Robot Navigation: From Abilities to Capabilities

Machine Learning in Robot Motion Planning Workshop @ IROS 2018

Aleksandra Faust, Ph.D

Google Brain Robotics

October 2018

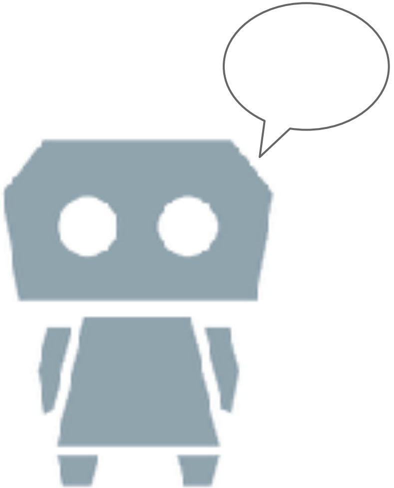

Google

Navigation

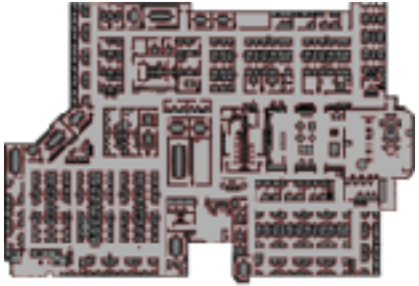

Go down the
hallway and take
the second right.

Google

Navigation

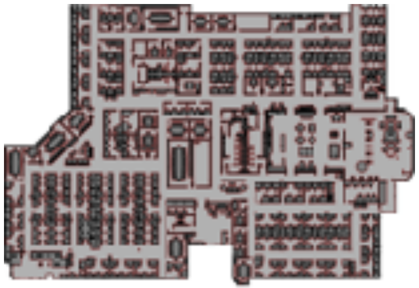

Go down the hallway and take the second right.

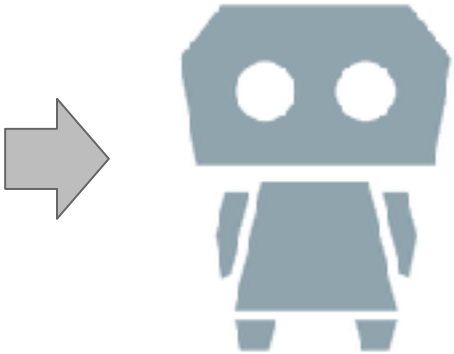

Perception

Slam

Planning

Controls

Google

Navigation

Sensors, hardware, geometry, determine how robot perceives the world

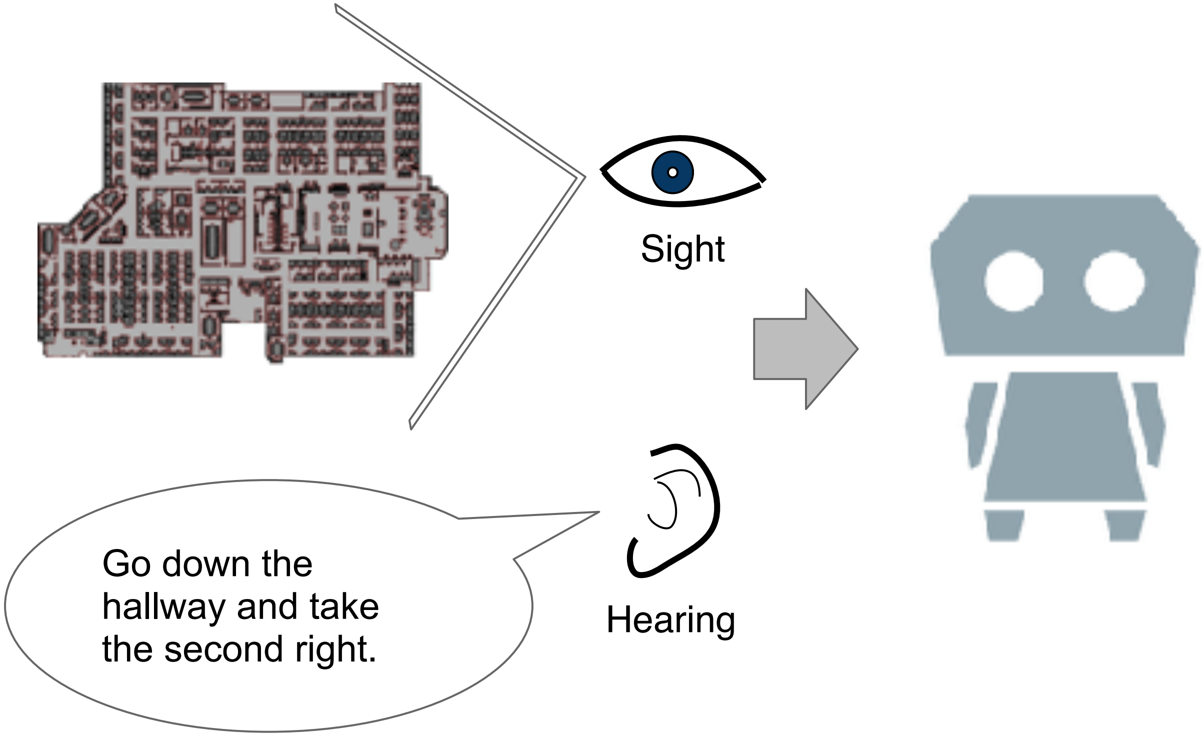

Navigation

And what it can do. How it can communicate, move, and interact.

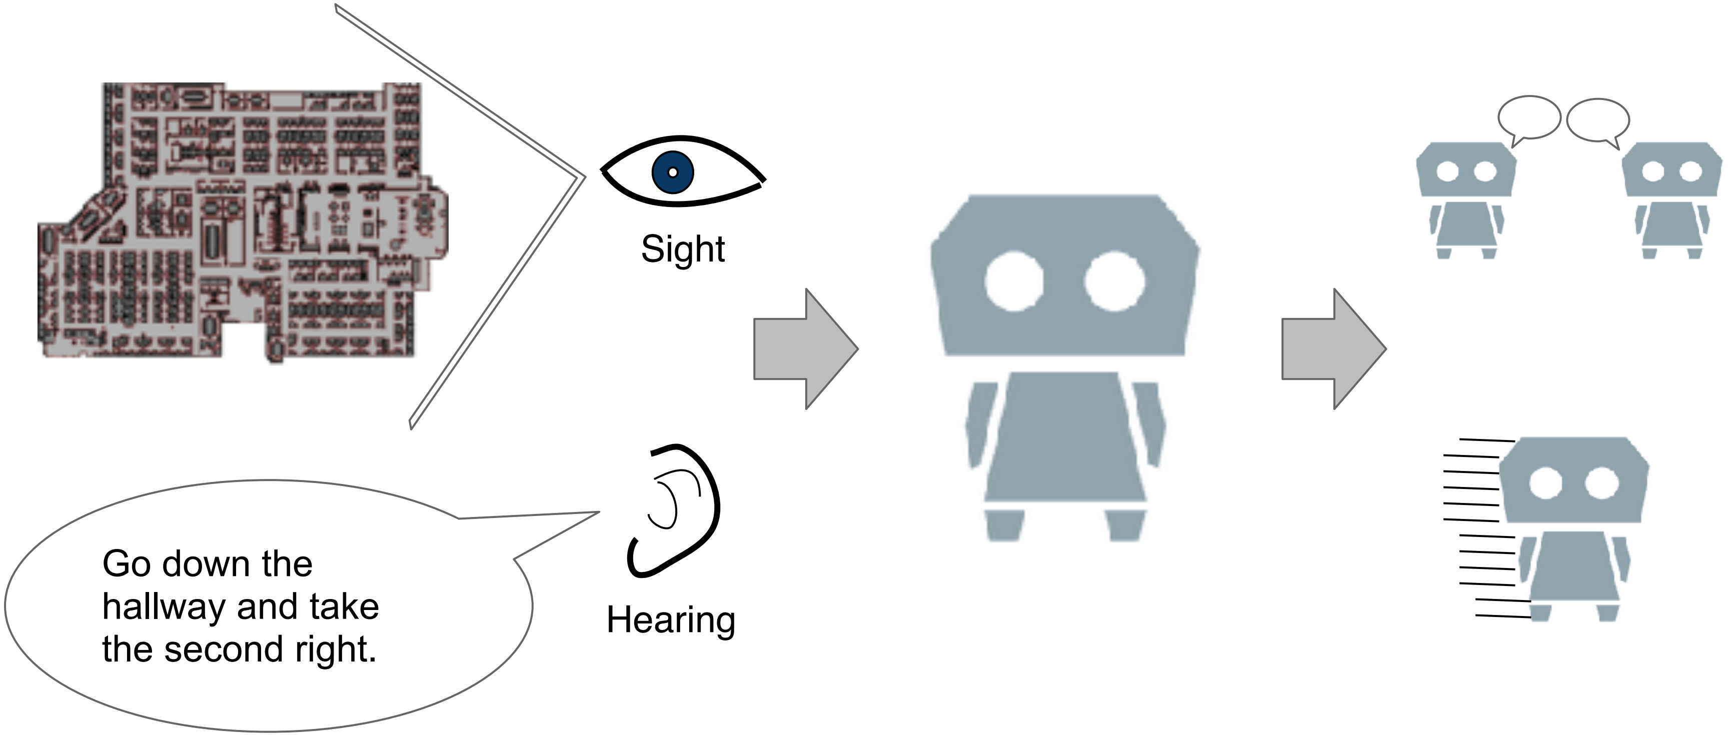

Abilities

Let’s focus on the robot abilities, and learn the foundational motion behaviors.

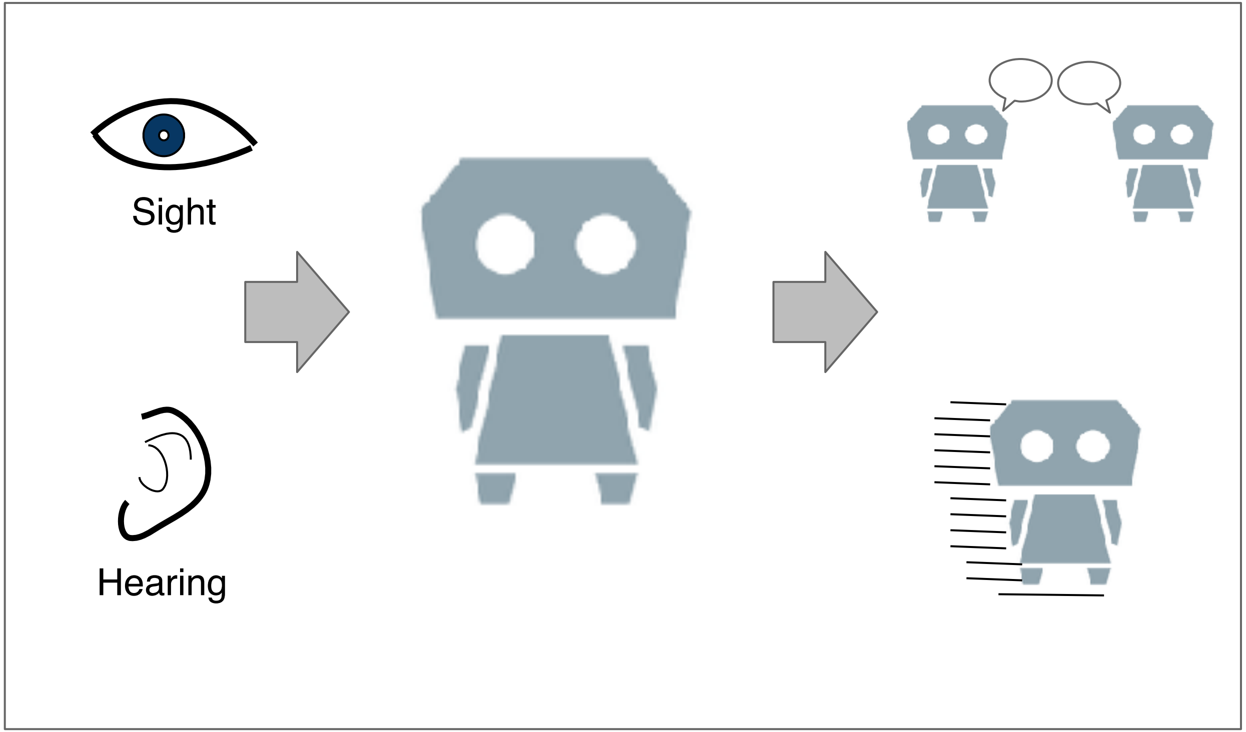

Google

Essential Navigation behavior

Moving obstacle avoidance

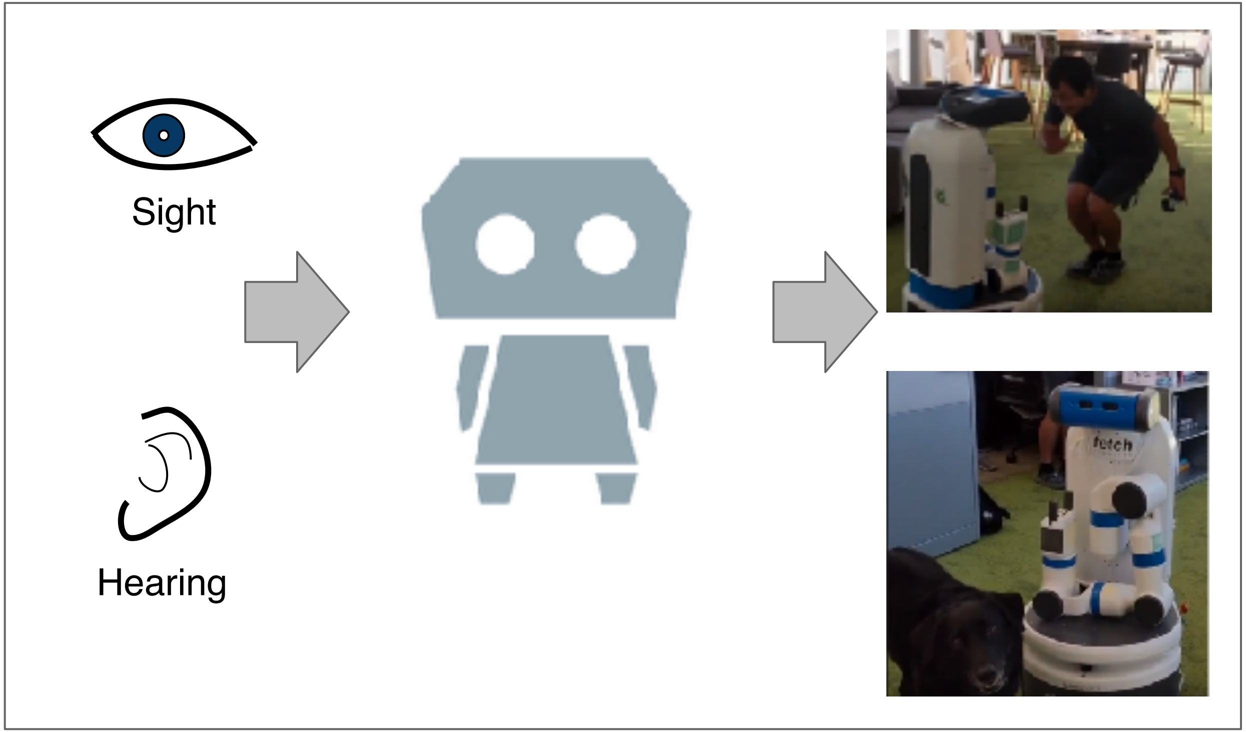

Navigation behaviors based on robot's abilities

Moving obstacle avoidance for real robots

- With primitive sensors
- Robust to noise
- Dynamically feasible
- Transfers between environments

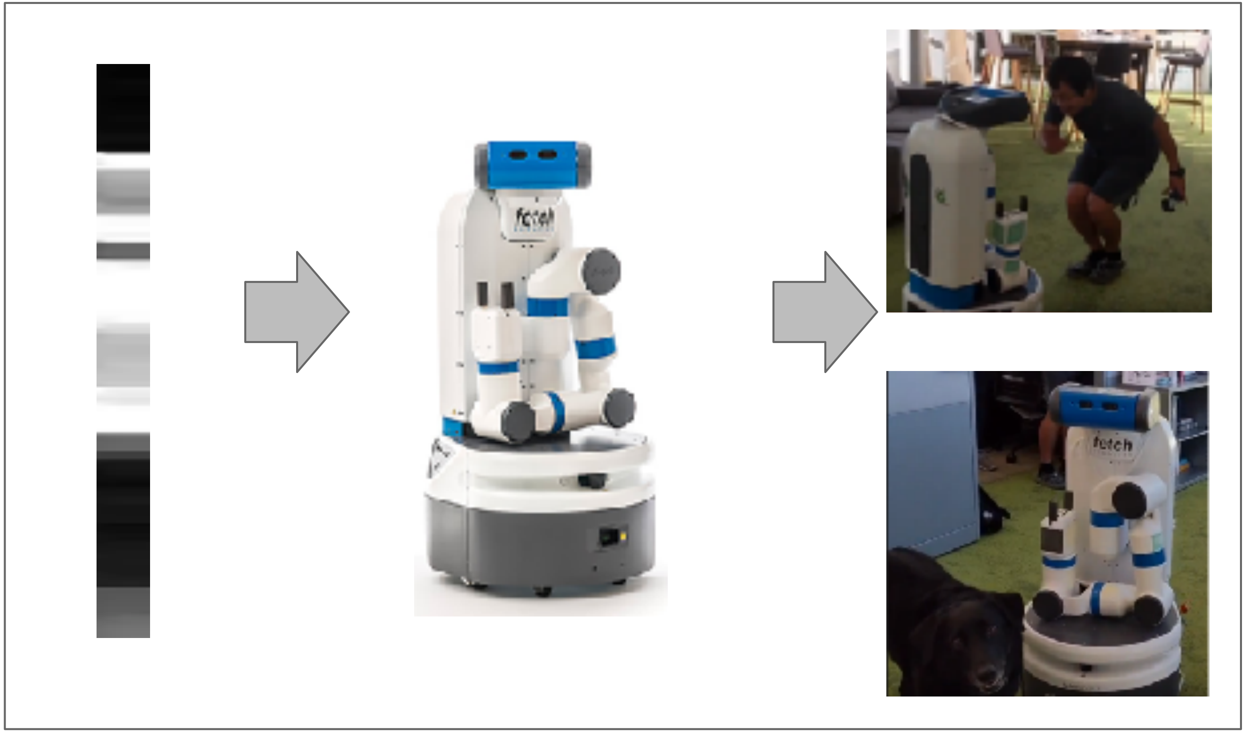

Google

Navigation behaviors based on robot's abilities

Moving obstacle avoidance for real robots

- With primitive sensors
- Robust to noise
- Dynamically feasible
- Transfers between environments

Behaviors to learn:

- Point to point navigation
- Path following

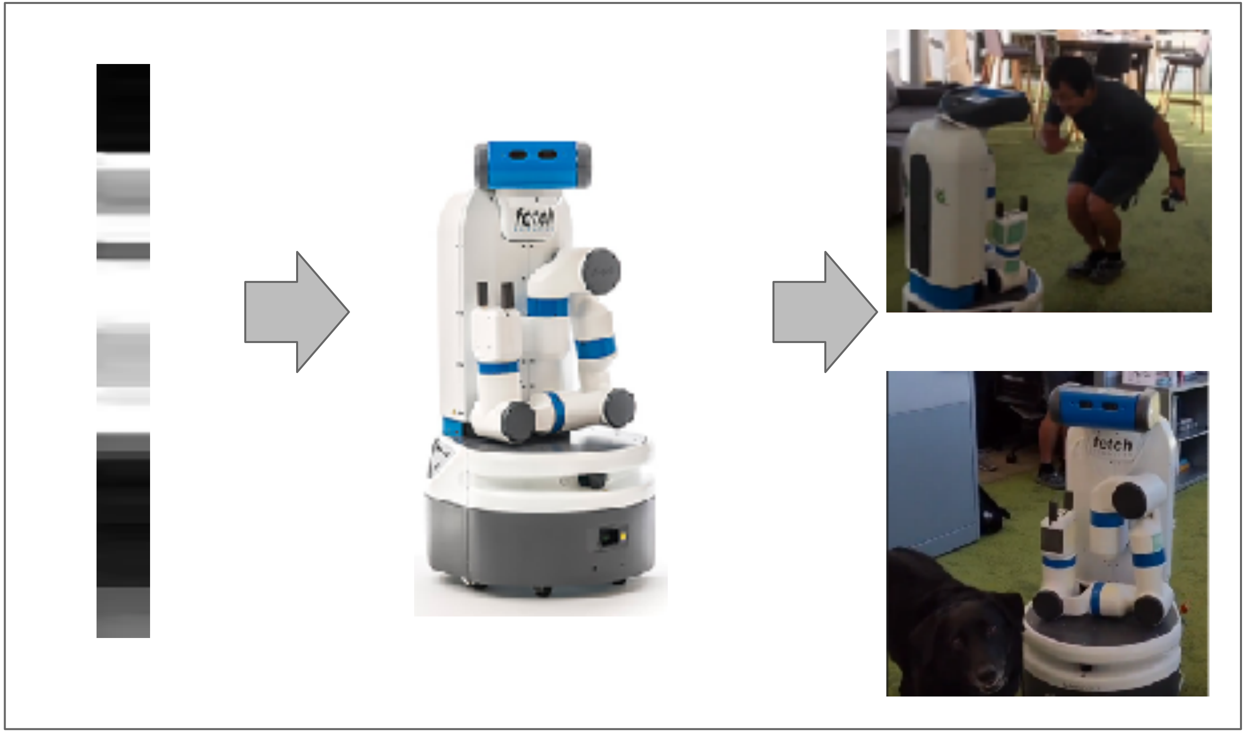

Learning navigation behaviors end to end

Under submission, https://arxiv.org/abs/1809.10124

Hao-Tian Lewis Chiang*, Aleksandra Faust*, Marek Fiser, Anthony Francis

*Equal contributions

Hao-Tien Chiang

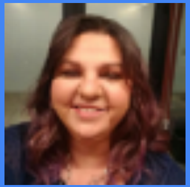

Aleksandra Faust

Marek Fiser

Anthony Francis

Learning navigation task

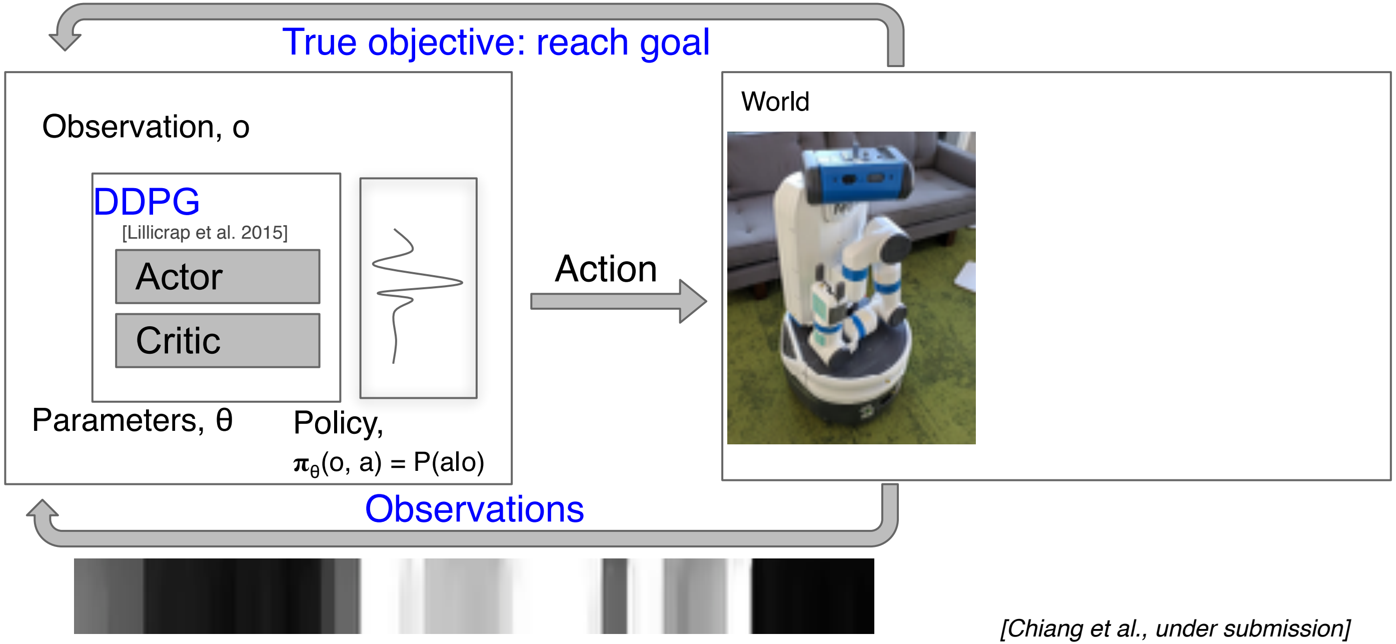

[Chiang et al., under submission]

Learning navigation task

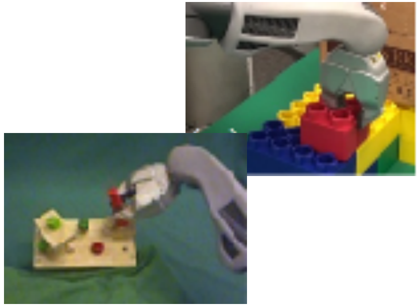

[Levine et al., ‘16]

[Yahya et al., ‘16]

[Mülling et al., ‘11]

[Faust et al, ‘14]

[Bagnell and Schneider ‘01]

Handles sensor input
Handles task and robot dynamics

Helicopter image from [Kober et al, ‘13]

Google

Learning Navigation Task

[Chiang et al., under submission]

RL Setup is Hard

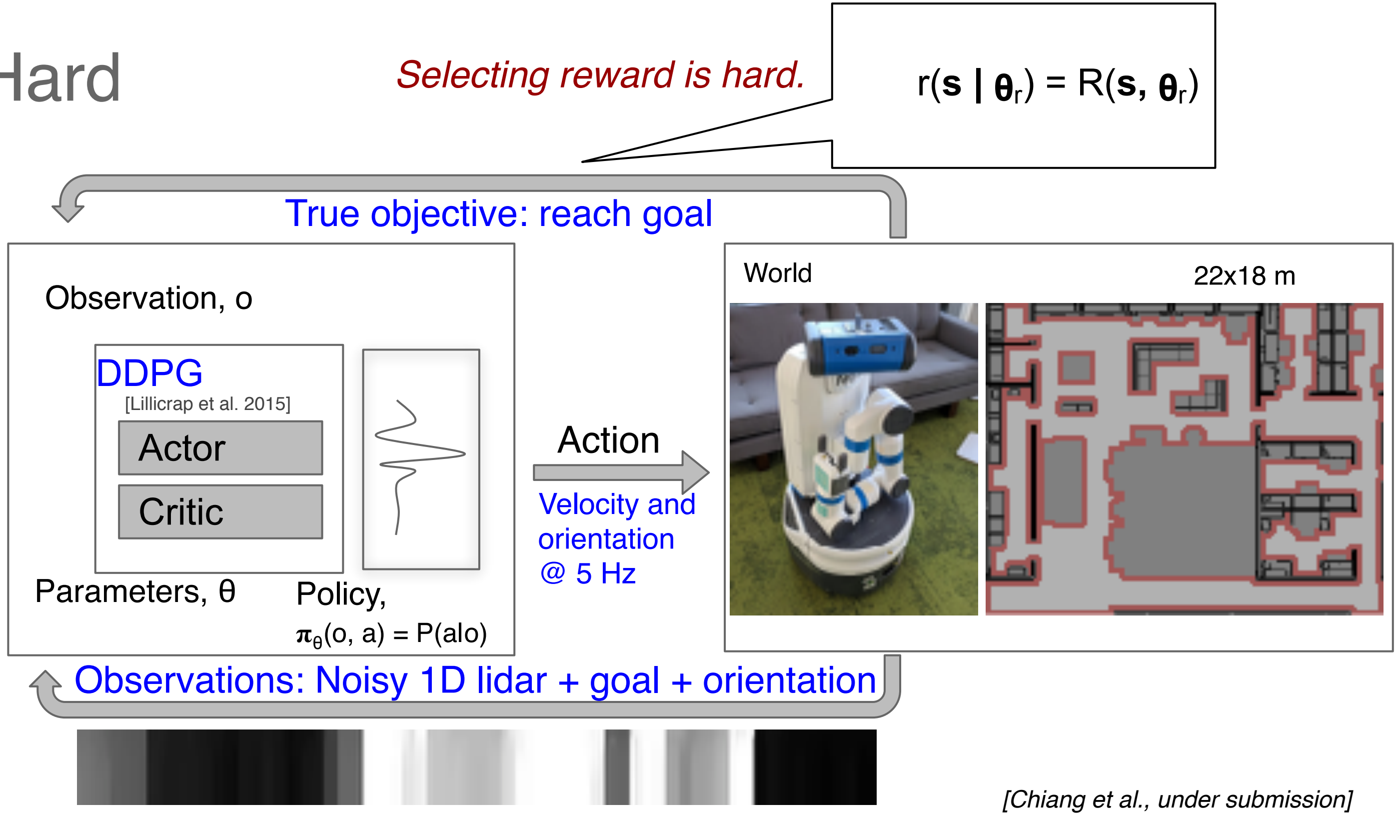

[Chiang et al., under submission]

RL Setup is Hard

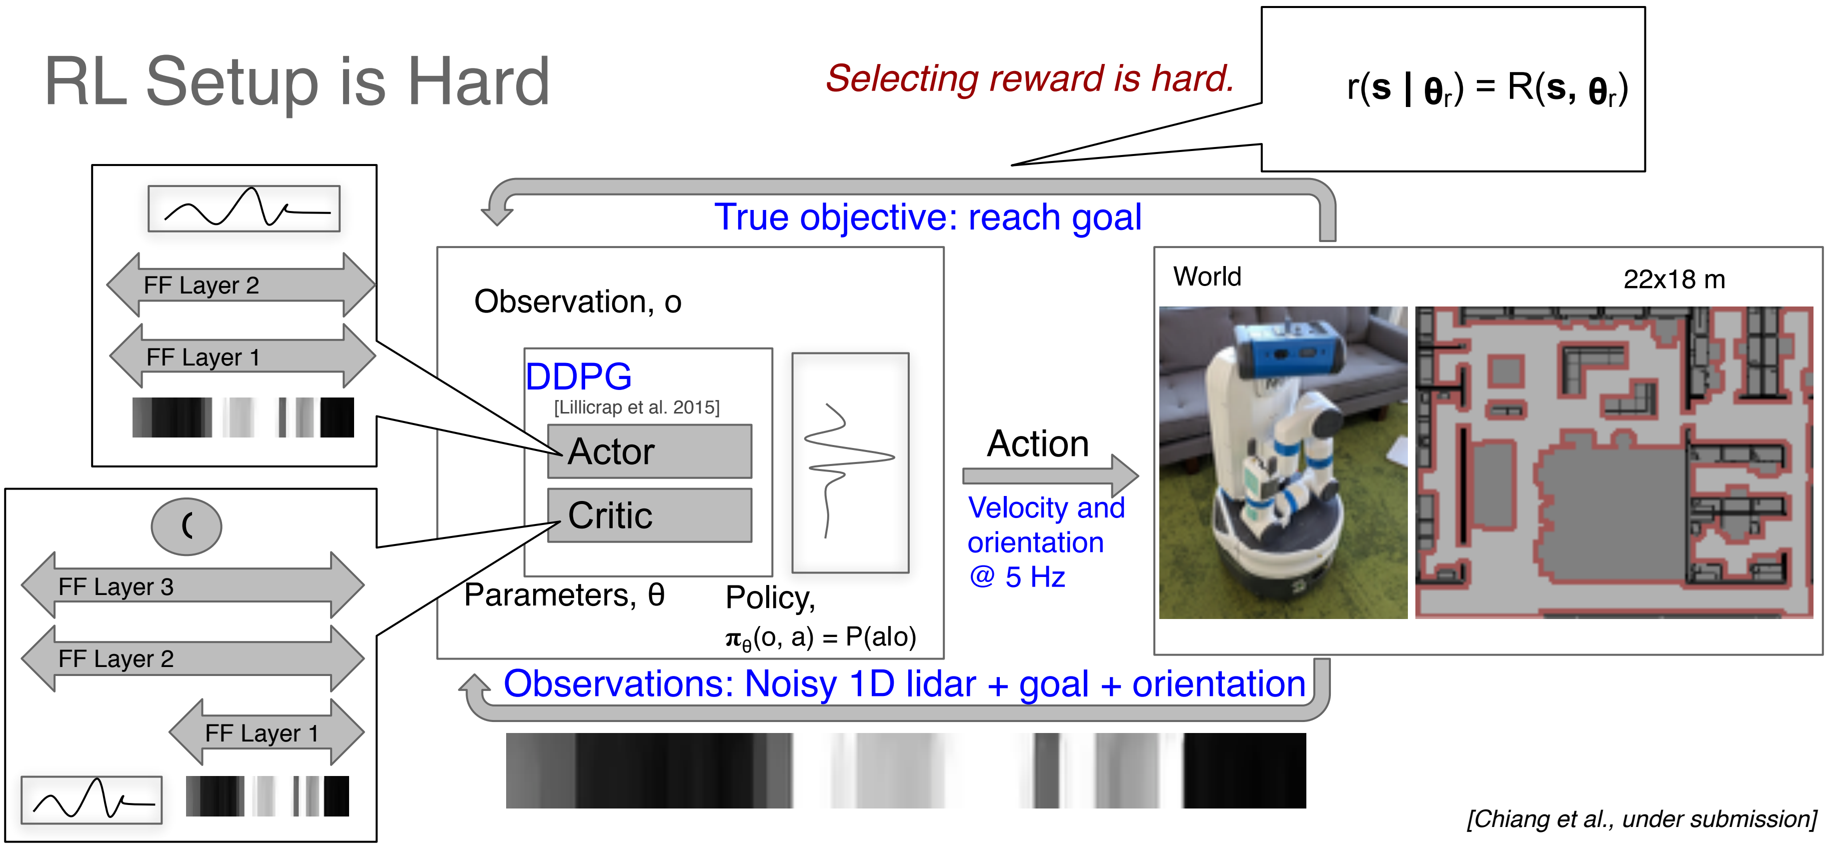

Network architecture selection is hard.

Shaped-DDPG

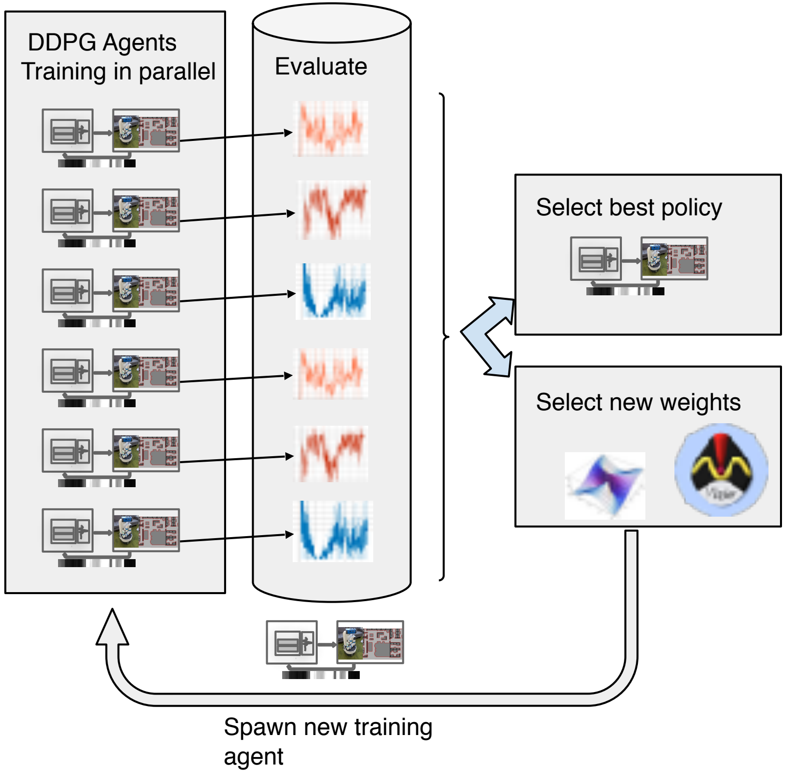

Solution: large-scale gradient-free hyper-parameter optimization.

Shaped-DDPG

[Chiang et al., under submission]

Find the **best reward** function that maximizes the **true objective**

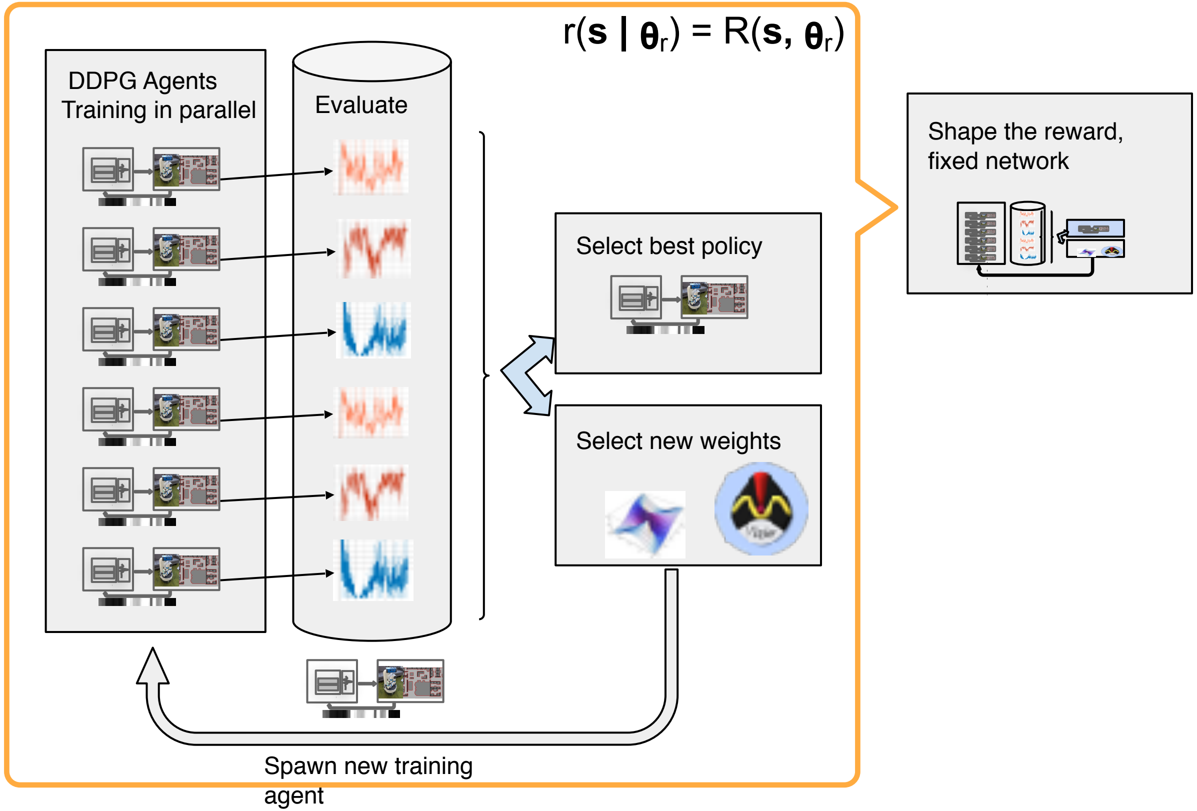

Solution: large-scale gradient-free optimization.

Shaped-DDPG

[Chiang et al., under submission]

Find the **best reward** function that maximizes the **true objective**

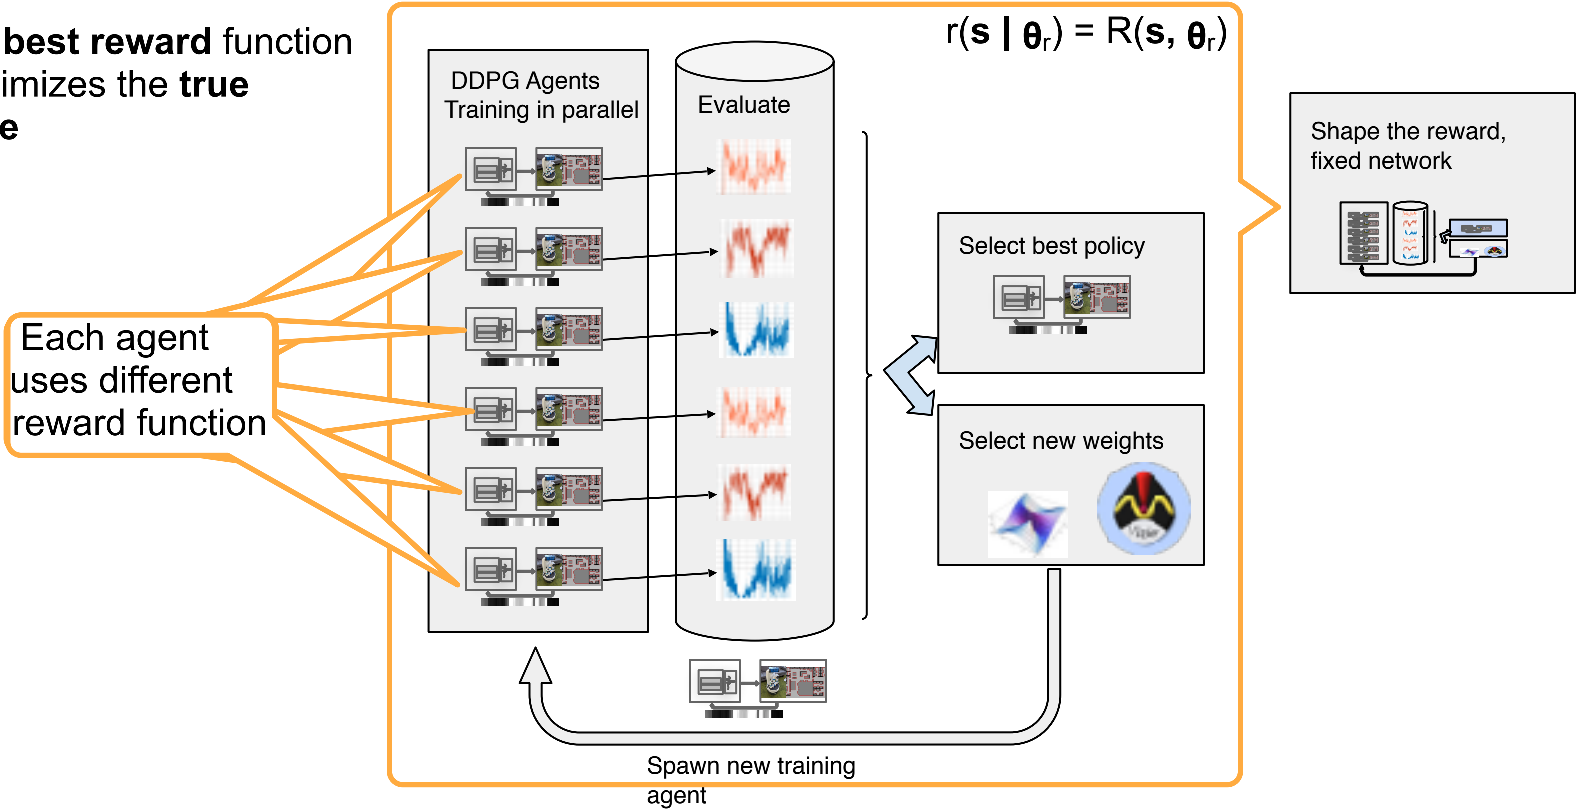

Solution: large-scale gradient-free optimization.

Shaped-DDPG

[Chiang et al., under submission]

Find the **best reward** function that maximizes the **true objective**

Find the **best NN** architecture, that maximizes the **reward**

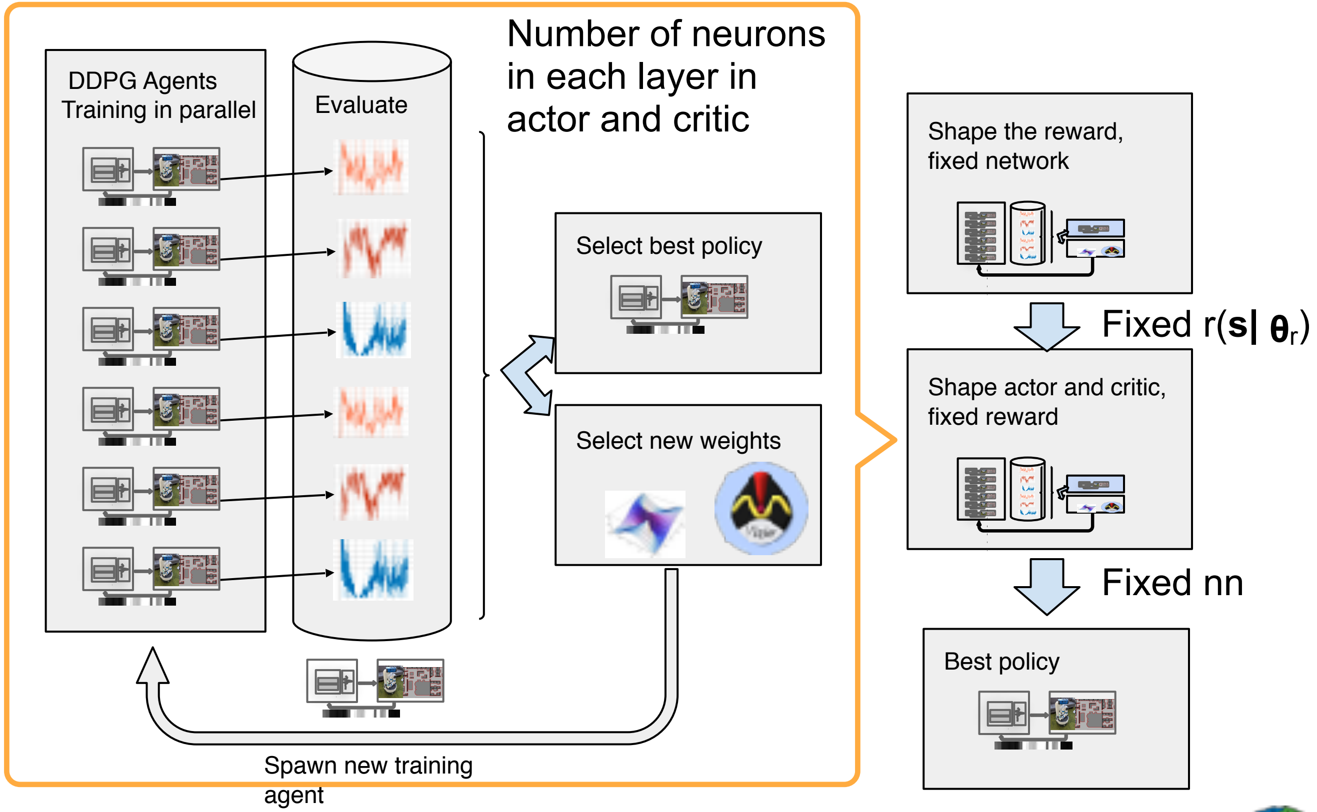

Solution: large-scale gradient-free optimization.

Shaped-DDPG

[Chiang et al., under submission]

Find the **best reward** function that maximizes the **true objective**

Find the **best NN** architecture, that maximizes the **reward**

Each agent uses different neural network architecture.

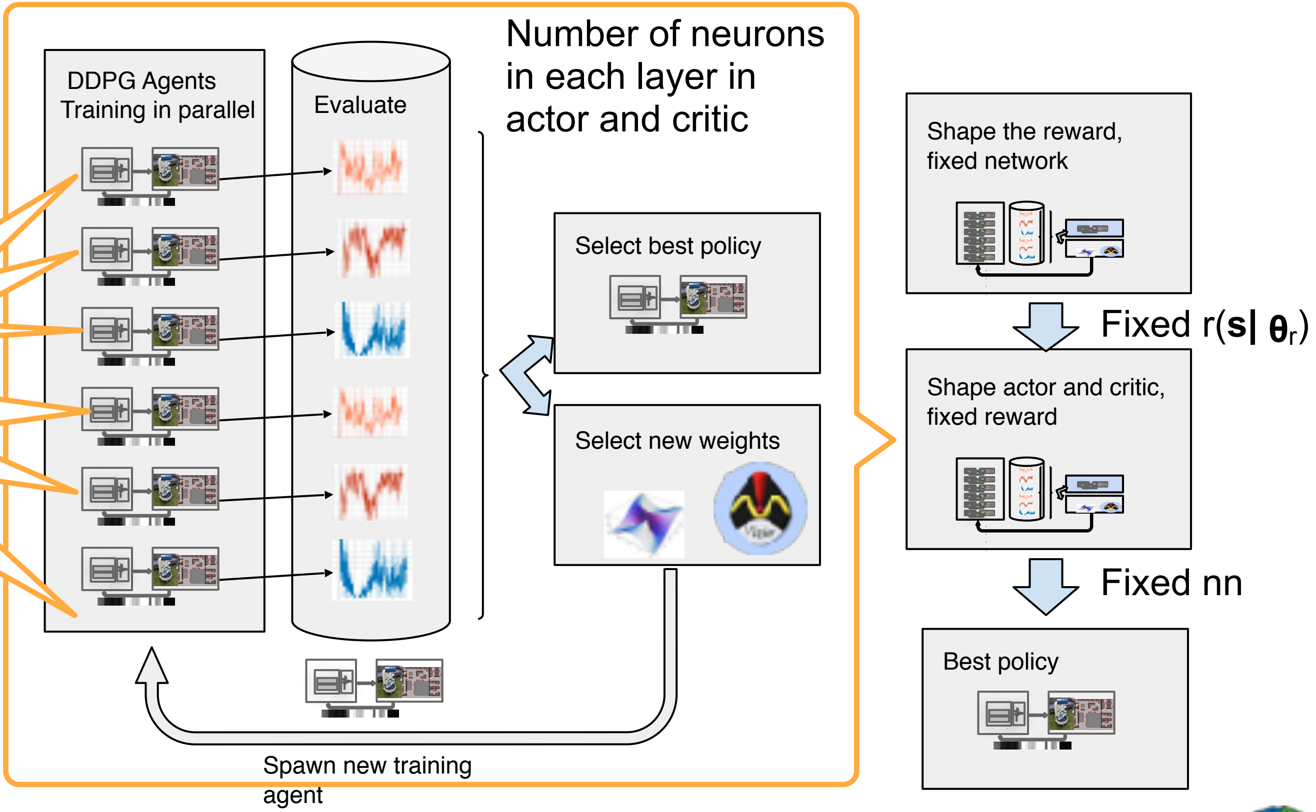

Solution: large-scale gradient-free optimization.

Google

Shaped-DDPG Learning Results for Path Following

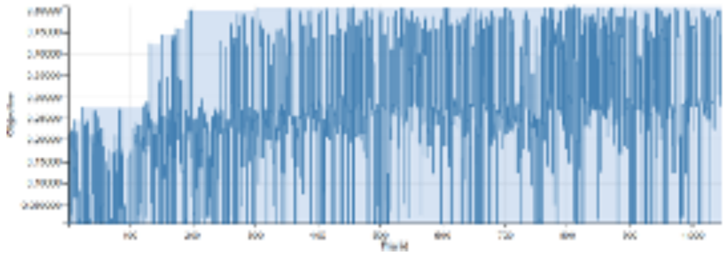

Reward only shaping

1000 trails, 5 million steps each @ 5Hz - trains in a week

[Chiang et al., under submission]

Google

Shaped-DDPG Learning Results for Path Following

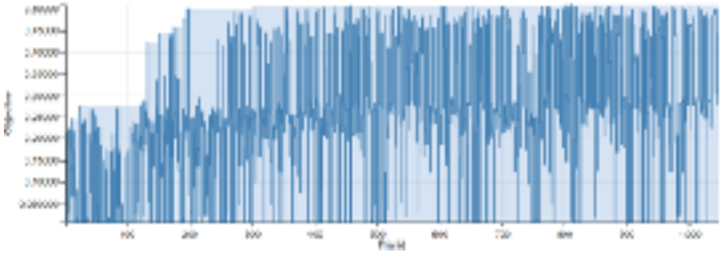

Reward only shaping

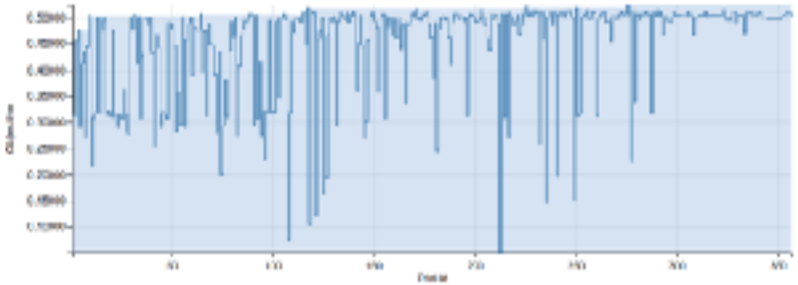

Reward and NN shaping

1000 trails, 5 million steps each @ 5Hz - trains in a week.

[Chiang et al., under submission]

Stable learning, consistent trials

Shaped-DDPG Learning Results for Path Following

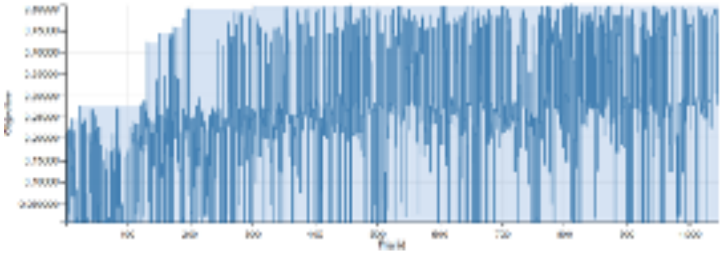

Reward only shaping

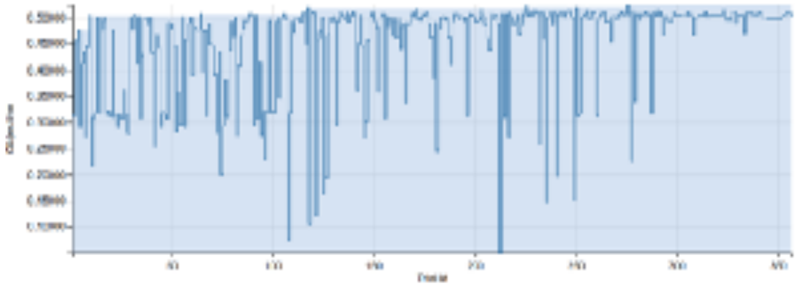

Reward and NN shaping

1000 trails, 5 million steps each @ 5Hz - trains in a week

[Chiang et al., under submission]

Equivalent of 32 years of collective experience.
12 days each trial, learning from previous generations

Stable learning, consistent trials

Google

Shaped-DDPG Learning Results for Path Following

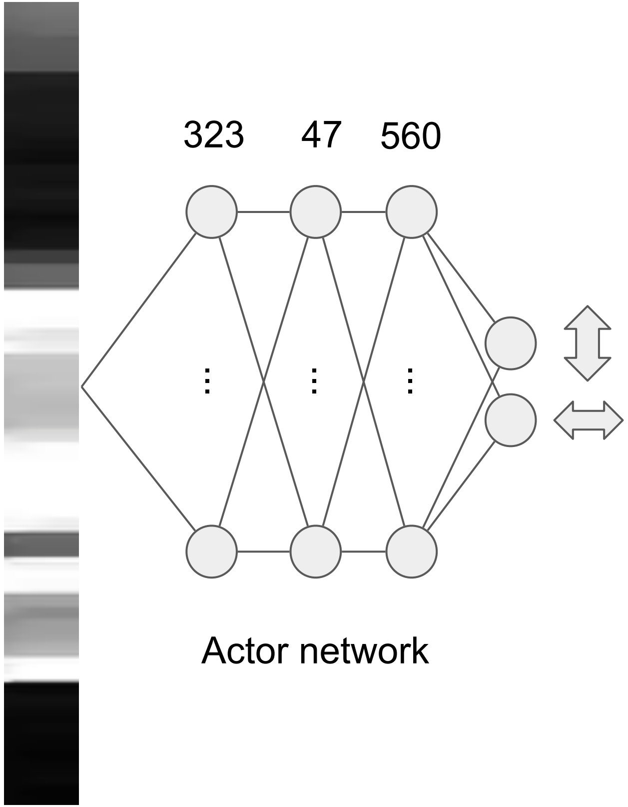

Actor network

[Chiang et al., under submission]

Google

Shaped-DDPG Evaluation

- Two baselines
  - Learned: Vanilla DDPG
  - Classic: Artificial potential fields
- Evaluation environments
  - 3 large buildings
  - With moving obstacles

Building 1, 183 by 66m

Training, 23 by 18m

Building 2, 60 by 47m

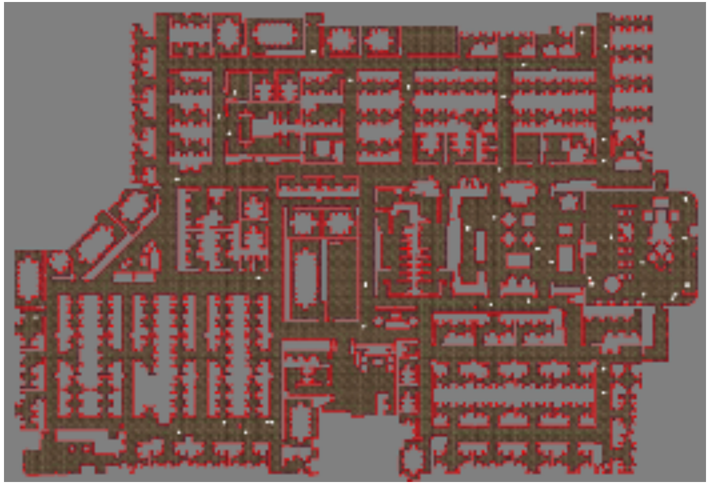

Building 3, 134 by 93m

[Chiang et al., under submission]

Google

Shaped-DDPG Evaluation: Success Rate

Higher success rate across all buildings

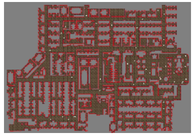

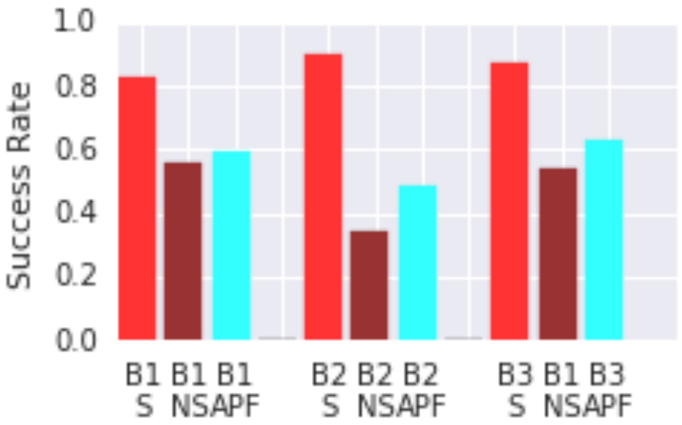

Point to point

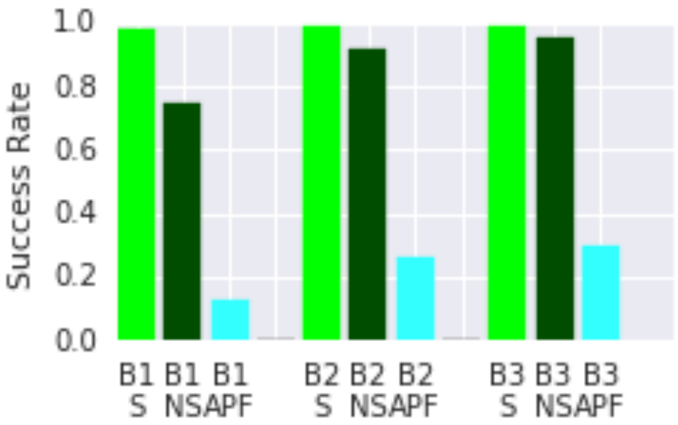

Path Following

Success rate:
shaped-DDPG, vanilla DDPG, classic APF

[Chiang et al., under submission]

Google

Shaped-DDPG vs. vanilla DDPG

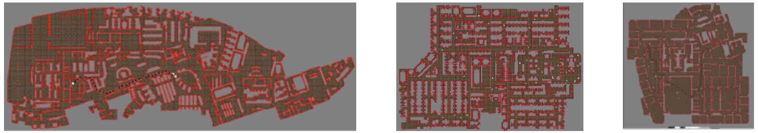

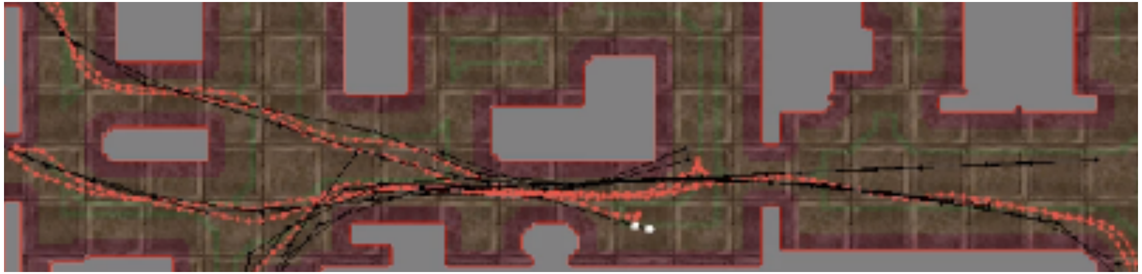

Shaped-DDPG: smooth trajectories

Vanilla DDPG: suboptimal behavior

[Chiang et al., under submission]

Google

Shaped-DDPG Evaluation: Impact of Noise

Shaped-DDPG policy is more robust to noise

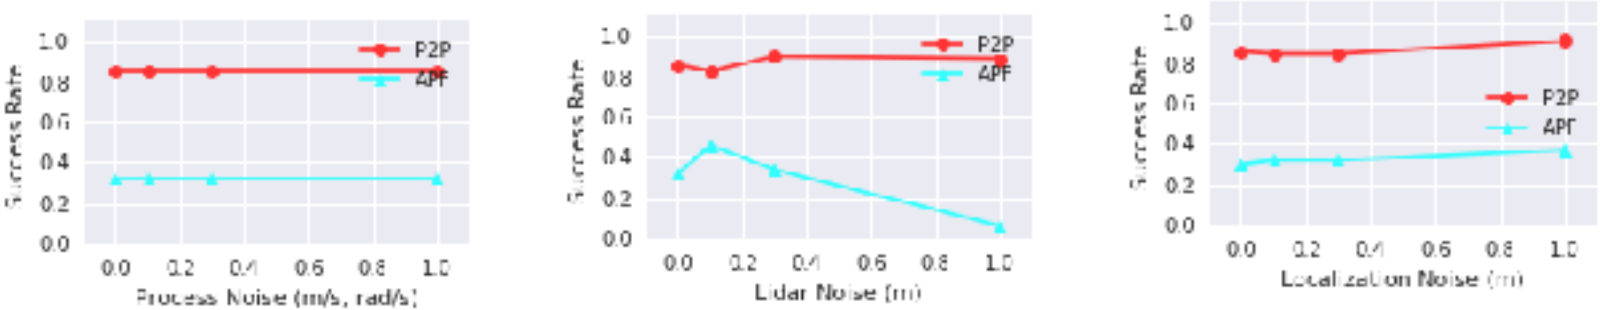

Point to Point

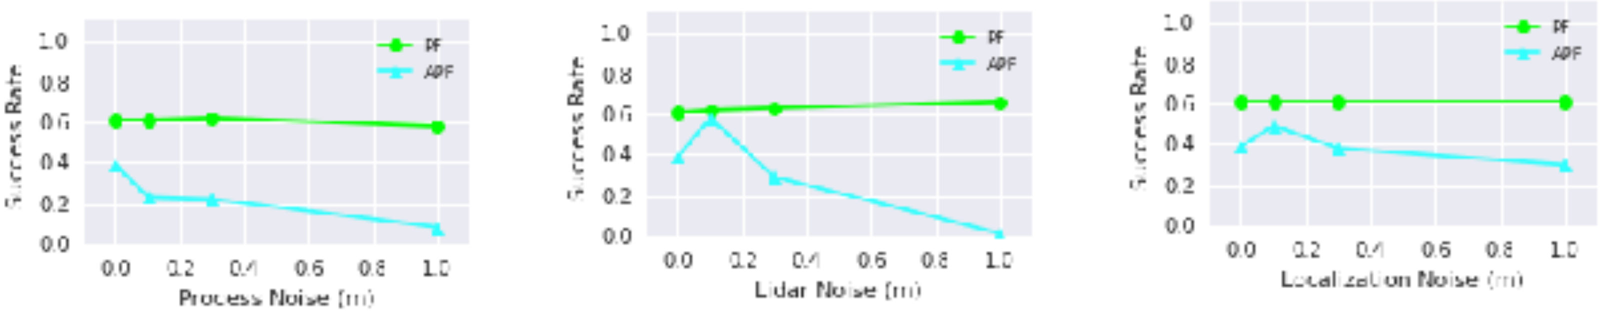

Path Following

[Chiang et al., under submission]

Google

Shaped-DDPG: On-robot experiments

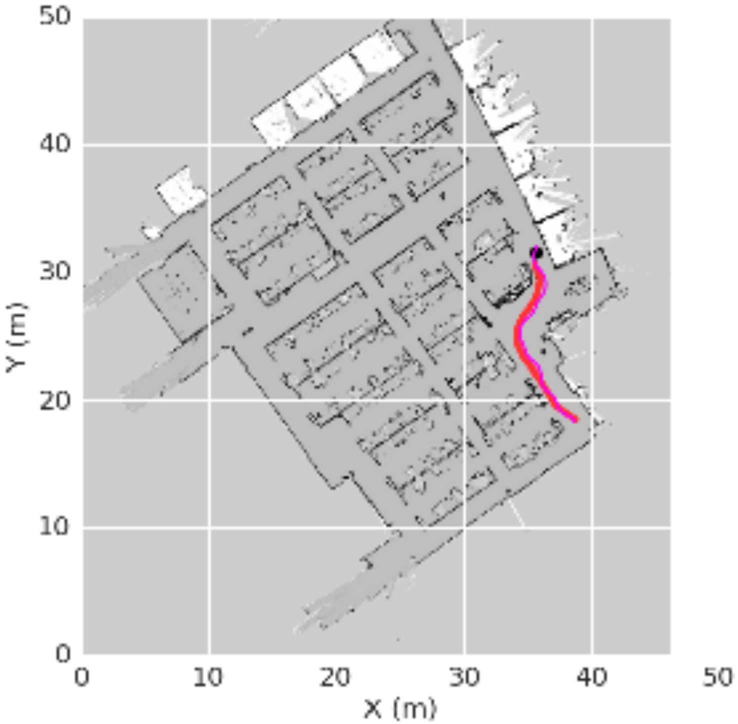

Point to point

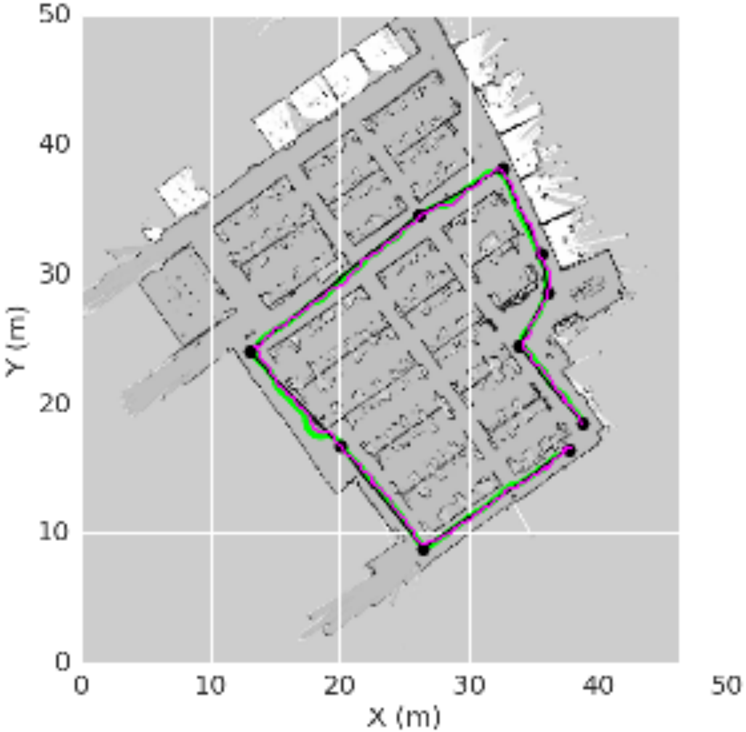

Path following

[Chiang et al., under submission]

Google


CRAFTSMAN
Corovan

Impact of Number of Obstacles

No Moving obstacle

30 moving obstacles

[Chiang et al., under submission]

Google


P2P In Dynamic Environments
PRIVATE ROAD NO THRU TRAFFIC
Real-Time

Navigation behaviors based on robot’s abilities

Learned end-end methods handle noise.

Shaping optimizes the trajectories.

Traditional methods: well behaved, brittle.

Transferable to new environments
Easy sim2real

Handles:
Sensors to controls, dynamics, noise
Obstacle avoidance

Google

Navigation capabilities:
Learn to navigate by looking at a map

PRM-RL: Long-range Robotic Navigation Tasks
by Combining Reinforcement Learning and Sampling-based Planning
ICRA 2018, https://arxiv.org/abs/1710.03937, *Best paper in Service Robotics*

Aleksandra Faust, Oscar Ramirez, Marek Fiser, Kenneth Oslund, Anthony Francis, James Davidson
Lydia Tapia

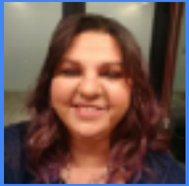
Aleksandra Faust


Oscar Ramirez


Marek Fiser


Ken Oslund


Anthony Francis


James Davidson


Lydia Tapia

How to navigate to a location on a map?

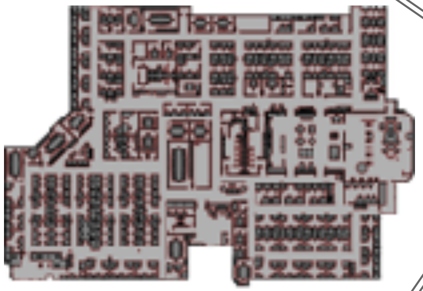

Transferable to new environments
Easy sim2real

Handles:
Sensors to controls, dynamics, noise
Obstacle avoidance

Lacks context needed for navigation.

Related work: Sampling-based planners

- Long-distance collision-free navigation
    - Approximate all possible robot motions
    - Sample and connect robot poses
    - Connect them with small, feasible, motion transitions
- Checking validity of pose transitions is expensive
- Sampling based planners that create reusable roadmaps
    - Often consider geometry only

[Hauser et al., '06]

[Hsu et al., '02]

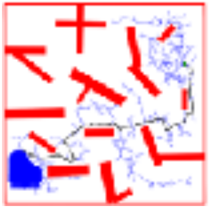

[LaValle & Kuffner, '01]

[LaValle & Kuffner, '01]

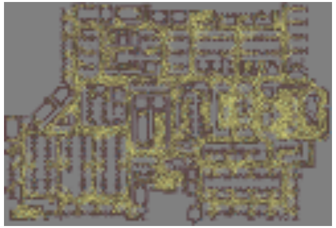

[Kavraki et al. '96]

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building

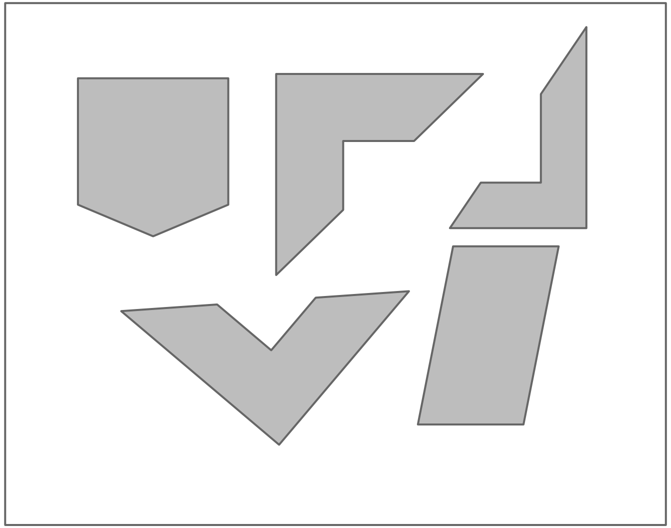

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building
  - Sample configuration space

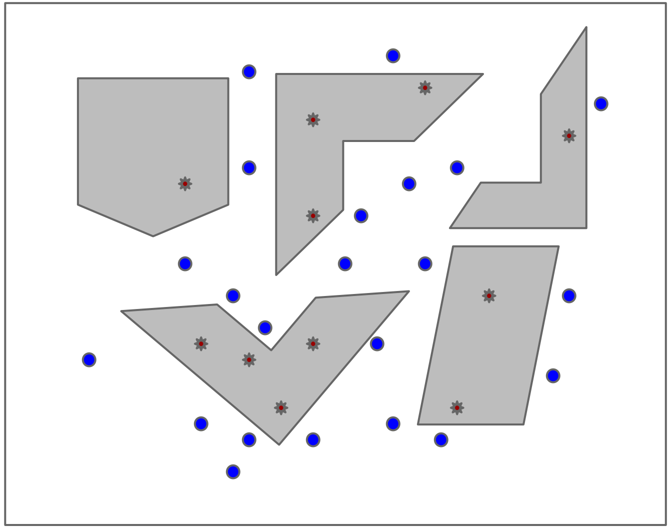

Google

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. ‘96]

- Building
  - Sample configuration space
  - Reject in-collision samples

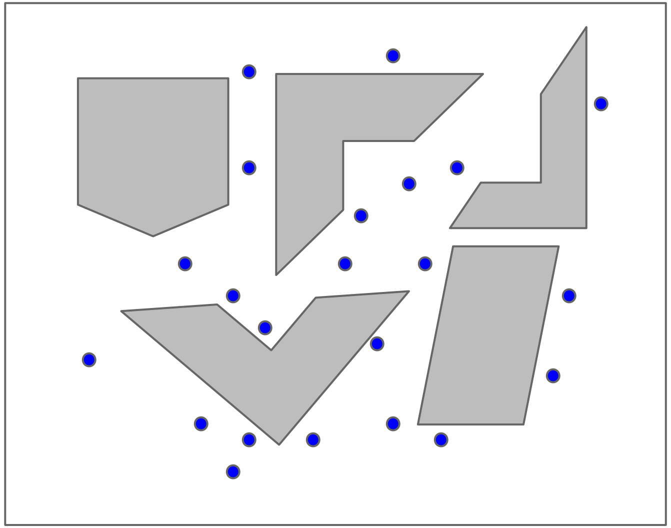

Google

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path

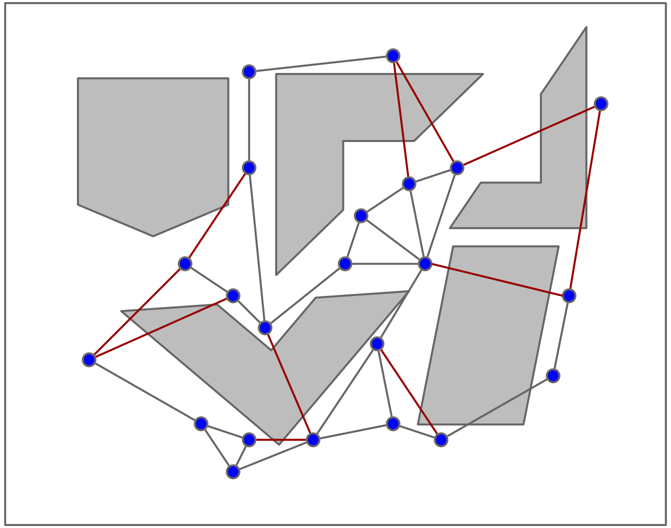

Google

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path

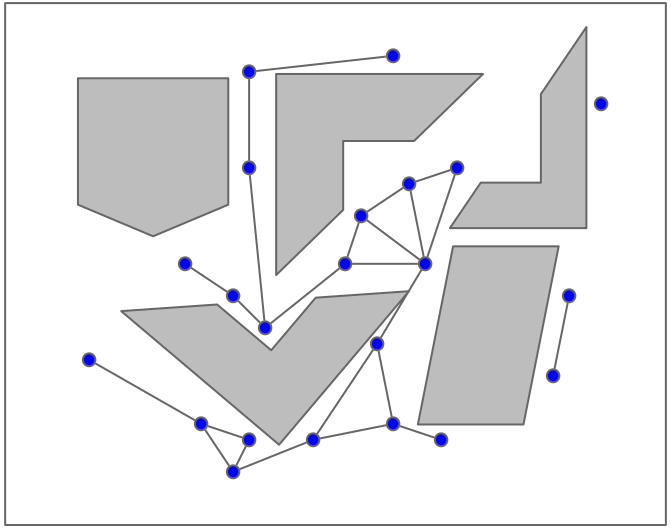

Google

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path
- Querying

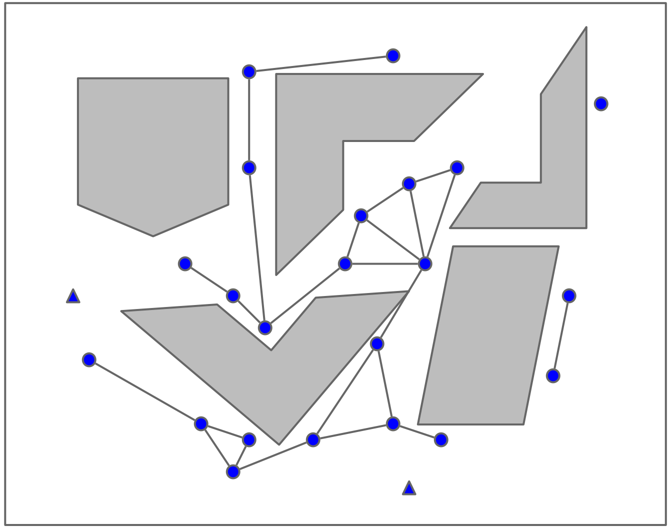

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. ‘96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path
- Querying
  - Add start and goal to the roadmap

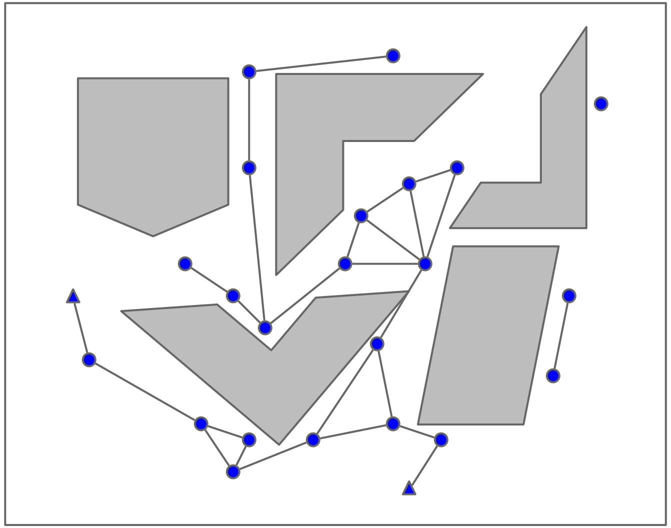

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. '96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path
- Querying
  - Add start and goal to the roadmap
  - Find the shortest path in the graph

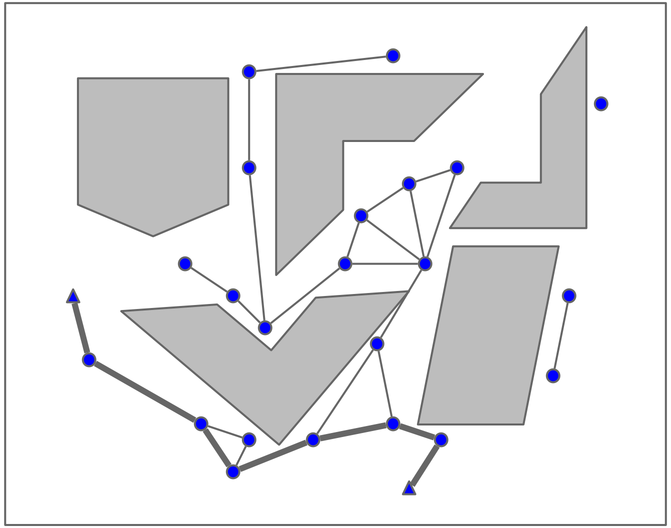

Google

Related Work: Probabilistic Roadmaps (PRMs)

[Kavraki et al. ‘96]

- Building
  - Sample configuration space
  - Reject in-collision samples
  - Connect samples only if a local planner finds a collision-free path
- Querying
  - Add start and goal to the roadmap
  - Find the shortest path in the graph
- Path following
  - Path guided artificial potential fields [Chiang et al. 2015]
  - Reinforcement learning [Faust et al. 2017]

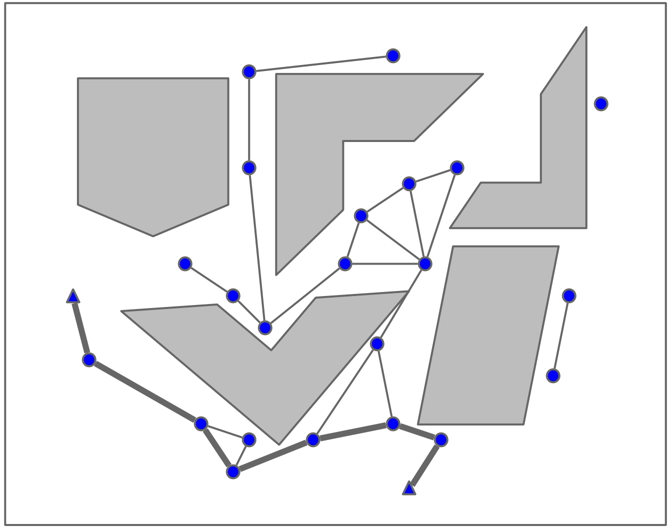

Google

PRM-RL Algorithm

Trained point to point agent - basic navigation behavior.

[Faust et al., ICRA 2018]

Google

PRM-RL Algorithm

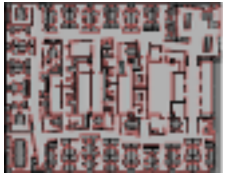

Trained point to point agent - basic navigation behavior.

One time setup

[Faust et al., ICRA 2018]

PRM-RL Algorithm

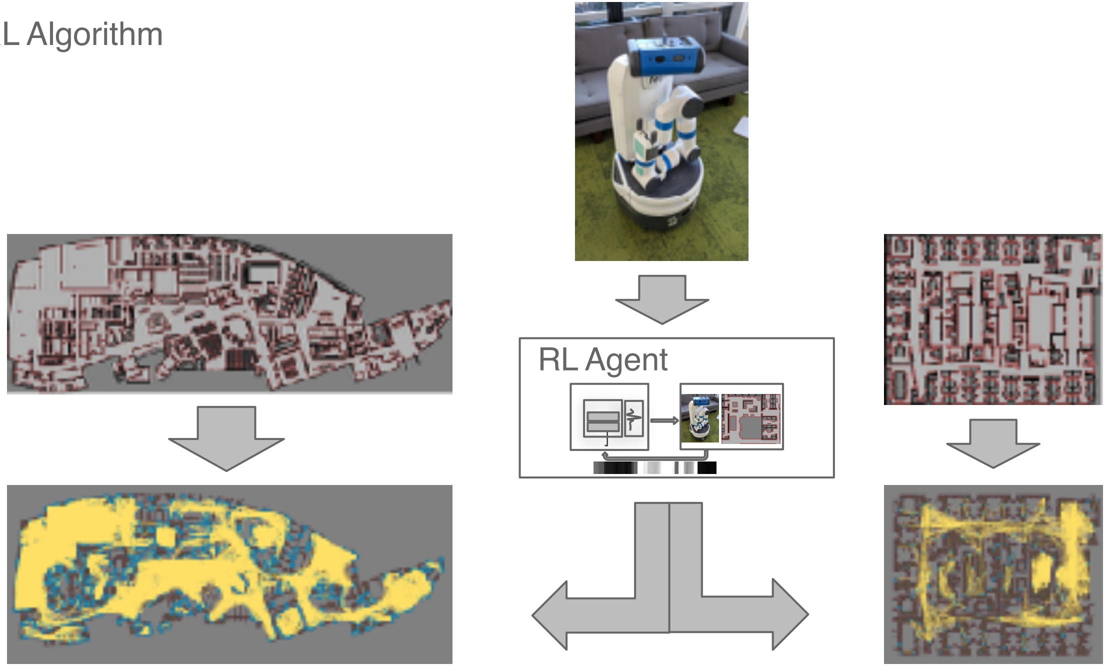

Trained point to point agent - basic navigation behavior.

One time setup

Add an edge only if RL agent can consistently navigate between two nodes

[Faust et al., ICRA 2018]

PRM-RL Algorithm

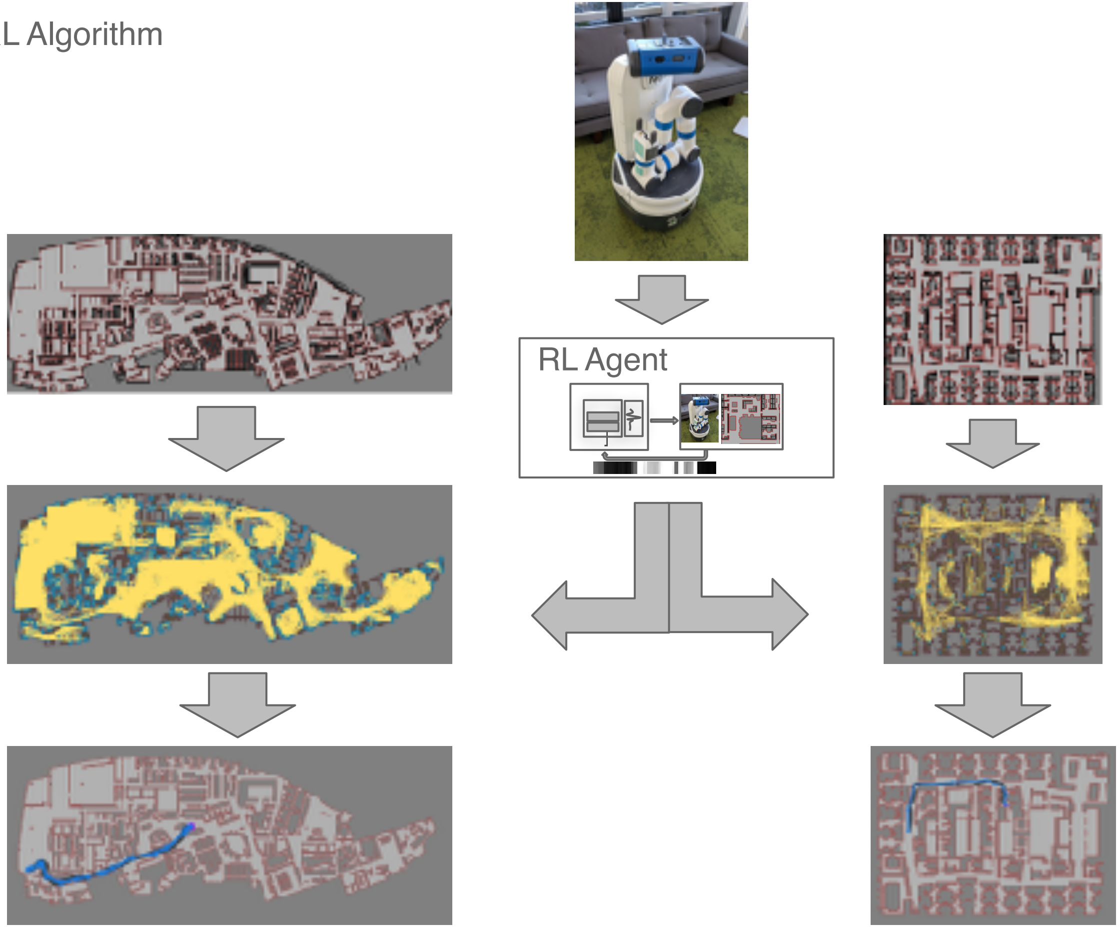

Trained point to point agent - basic navigation behavior.

One time setup

Add an edge only if RL agent can consistently navigate between two nodes

Execute long trajectories

[Faust et al., ICRA 2018]

PRM-RL: Indoor Navigation Building PRMs

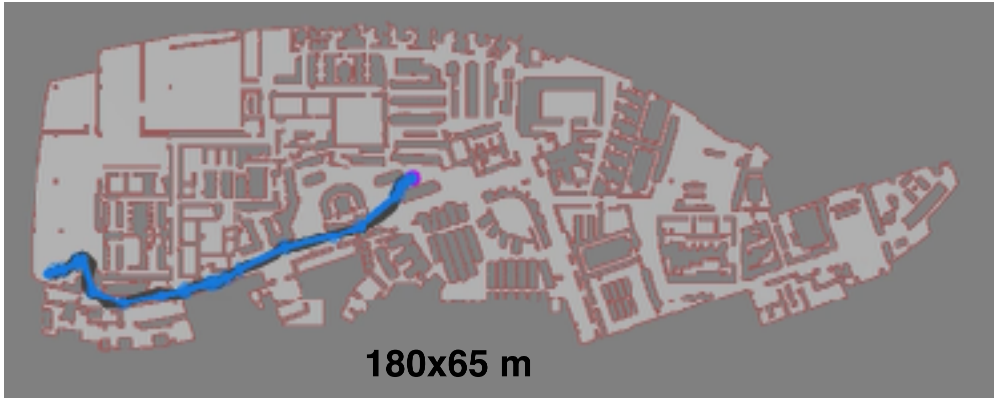

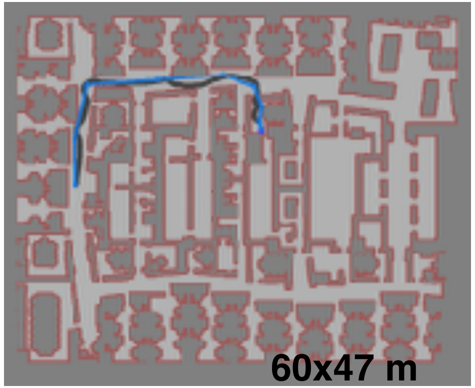

60x larger than the training

20 trials with 85% confidence

[Faust et al., ICRA 2018]

Google

PRM-RL: Indoor Navigation Building PRMs

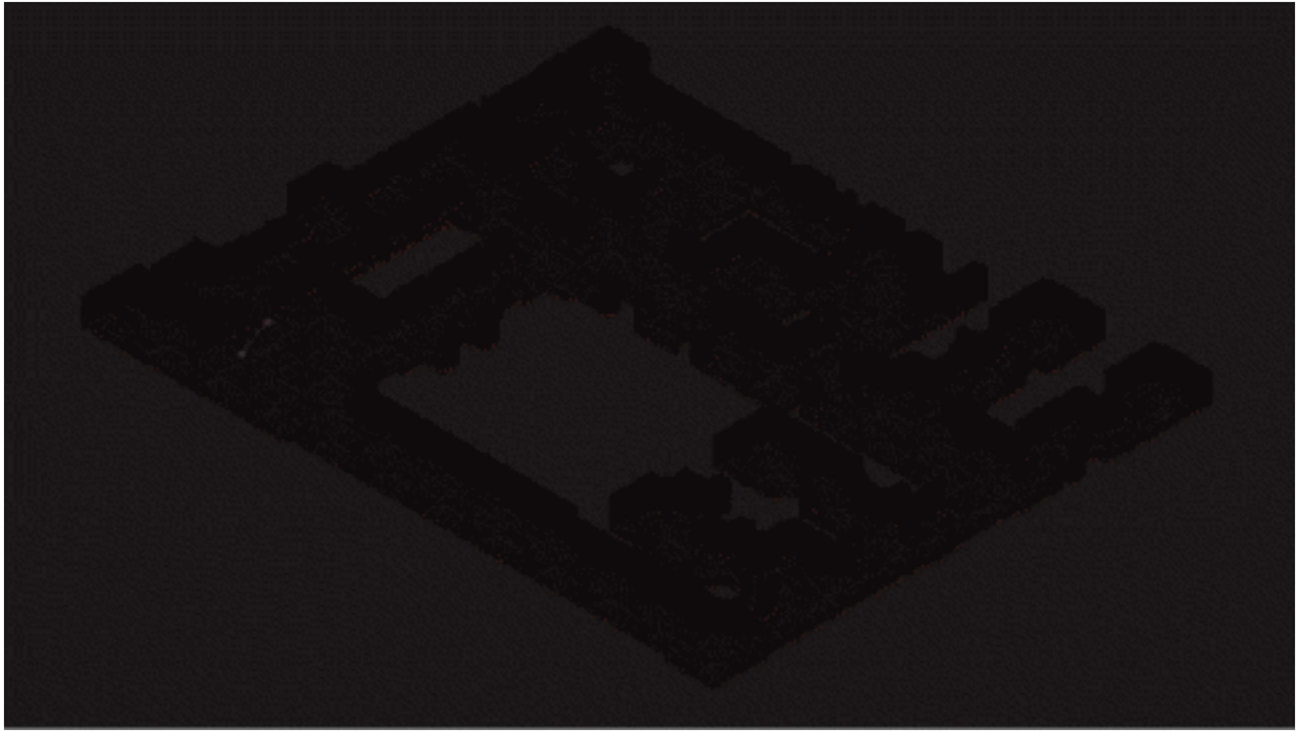

2 hours to build

Largest roadmap:

1700 nodes
60 000 edges
23 million collision checks

One-time set-up

[Faust et al., ICRA 2018]

Google

PRM-RL: Results for Indoor Navigation

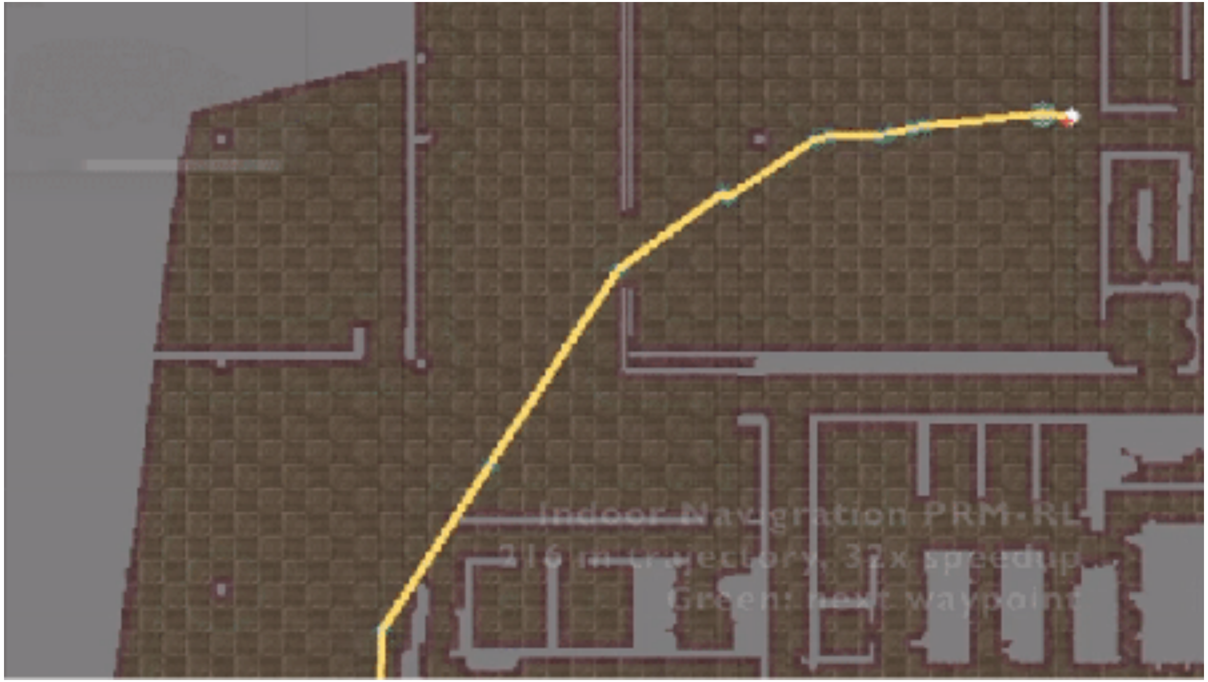

Longest trajectory 215 meters

45 waypoints

[Faust et al., ICRA 2018]

Google

PRM-RL Experimental Results

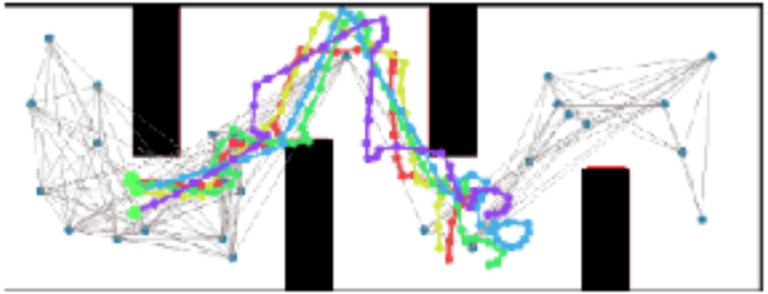

Four noisy trials

All successful, because the map is tuned to the robot's abilities

[Faust et al., ICRA 2018]

Google

How to navigate to a location on a map?

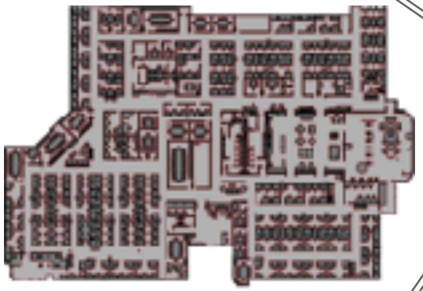

Navigates over long distances.

Requires one time set-up.

Transferable to new environments
Easy sim2real

Handles:
Sensors to controls, dynamics, noise
Obstacle avoidance

Navigation capabilities: Following Directions

Following Natural Language Navigation Instructions with Deep Reinforcement Learning
Under submission
Aleksandra Faust, Chase Kew, Dilek Hakkani-Tur, Marek Fiser, Pararth Shah

FollowNet: Towards Robot Navigation by Following Natural Language Directions with Deep Reinforcement Learning
MLPC @ ICRA 2018, https://arxiv.org/abs/1805.06150
Pararth Shah, Marek Fiser, Aleksandra Faust, Chase Kew, Dilek Hakkani-Tur


Pararth Shah


Marek Fiser

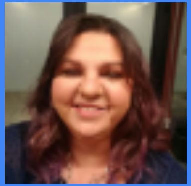
Aleksandra Faust


J. Chase Kew


Dilek Hakkani-Tur

How to follow directions?

One time building / robot setup.

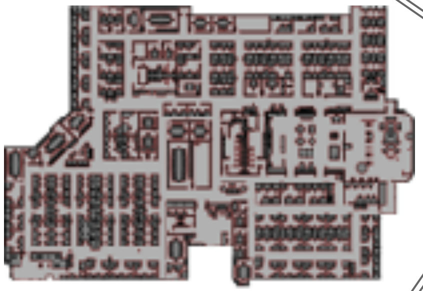

Go down the hallway and take the second right.

Transferable to new environments
Easy sim2real

Hearing

Sensor to controls, dynamics, noise
Obstacle avoidance

Google

Following instructions

Go down the hallway and take the second right.

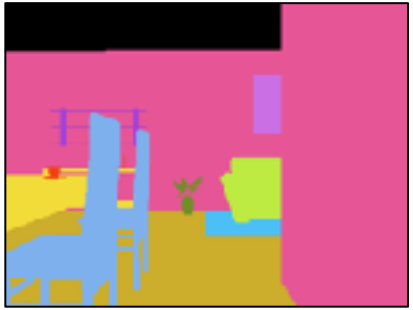

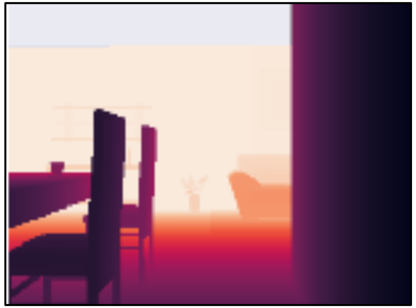

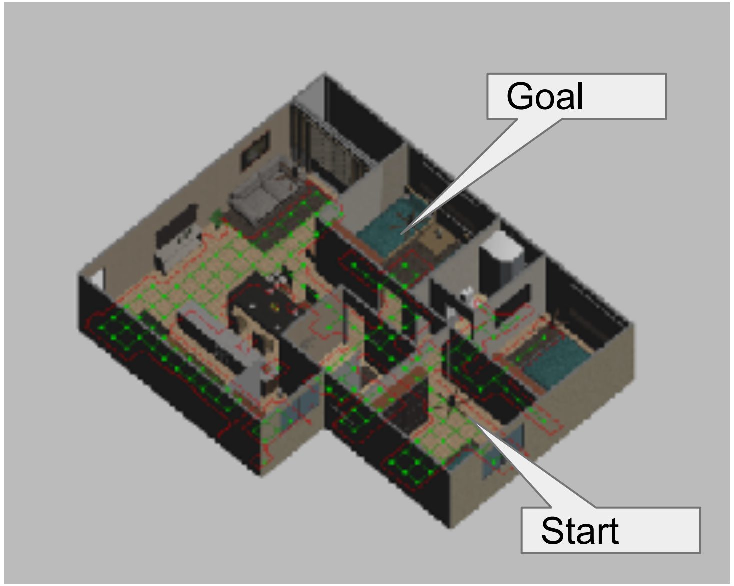

[Shah et al., 2018]

Following instructions

Go down the hallway and take the second right.

Dataset:
150 instructions
2 buildings

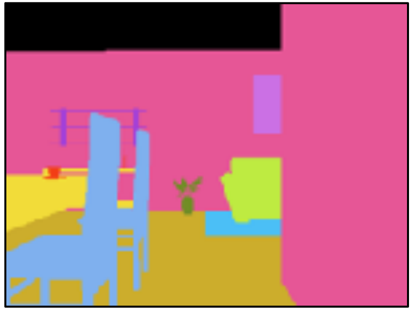

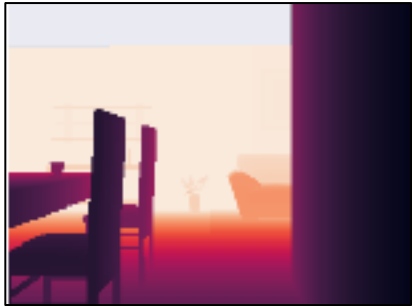

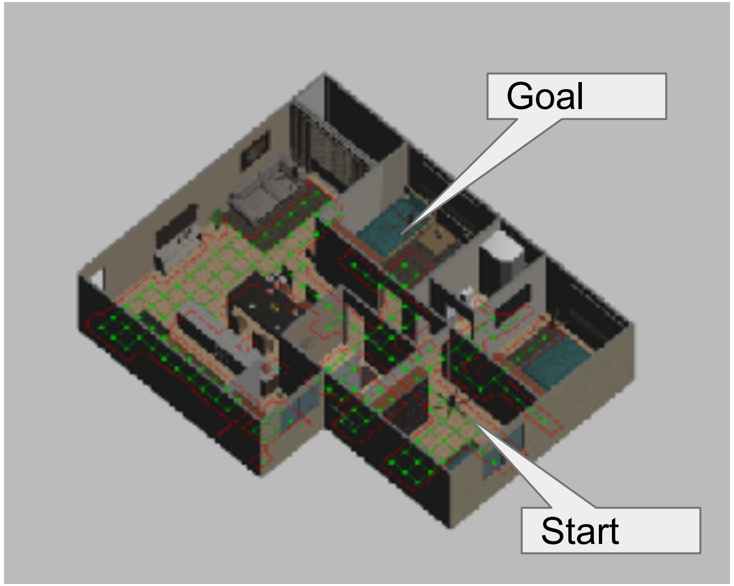

[Shah et al., 2018]

Following instructions: Related work

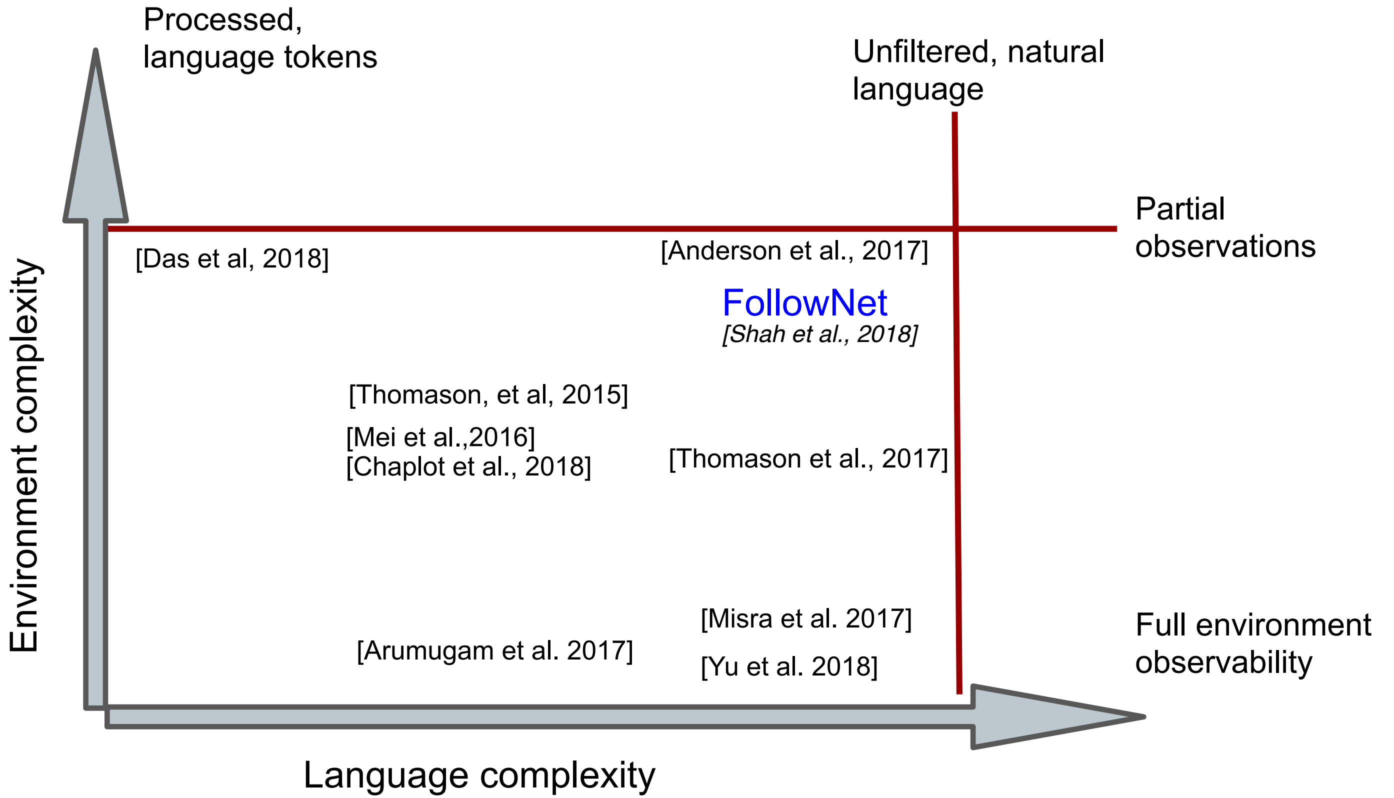

Google

Following instructions: RL Training

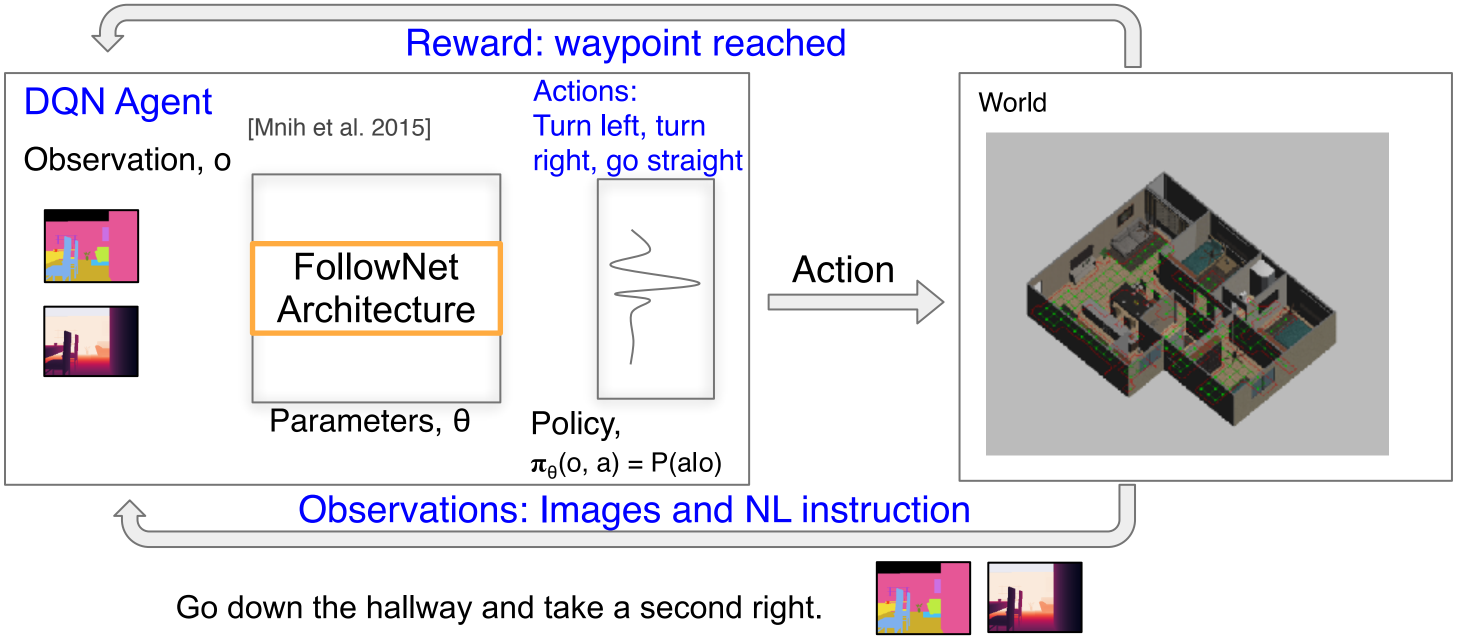

[Shah et al., 2018]

FollowNet: Architecture

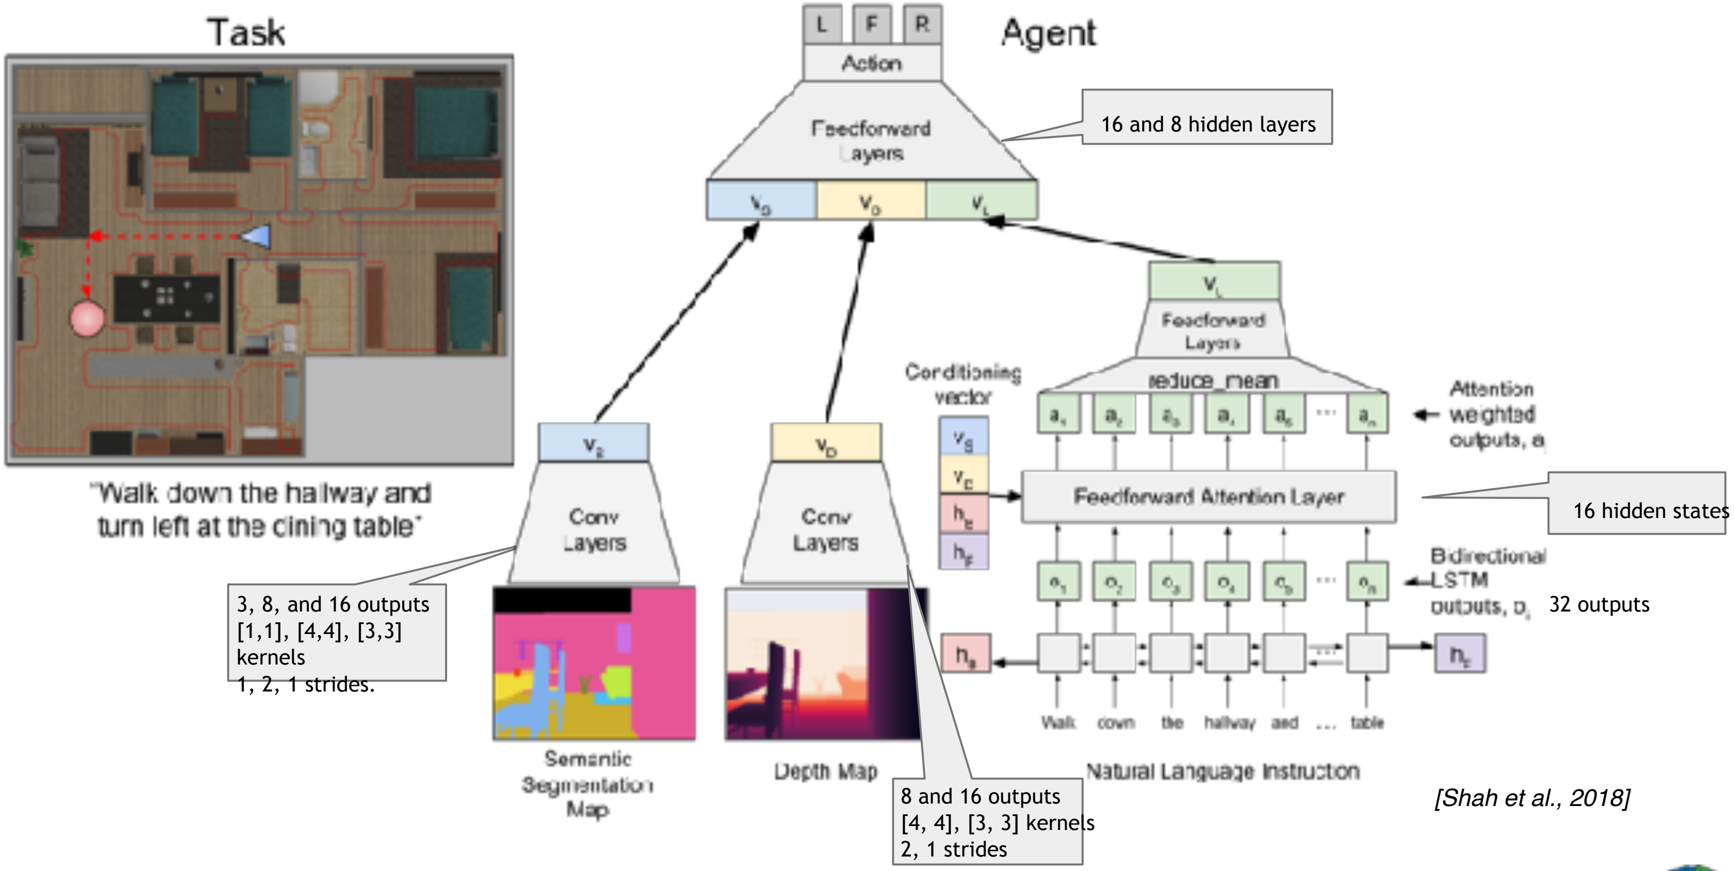

[Shah et al., 2018]

Google

FollowNet: Results

Baseline: model without attention

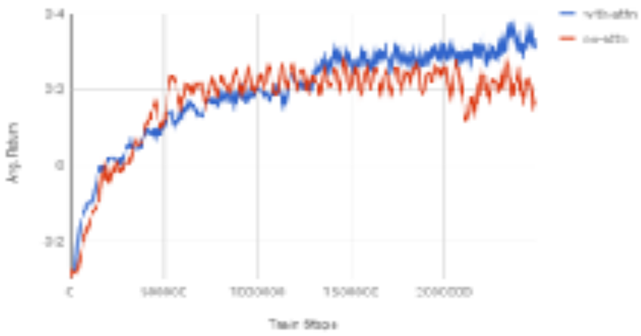

Better learning curve than the baseline

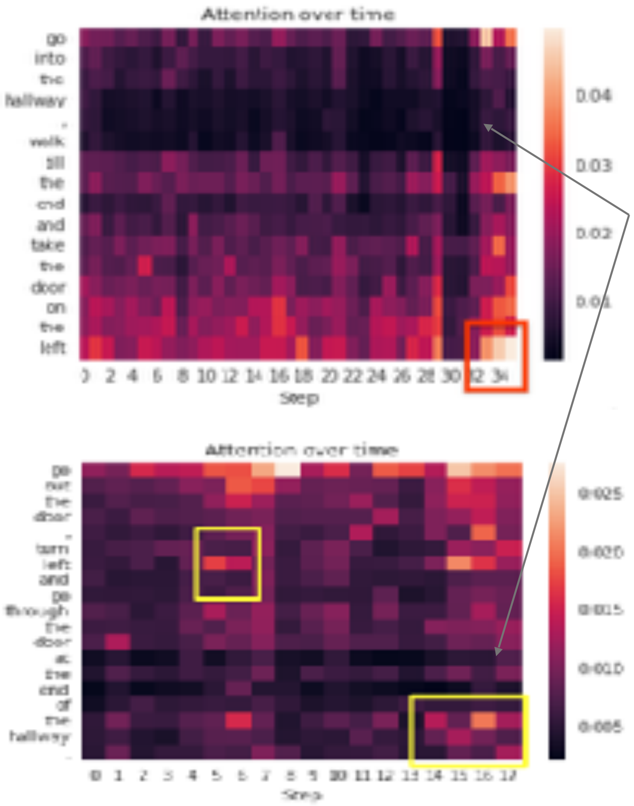

Learns what to ignore

Learns what is important

Attention over steps

[Faust et al., under submission]

Instruction complexity

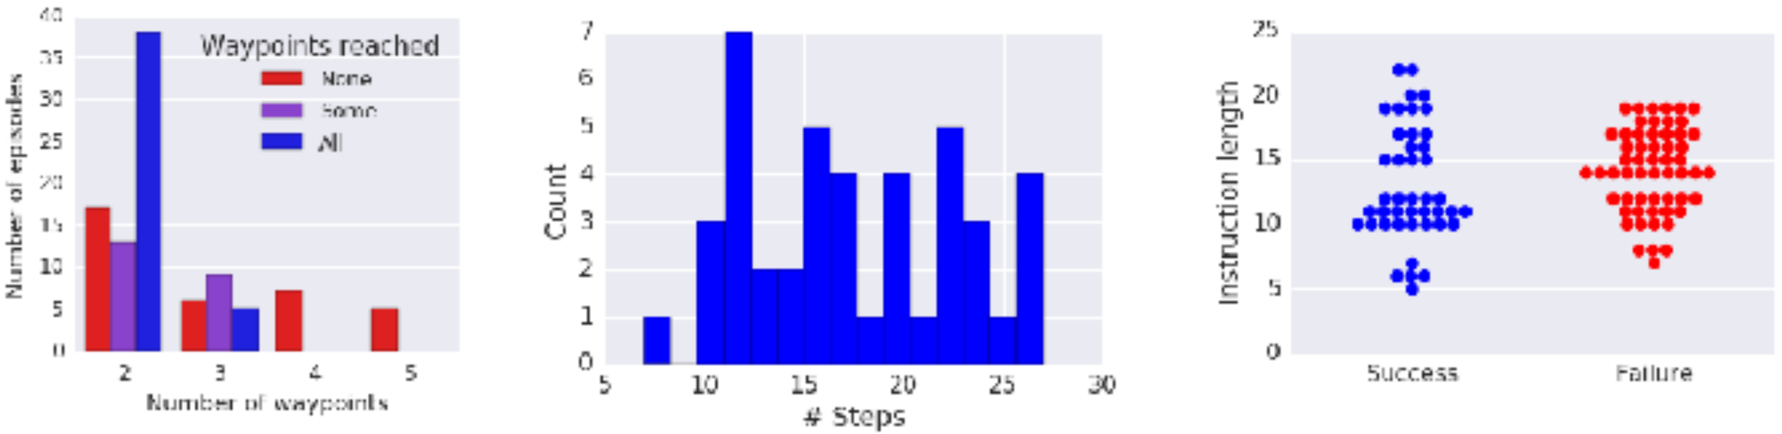

Number of waypoints measures the instruction complexity, not number of words or path length.

[Faust et al., under submission]

Google

FollowNet on New Instructions

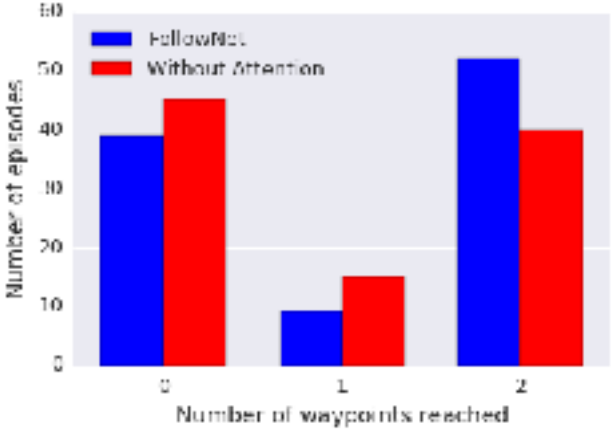

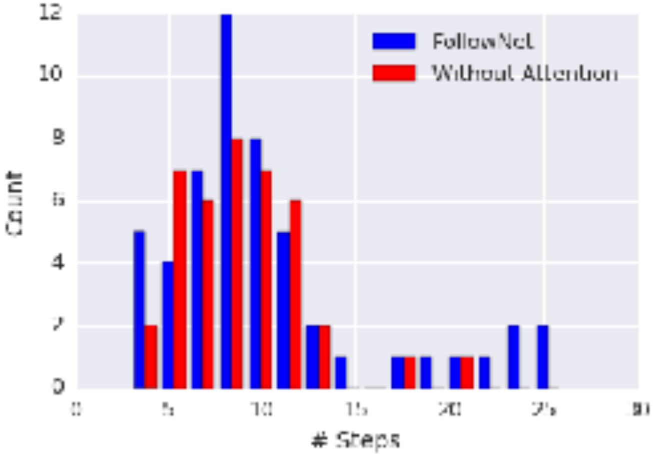

52% success on new instructions
67% at least partial success

30% increase over the baseline

Google

FollowNet
performance per word

Words with spatial semantics are more likely to be successful

Landmarks, less so

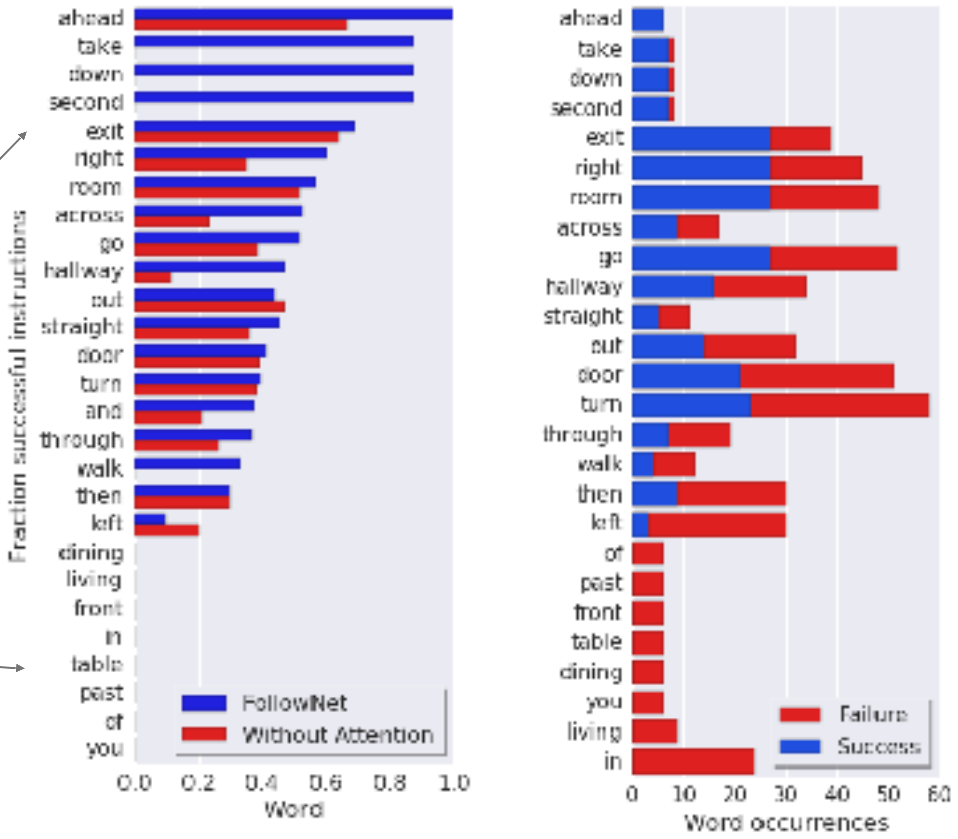

[Faust et al., under submission]

Google

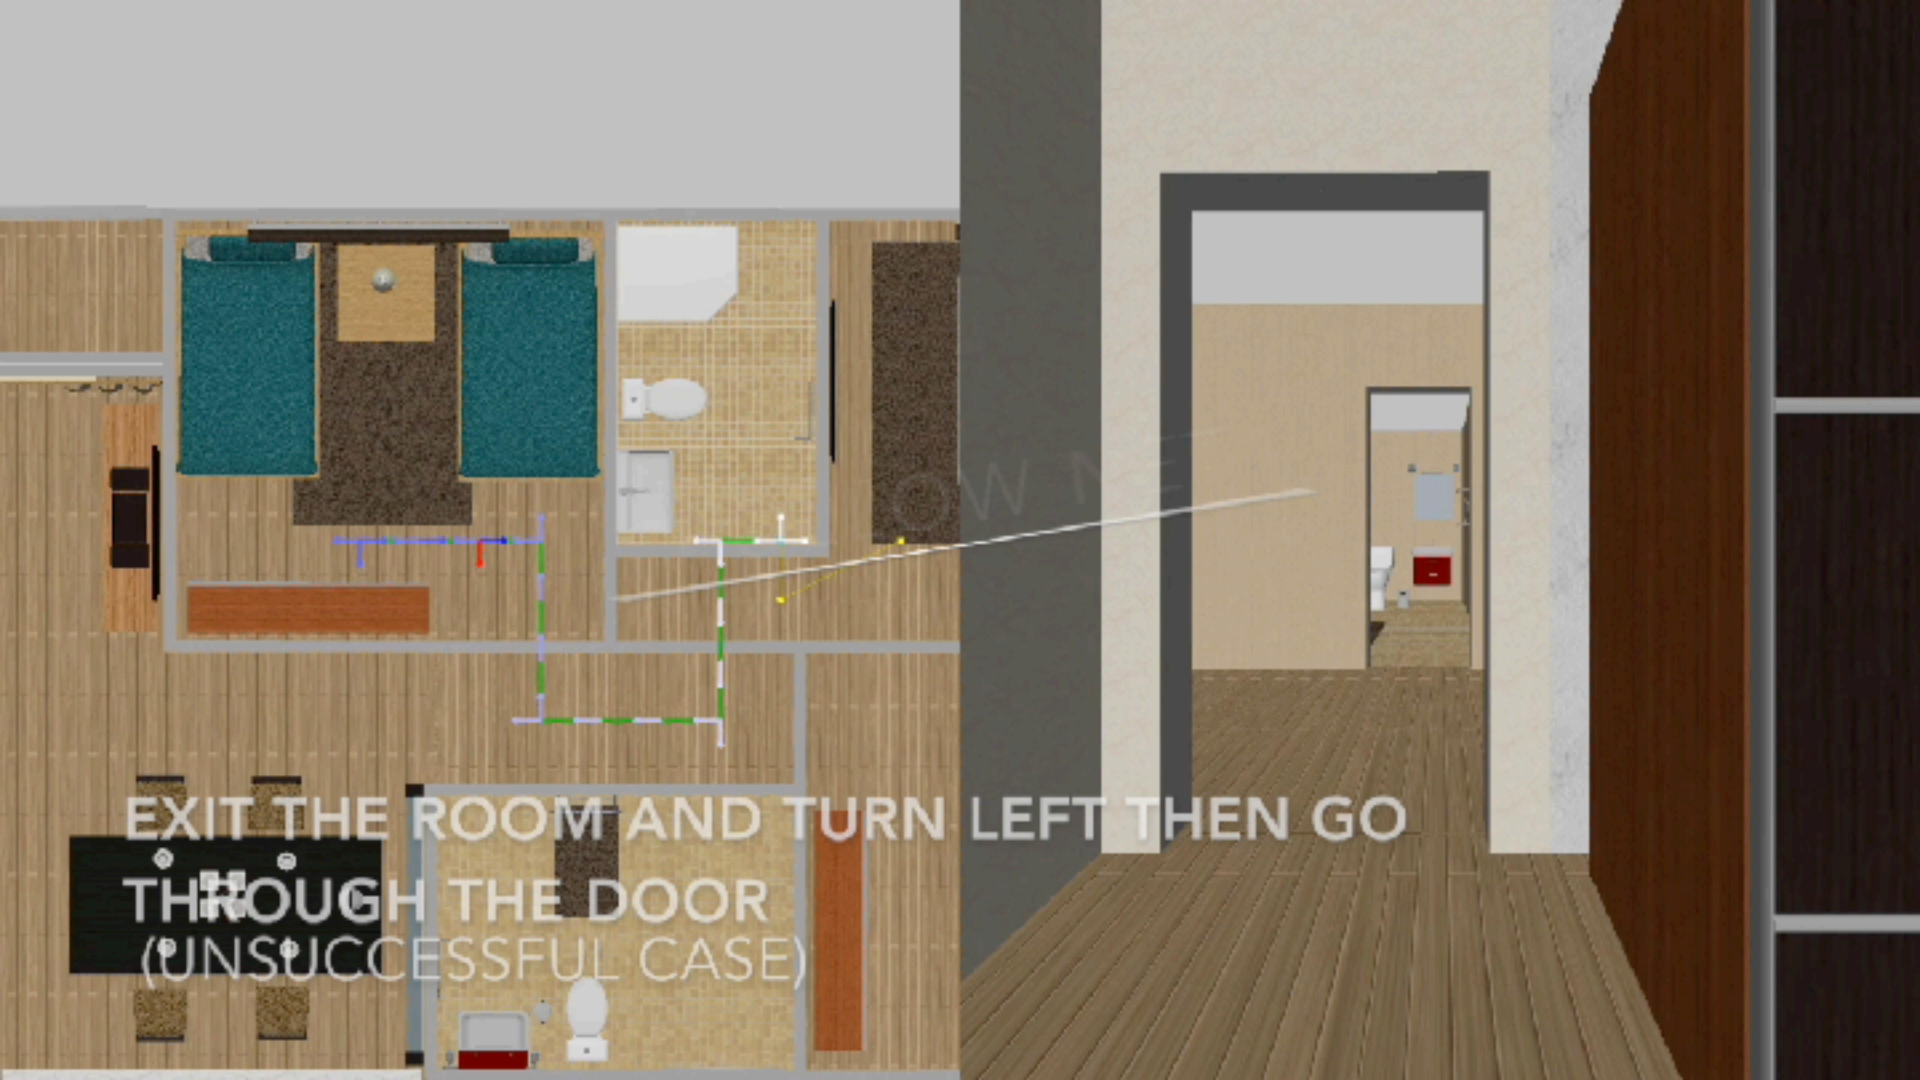
EXIT THE ROOM AND TURN LEFT THEN GO
THROUGH THE DOOR
(UNSUCCESSFUL CASE)

How to follow directions?

One time building / robot setup.

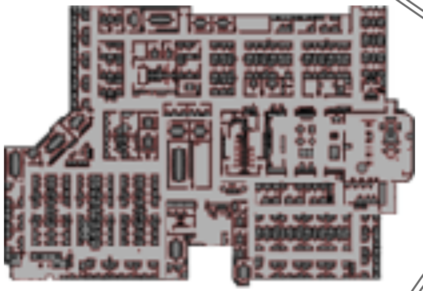

Go down the hallway and take the second right.

Transferable to new environments
Easy sim2real

Sensor to controls, dynamics, noise
Obstacle avoidance

Does not require building set-up.
Promising results.

From robot abilities to navigation capabilities

End to end complex navigation tasks

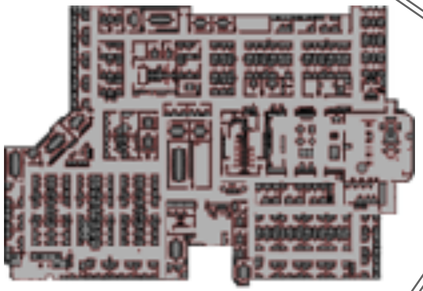

Go down the hallway and take the second right.

End to end basic navigation behaviors behaviors

Hearing

Navigation

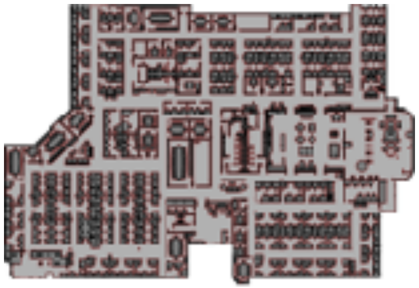

Go down the
hallway and take
the second right.

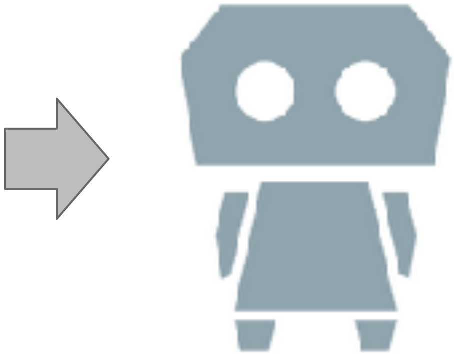

Perception

Slam

Planning

Controls

Navigation

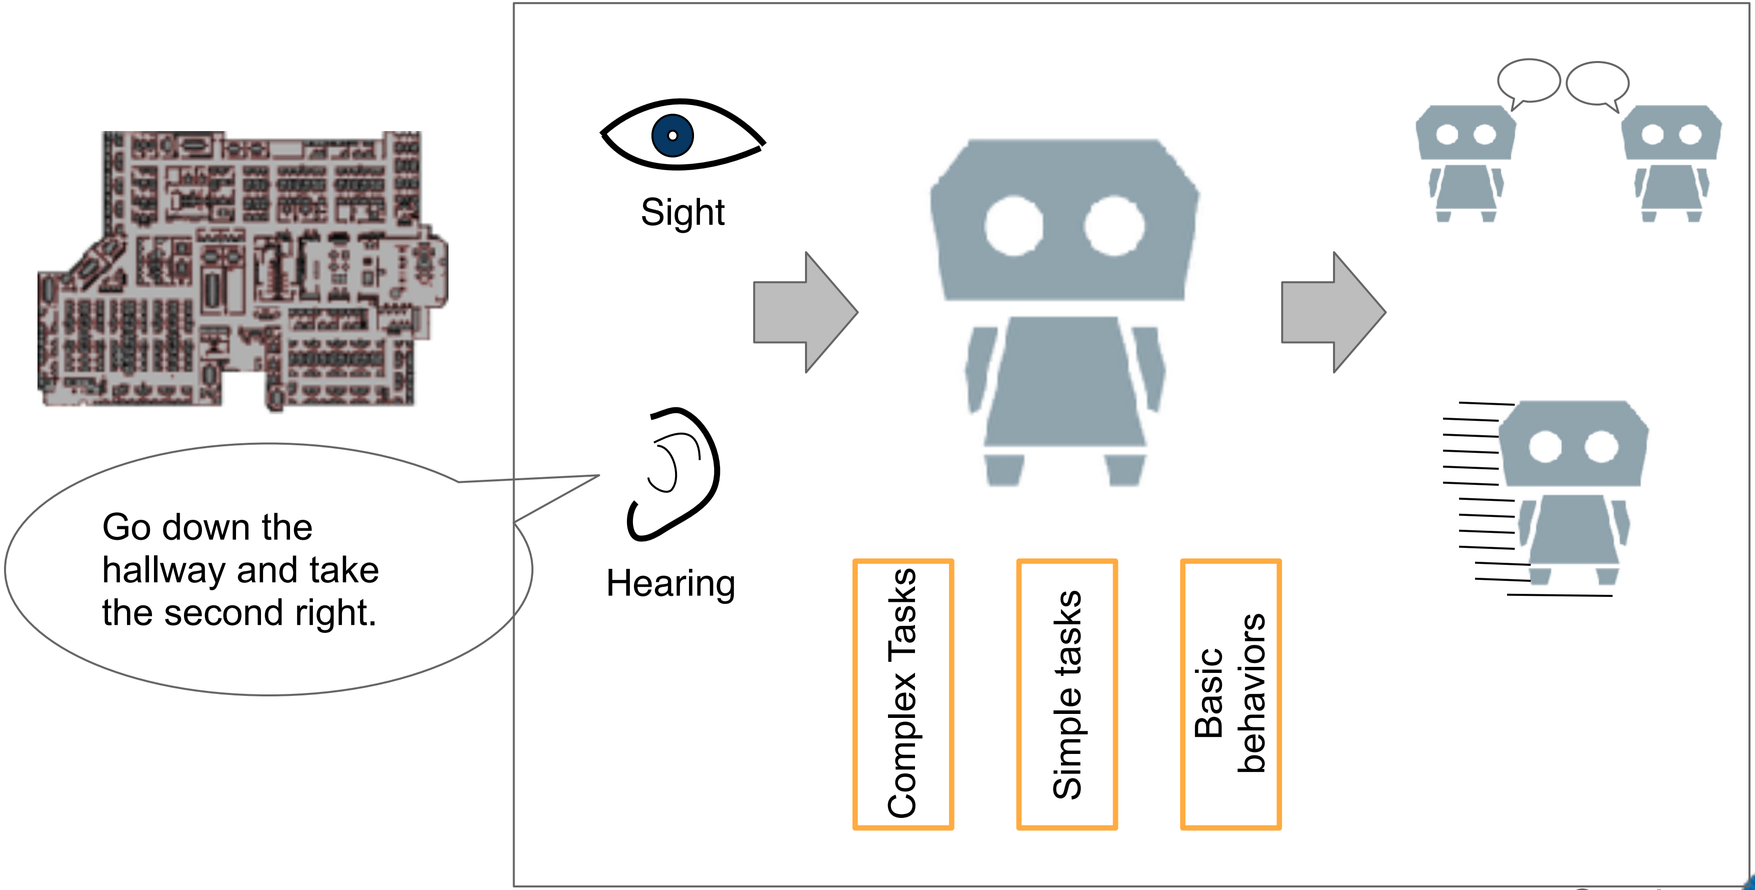
Go down the hallway and take the second right.
Sight
Hearing
Complex Tasks
Simple tasks
Basic behaviors

Navigation

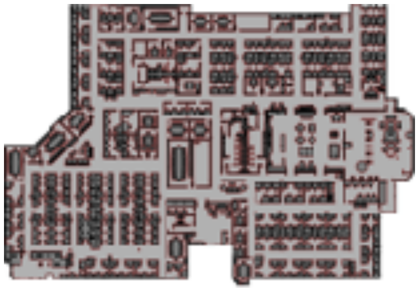

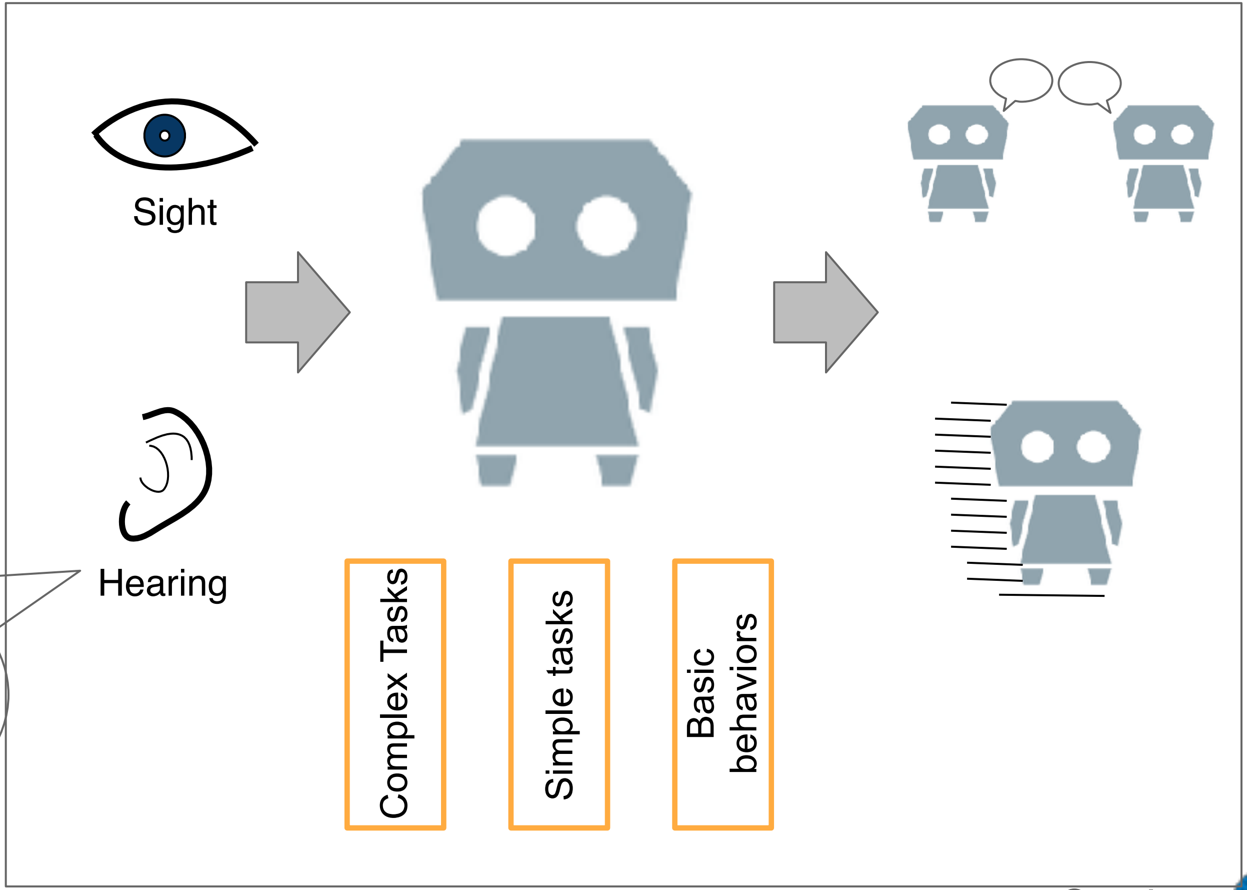

Thank you!
QUESTIONS
?

Thank you


Anthony Francis


Marek Fiser


J. Chase Kew


Lydia Tapia


Ken Oslund


Oscar Ramirez


Hao-Tien Chiang


Dilek Hakkani-Tur


Pararth Shah


James Davidson

UNM

Google