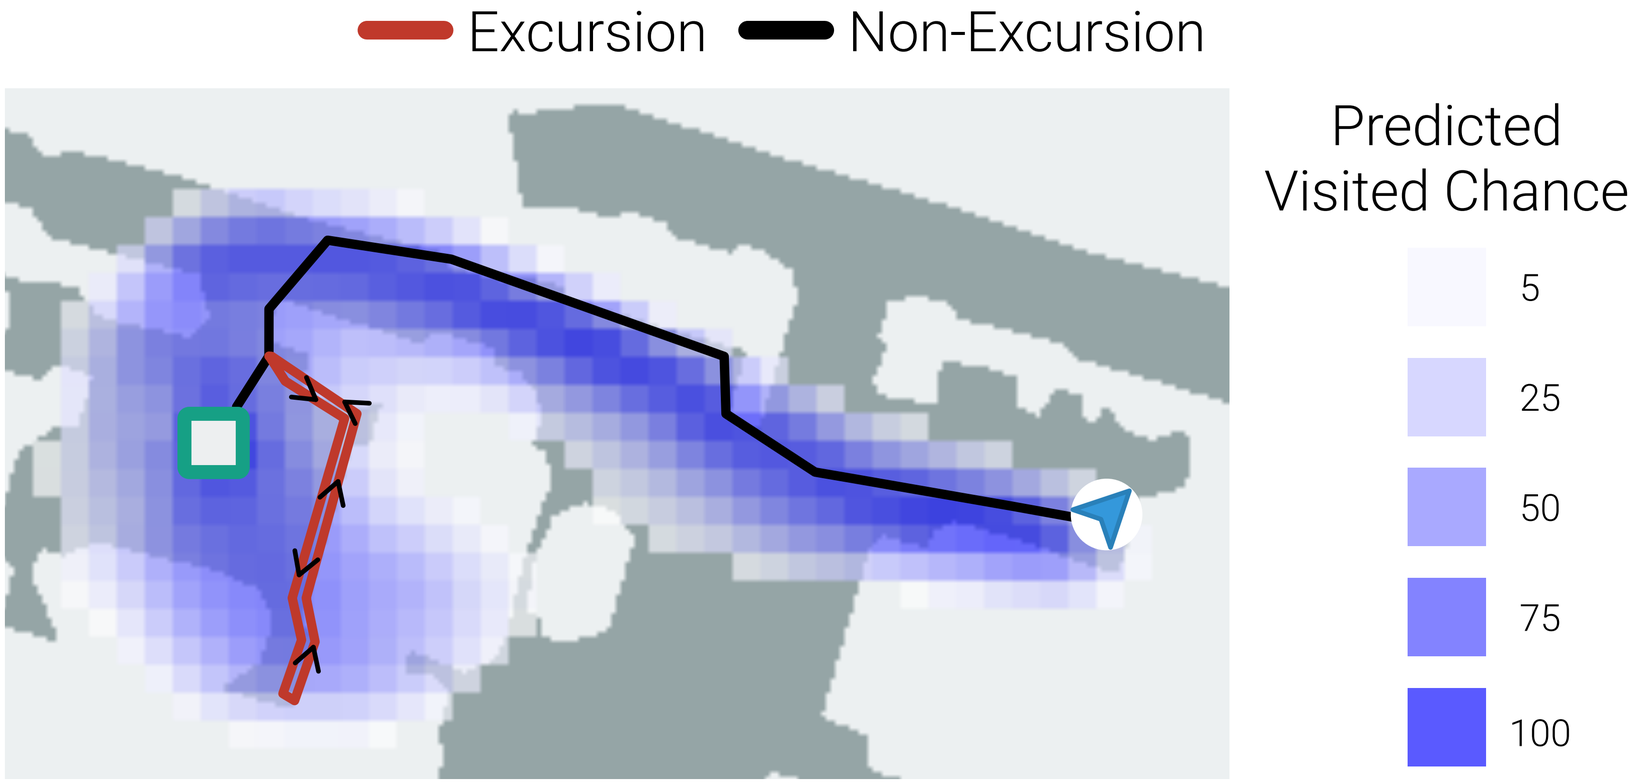

Discussion

Discussion: Limitations

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors
- Blind agents: Inconclusive results with sight. Conjecture that blind requires high-level solutions

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors
- Blind agents: Inconclusive results with sight. Conjecture that blind requires high-level solutions
- Noiseless observations and actuations

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors
- Blind agents: Inconclusive results with sight. Conjecture that blind requires high-level solutions
- Noiseless observations and actuations
- Only examine agents with an implicit map-building mechanism

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors
- Blind agents: Inconclusive results with sight. Conjecture that blind requires high-level solutions
- Noiseless observations and actuations
- Only examine agents with an implicit map-building mechanism
- Deployment for a short period of time (order of minutes)

Discussion: Limitations

- Real houses: 2D manifold and strong structural priors
- Blind agents: Inconclusive results with sight. Conjecture that blind requires high-level solutions
- Noiseless observations and actuations
- Only examine agents with an implicit map-building mechanism
- Deployment for a short period of time (order of minutes)
- No mechanistic account nor complete account of all in memory

Discussion: World Models

Discussion: World Models

- Are world models emergent phenomena?

Discussion: World Models

- Are world models emergent phenomena?
- Or should we design them?

Discussion: World Models

- Are world models emergent phenomena?
- Or should we design them?
- Do large language models (LLMs) have world models?

Discussion: World Models

- Are world models emergent phenomena?
- Or should we design them?
- Do large language models (LLMs) have world models?
- We show that mapping is an emergent phenomena

Discussion: World Models

- Are world models emergent phenomena?
- Or should we design them?
- Do large language models (LLMs) have world models?
- We show that mapping is an emergent phenomena
  - A map is a key component

Discussion: World Models

- Are world models emergent phenomena?
- Or should we design them?
- Do large language models (LLMs) have world models?
- We show that mapping is an emergent phenomena
  - A map is a key component
  - and decodable

Poster: #106

Conclusion

Emergence of Maps in the Memories of Blind Navigation Agents

Effective blind navigation is possible
“Blind” (localization-only) navigation can be performed effectively, but not efficiently

Enabled by memory and collision detection neurons
These agents rely heavily on memory and collision detection neurons emerge

Emergence of maps
They use their memory to build a map of their environment

Maps are task specific
There is less detail about excursions