# Human-Interactive Mobile Robots: from Learning to Deployment

Xuesu Xiao, Ph.D.

Assistant Professor

Computer Science, George Mason University

# RobotiXX


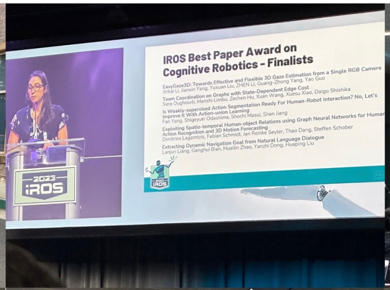
IROS Best Paper Award on
Cognitive Robotics – Finalists
2023
IROS

Toward Human-Like Social Robot Navigation:
A Large-Scale, Multi-Modal, Social Human
Navigation Dataset








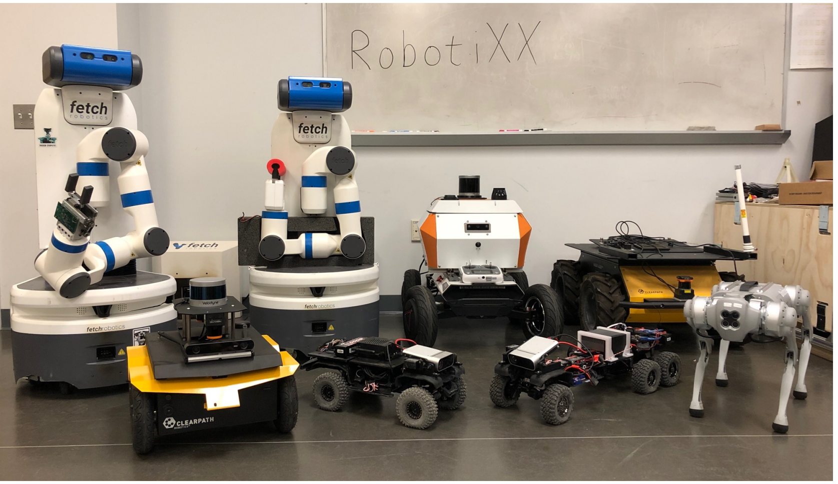
RobotiXX
fetch
fetch
CLEARPATH

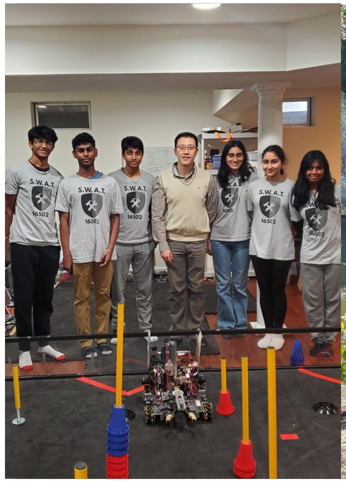
S.W.A.T.
16502

**Research Goal**: Develop highly capable and intelligent mobile robots that are robustly deployable in the real world with minimal human supervision

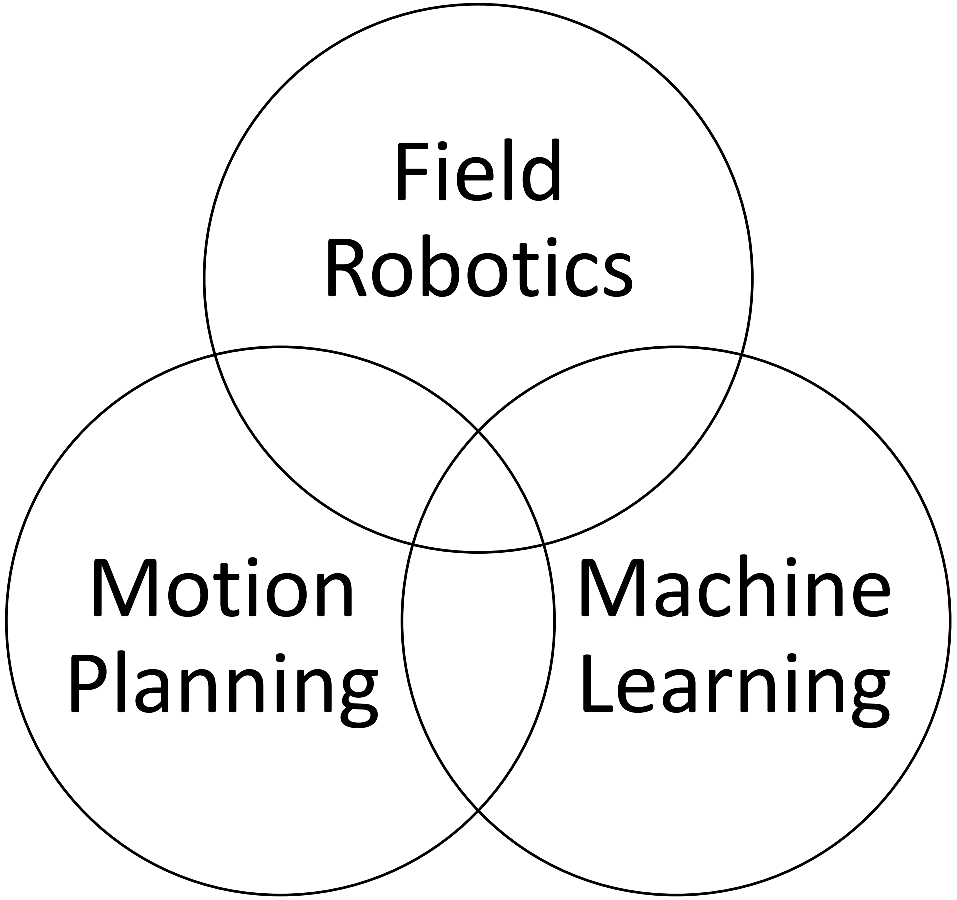

# Disaster Robotics

Modular Snake (Mexico City Earthquake)

EMILY (Greece Refugee Crisis)

PackBot (Fukushima Daiichi)

[***Xiao*** *et al., ICRA15*]

- Overhead Cameras
- Locomotive Reduction

[***Xiao*** *et al., IROS17, Dufek,* ***Xiao****, Murphy, SSRR17*]

- UAV-USV Team
- Visual Pose Stabilization
- Visual Navigation

[*Dufek,* ***Xiao****, Murphy, THMS21,*
***Xiao*** *et al., RA-L20,*
***Xiao*** *et al., FSR19*
***Xiao*** *et al., SSRR19a,*
***Xiao*** *et al., SSRR19b,*
***Xiao*** *et al., IROS18,*
***Xiao*** *et al., SSRR18*
*(Best Paper Finalist),*
***Xiao*** *et al., SSRR17*]

- Viewpoint Theory
- Risk-Awareness
- Tethered Flight

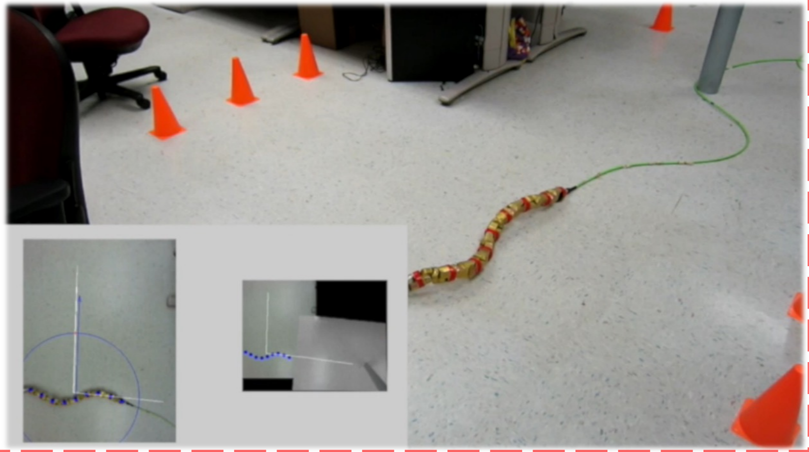

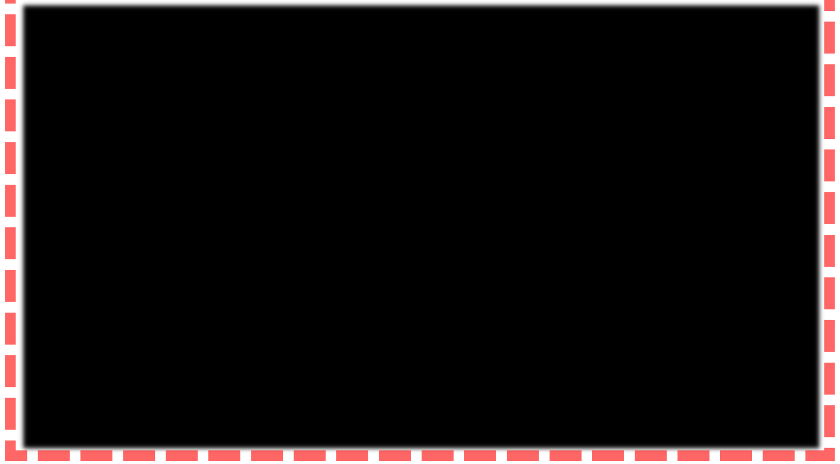

# Off-Road Mobility

High Speed

[***XX** et al., RA-L21,*
*K, S, A, R,* ***XX** et al., IROS22*]

Vertically Challenging Terrain

[*D, P, N,* ***XX**, ICRA24,*
*D, P,* ***XX**, under review*]

# Mobility in Highly-Constrained Environments

# The BARN Challenge and Datasets

ICRA 2022
Philadelphia
[***XX*** *et al. RAM22*]

ICRA 2023
London
[***XX*** *et al. RAM23*]

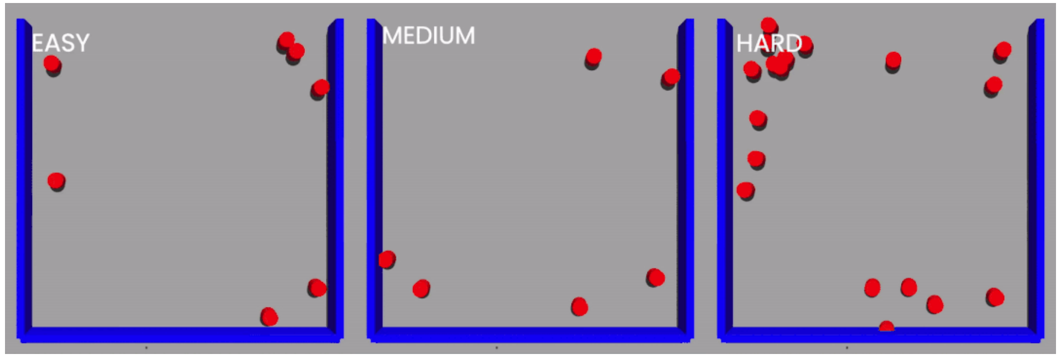


DynaBARN [N, J, H, X, L, ***XX***, *S, SSRR22*]

Benchmark Autonomous Robot Navigation (BARN)
[P, T, ***XX***, *S. SSRR20*]

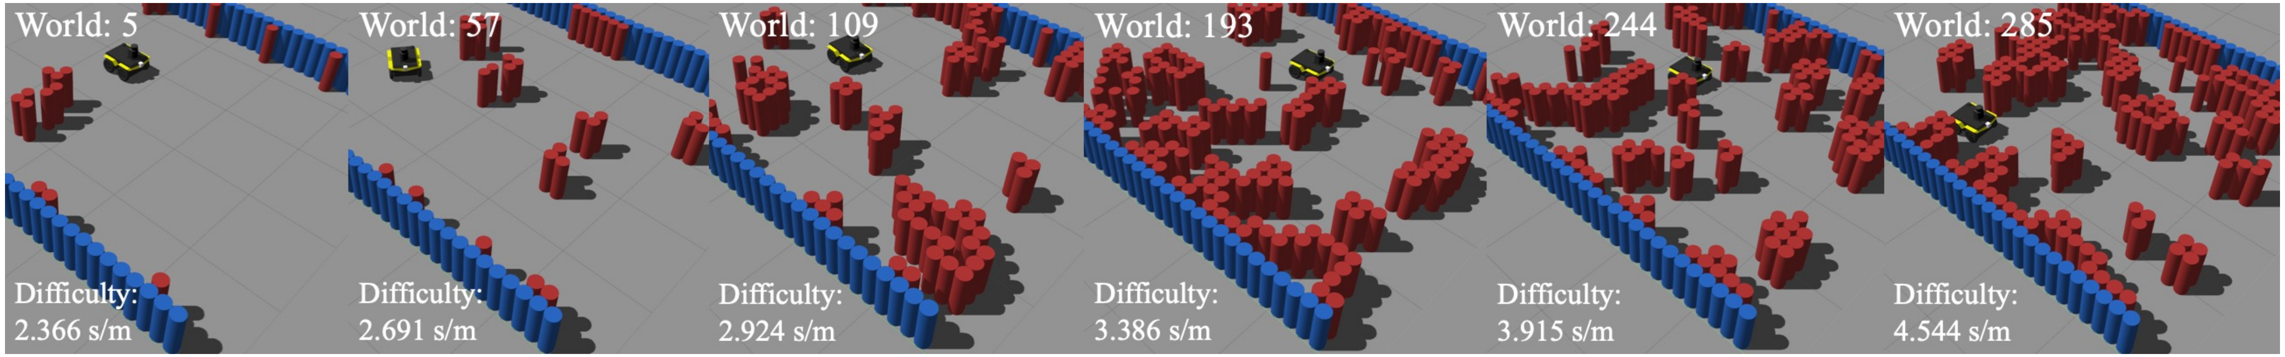

# Where are robots currently deployed?

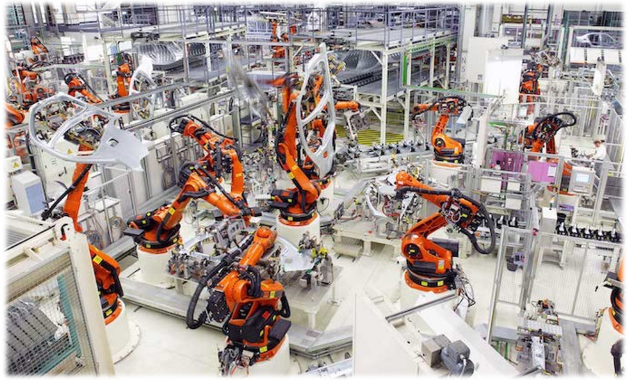

Manufacture (Kuka)

**Highly Controlled Workspace**

**Do not Learn**

Logistics (Amazon)

Home (iRobot)

**Preprogrammed Single Task**

**Do not Learn**


Entertainment (Intel)

Healthcare (Da Vinci)

**Fully Piloted by Skilled Humans**

**Not Autonomous**

Safety (Endeavor)

# Where do we want robots to be deployed?

[Times Square, New York]

[George Mason University]

# What has changed?

- Humans!

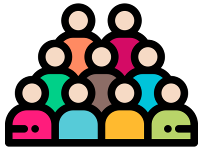

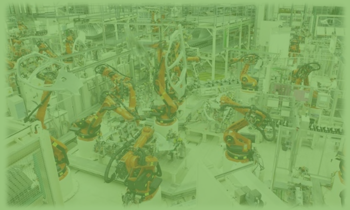

Manufacture (Kuka)

Home (iRobot)

**No need to Deal with Humans**

Logistics (Amazon)

Entertainment (Intel)

Healthcare (DaVinci)

**Utilize Expert Human Help**

Safety (Endeavor)

**??!!**

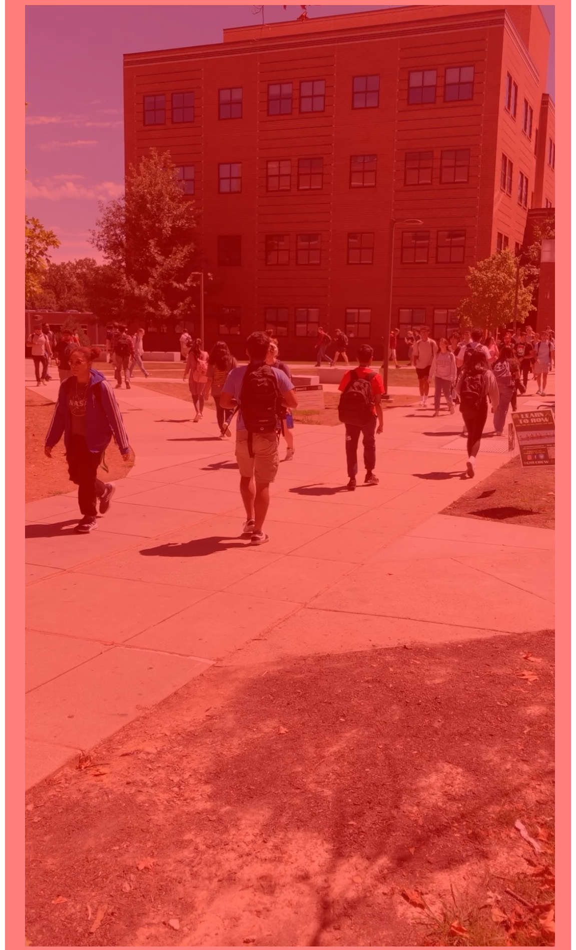

# Human-Interactive Mobile Robots: from Learning to Deployment

- Humans pose both challenges and opportunities:
  - Challenges: Diverse and uncertain human-robot interactions in the wild.
  - Opportunities: A wealth of diverse (non-expert) knowledge.

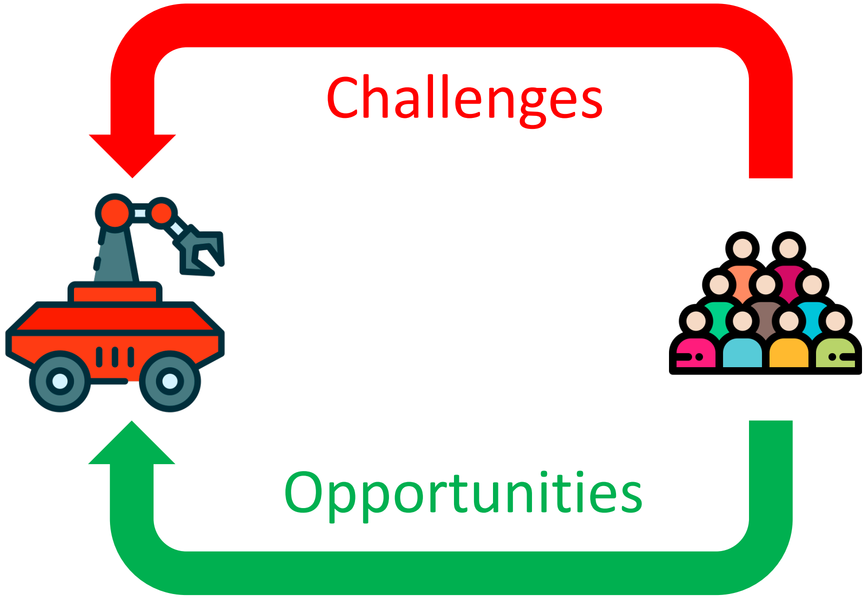

# Human-Interactive Mobile Robots: from Learning to Deployment

- This talk: Human-interactive mobile robots that efficiently learn from and harmoniously deploy among humans:
  - Adaptive Planner Parameter Learning (APPL) to utilize the opportunities from easily available non-expert human interactions.
  - Datasets, protocols, principles, guidelines, and learning methods to address social robot navigation challenges.

# Adaptive Planner Parameter Learning (APPL)

**Motivation:**

Deploying an autonomous navigation system in a new environment is not as straightforward as it may seem.

During an existing deployment, autonomous mobile robots will keep repeating the same mistake until **a roboticist** reprogram the robot.

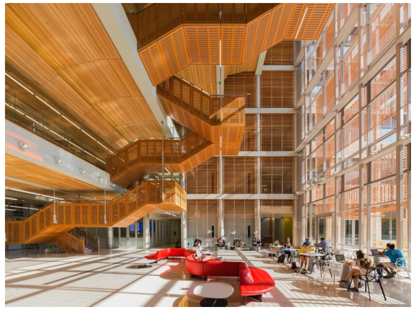

# Adaptive Planner Parameter Learning (APPL)

**Inspiration:** (Non-expert) Humans can do this effortlessly.

# Adaptive Planner Parameter Learning (APPL)

**Central Question:** Can we squeeze more robust performance out of our existing navigation systems using limited human interaction and learning?

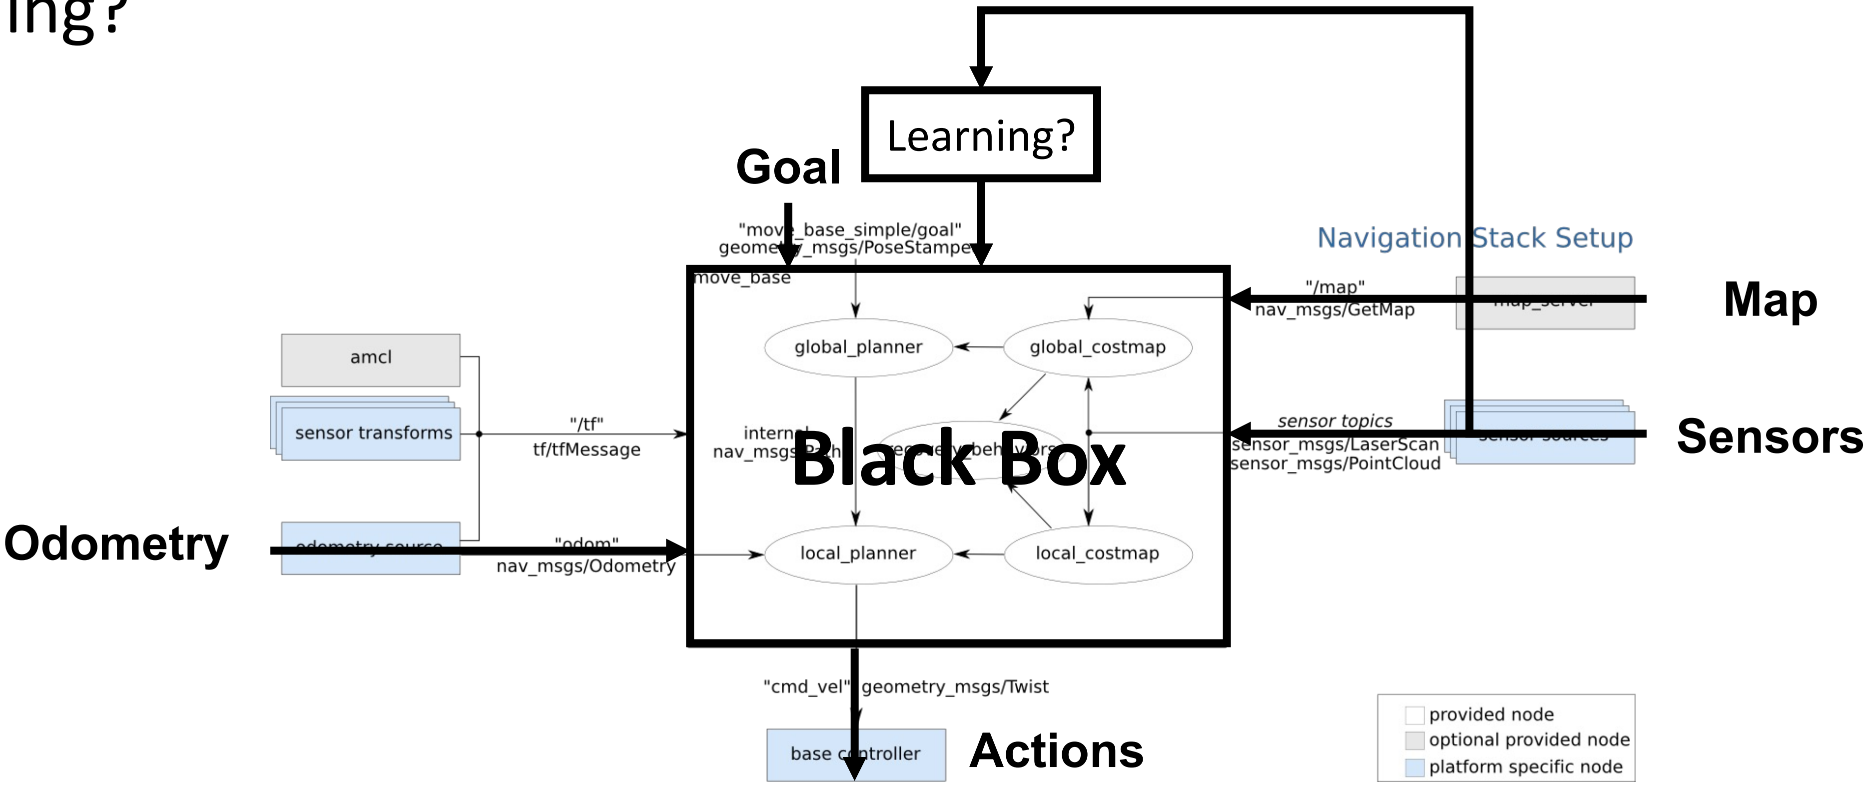


ROS **`move_base`** navigation stack

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

**Proposed:** Use behavioral cloning to “tune” any navigation stack.

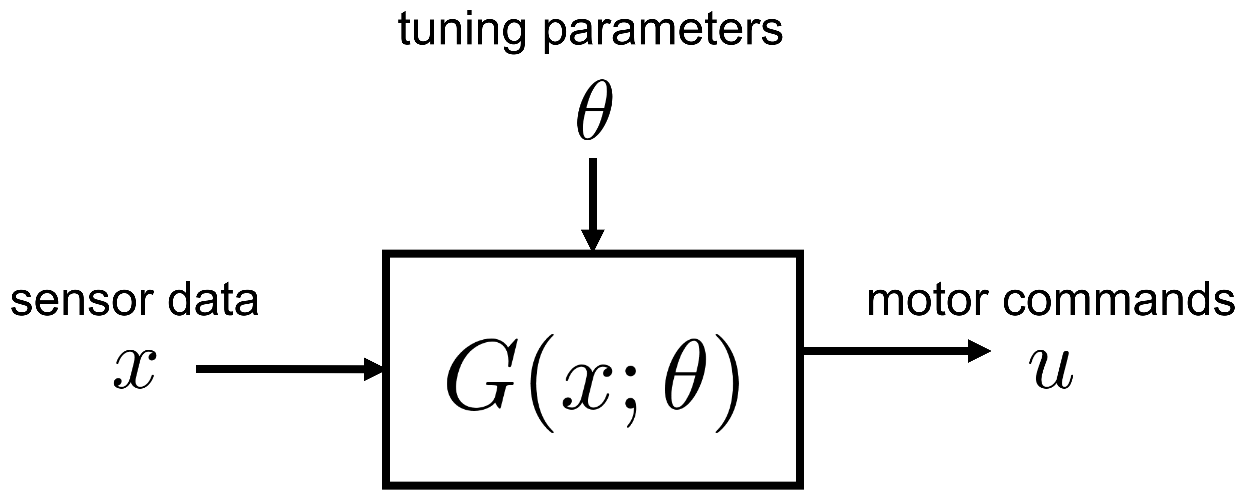


***max_vel_x***: 0.5
***min_vel_x***: 0.1
***max_vel_theta***: 1.57
***min_vel_theta***: -1.57
***vx_samples***: 6
***vtheta_samples***: 20
***occdist_scale***: 0.1
***pdist_scale***: 0.75
***gdist_scale***: 1.0
***Inflation_radius***: 0.30
......

**Behavioral Cloning**: Learn parameters from a demonstration using supervised learning.

$$\theta^* = \arg\min_\theta \sum_i \ell(G(x_i;\theta), u_i)$$

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

**Rough Procedure:**

1. Collect demonstration.

$$\mathcal{D} = \{(x_1, u_1), \ldots, (x_N, u_N)\}$$

2. Use black-box optimization to solve for planner parameters.

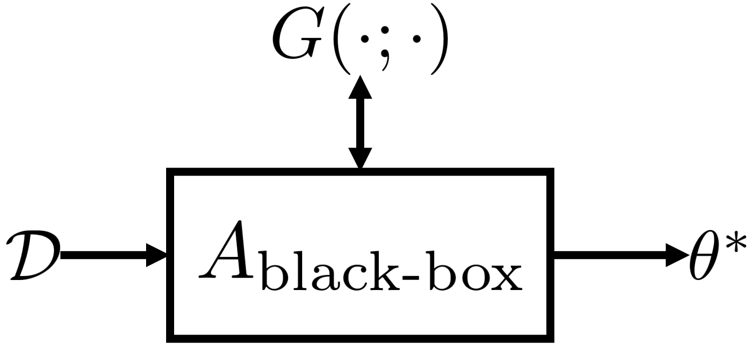

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

**Context Problem:** Humans exhibit qualitatively different navigation behaviors in qualitatively different environments.

$$\cancel{\theta^*}$$

$$\{\theta_1^*, \ldots, \theta_K^*\}$$

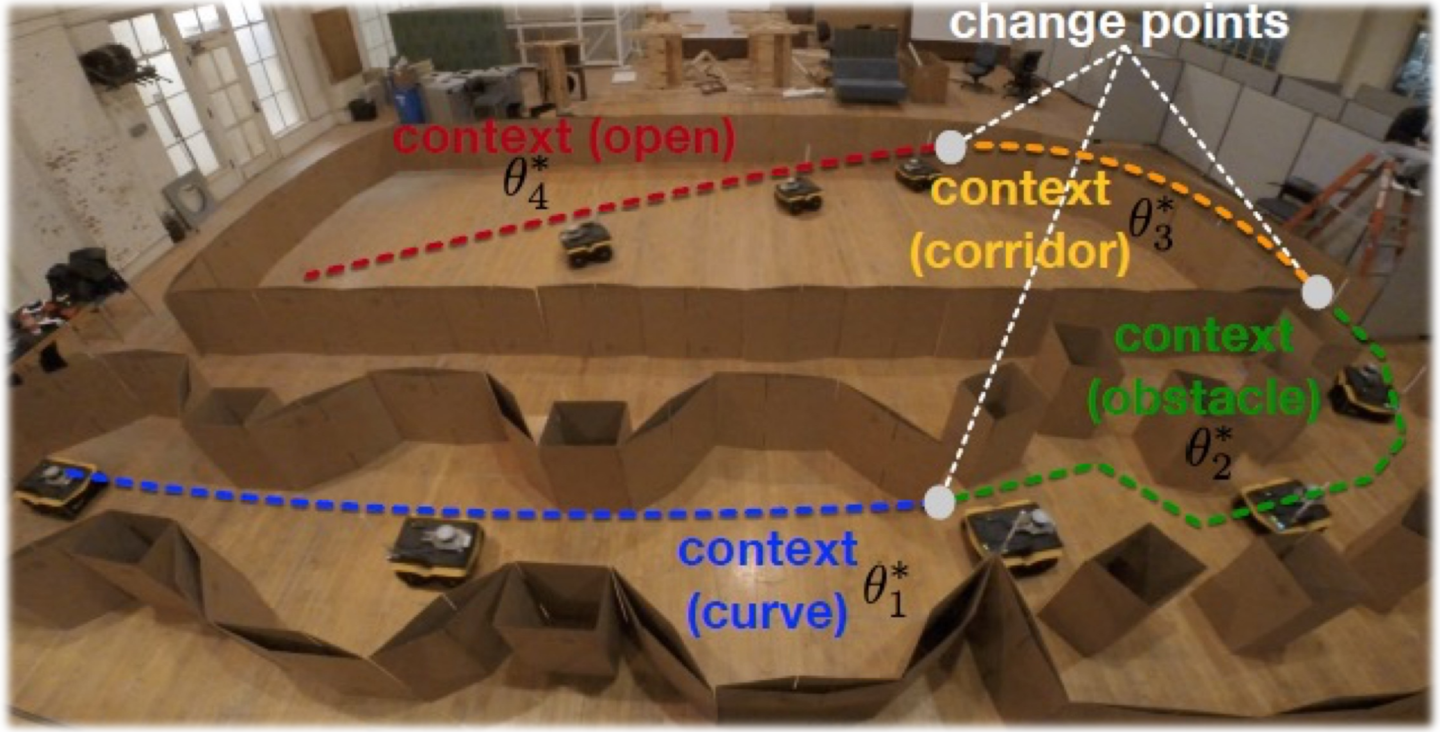

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

## APPLD Pipeline

1. Collect demonstration.  $\rightarrow \mathcal{D} = \{(x_1, u_1), \ldots, (x_N, u_N)\}$

2. Perform automatic demonstration segmentation.

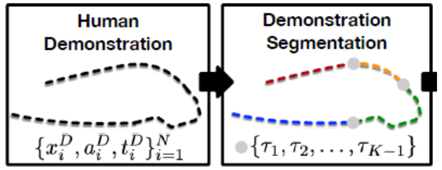


3. Use black-box optimization to find set of optimal parameters. $\mathcal{D}_k \rightarrow \boxed{A_{\text{black-box}}} \rightarrow \theta_k^*$, with $G(\cdot;\cdot) \leftrightarrow A_{\text{black-box}}$

4. Use supervised learning to train a context predictor. $x \rightarrow \boxed{B_\phi} \rightarrow k$

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

## APPLD Deployment

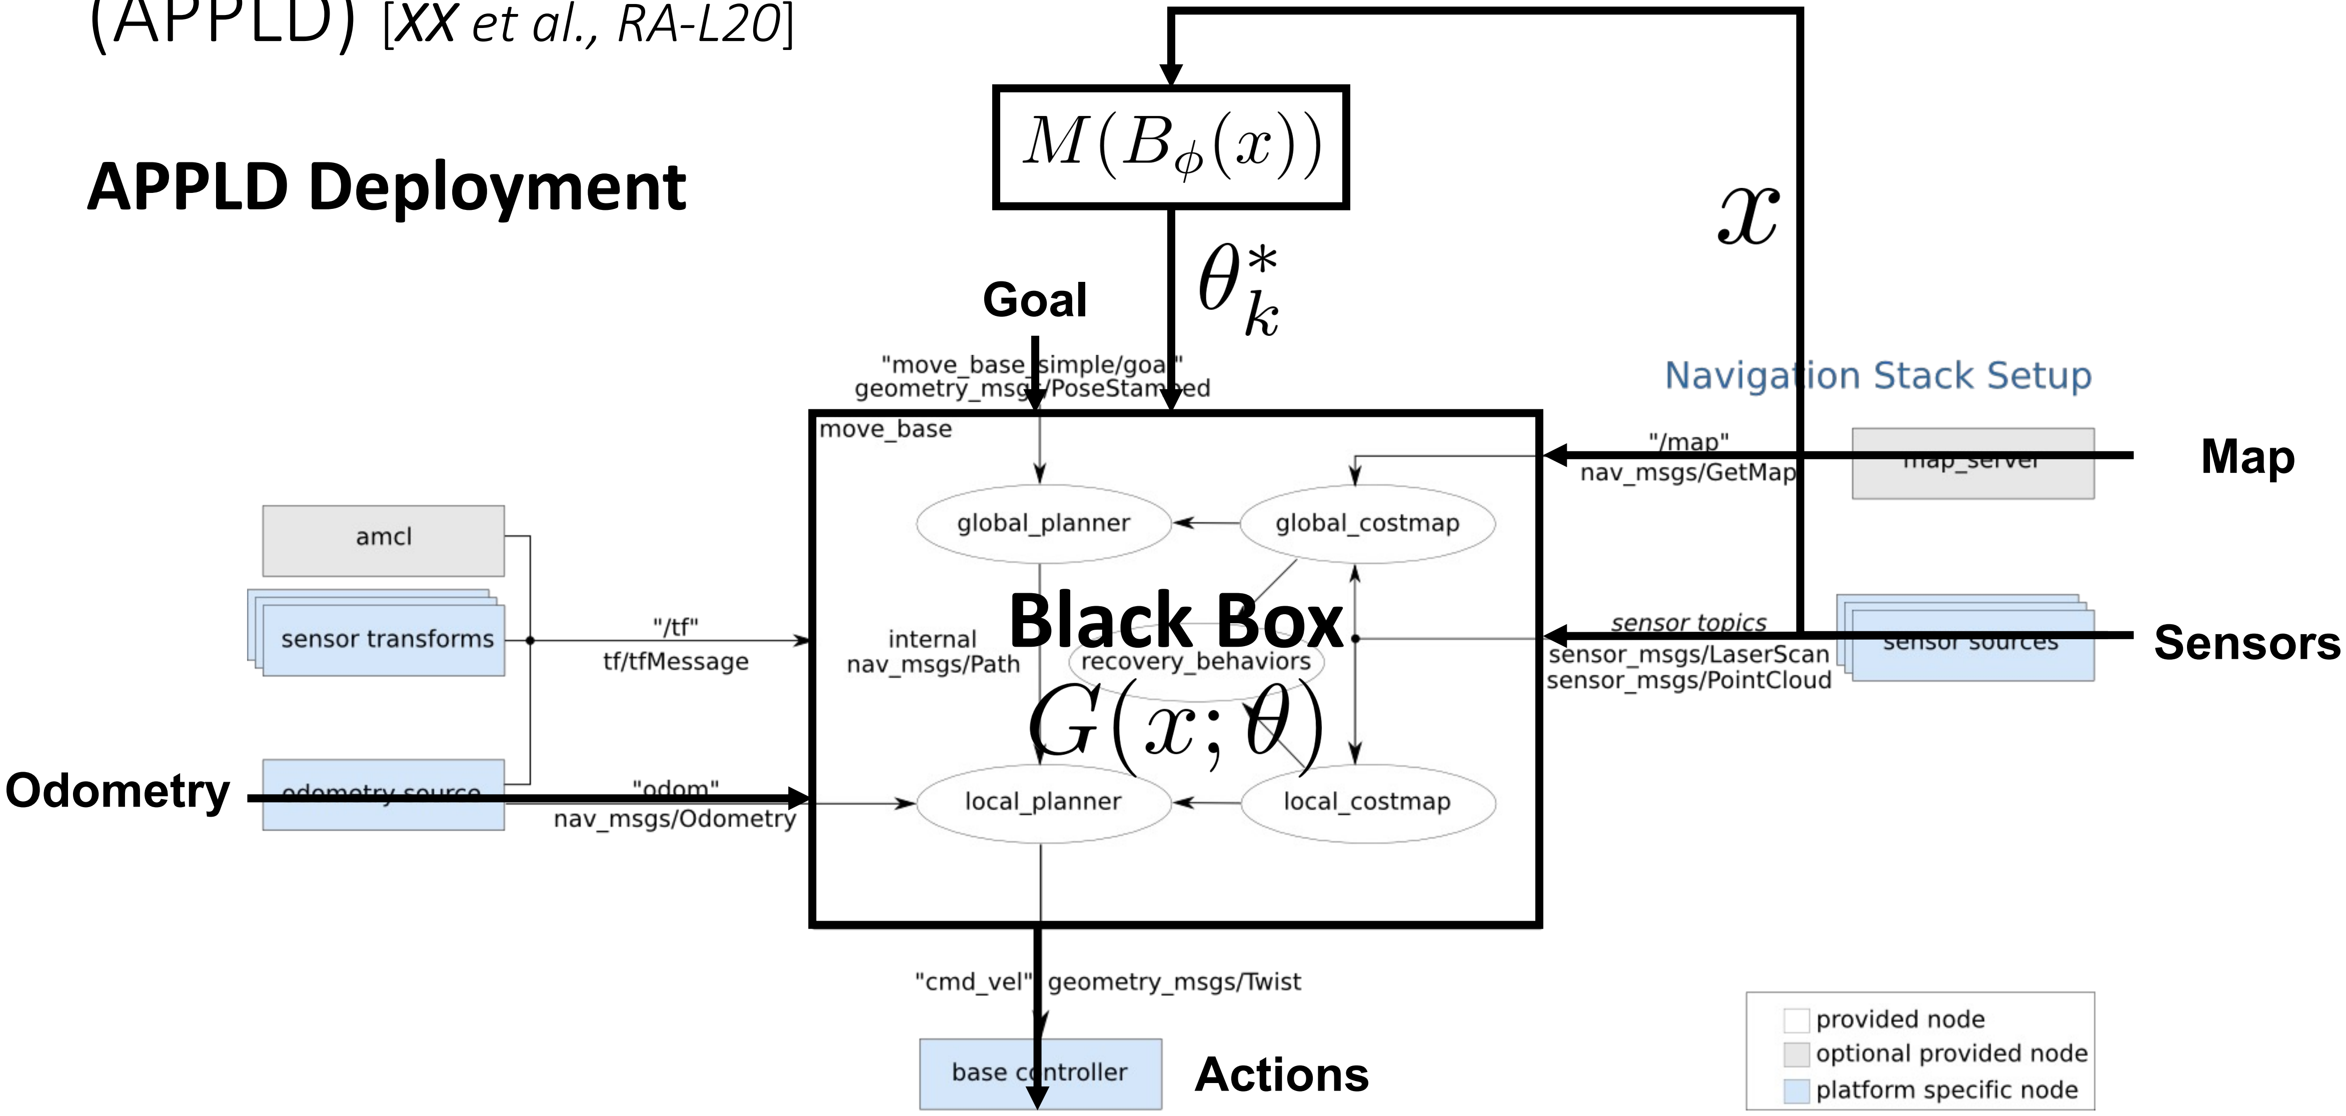

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

## Experiments

**Robot**: Clearpath Jackal (Velodyne Puck lidar)

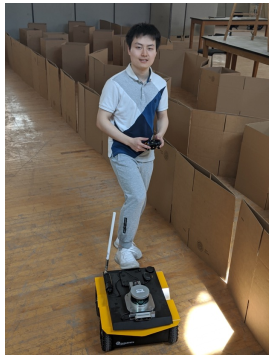

**Human**: An author (Xbox wireless controller)

**Environment**: Challenging obstacle course

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

## Deployment

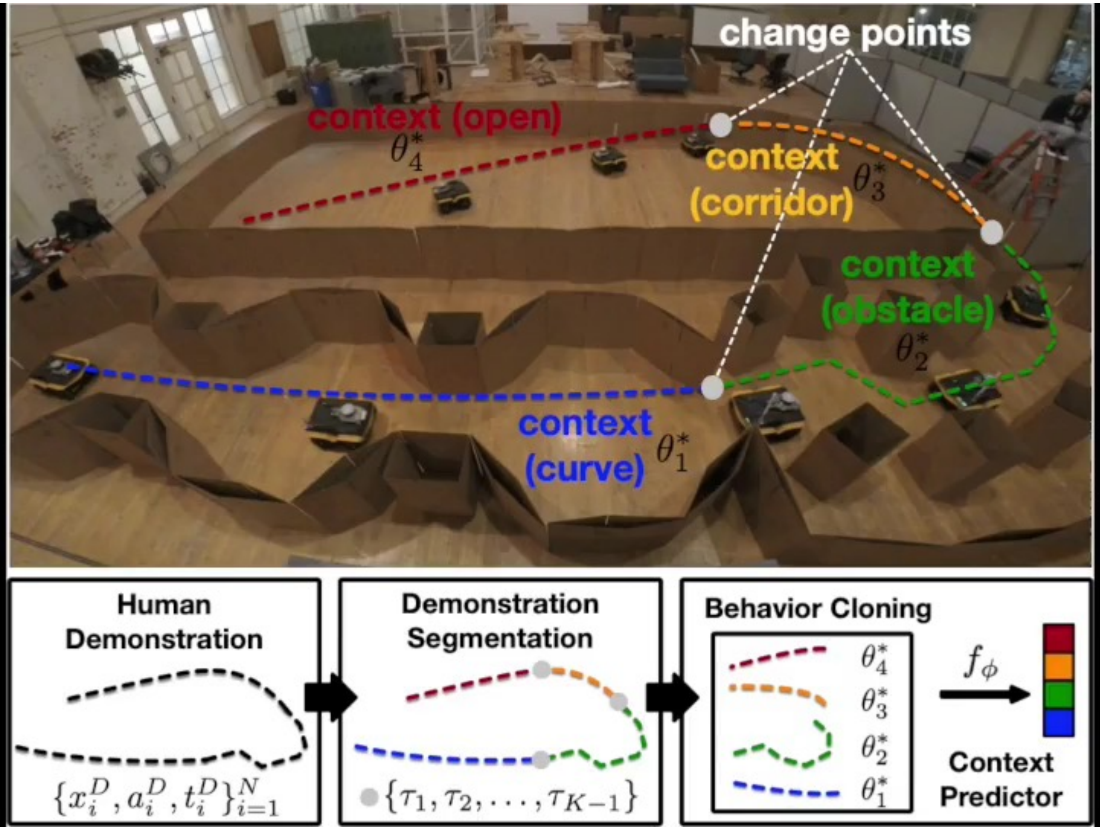

# Adaptive Planner Parameter Learning from Demonstration (APPLD) [*XX et al., RA-L20*]

Different robot, navigation stack, and environment

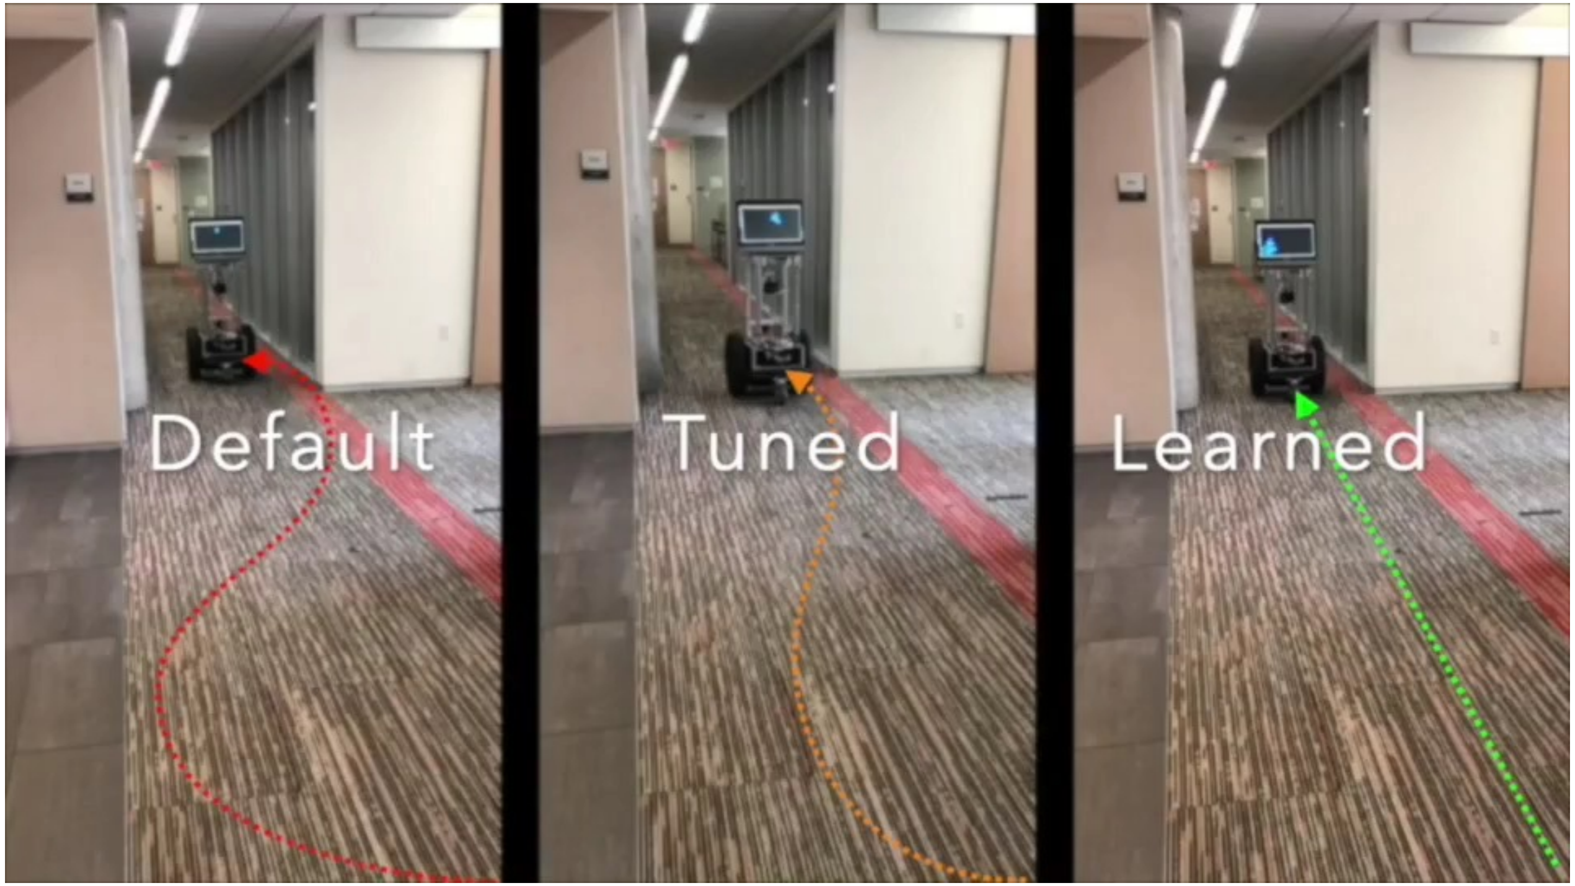

# APPL in Navigation Language

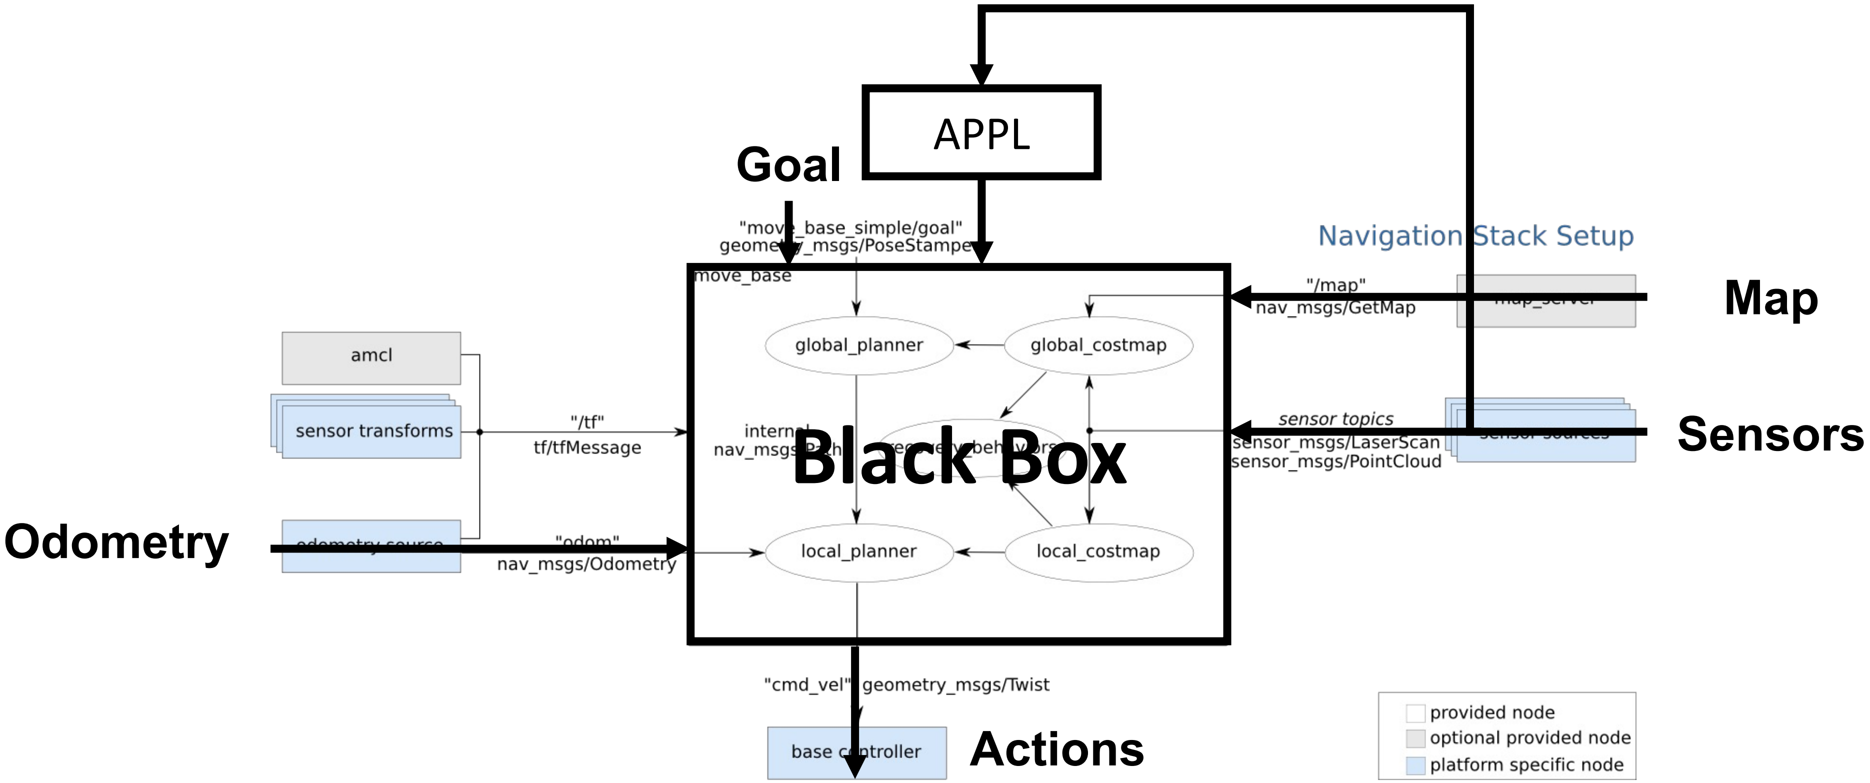


ROS **move_base** navigation stack

# APPL in Learning Language

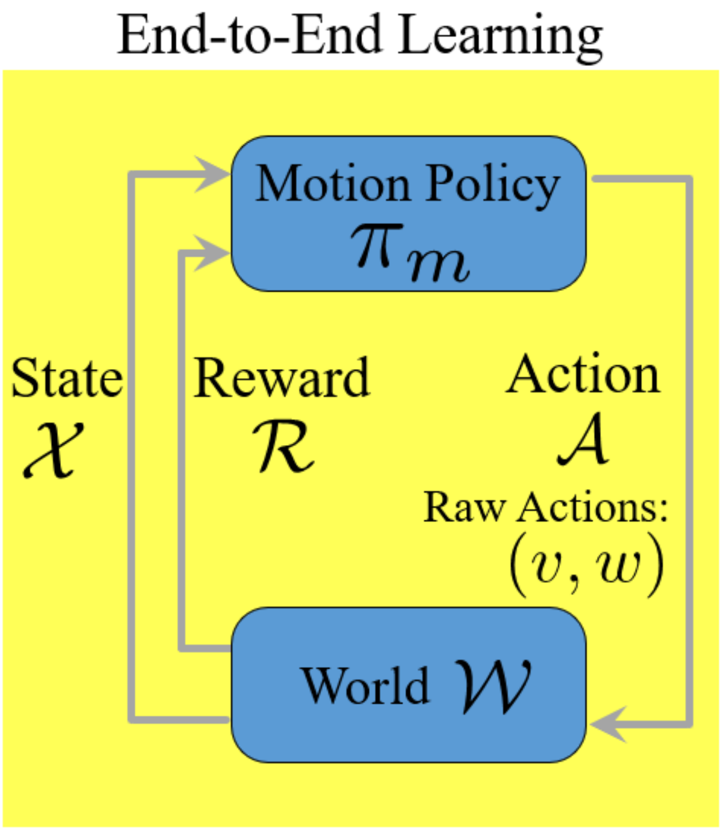


*[Gao et al., CoRL17, Pfeiffer et al., RA-L18, Faust et al., ICRA18, Chiang et al., RA-L19]*

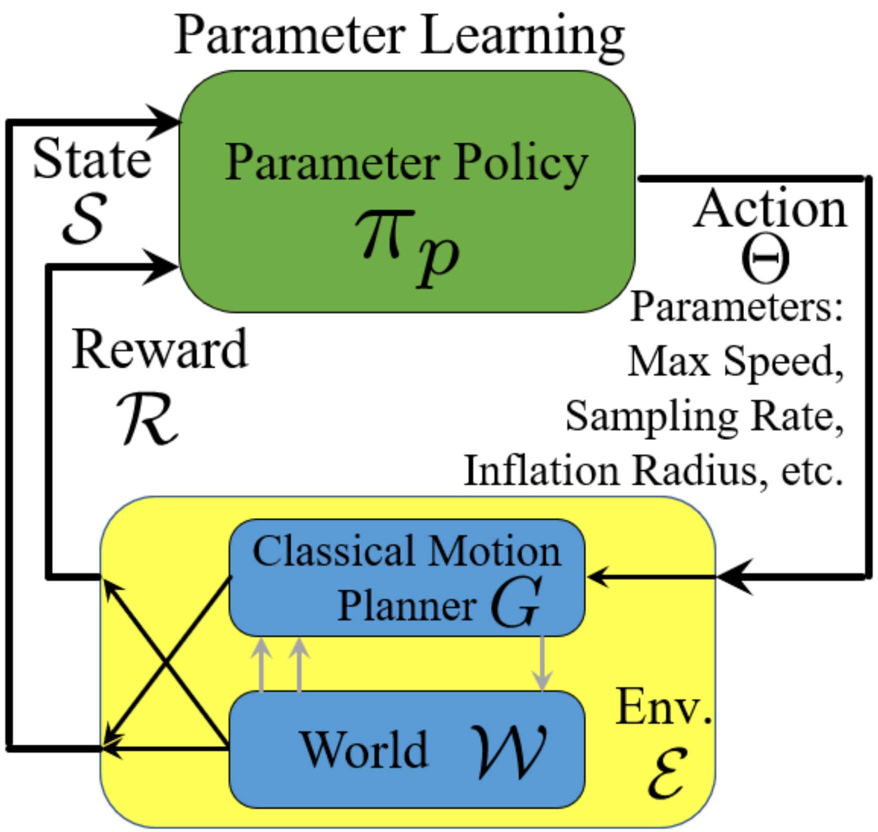


*[**XX** et al., RA-L20, W, **XX** et al., ICRA21, W, **XX** et al., RA-L21, X, D, N, **XX** et al., ICRA21]*

# APPL from Human Interactions *[Xiao et al., RAS22]*

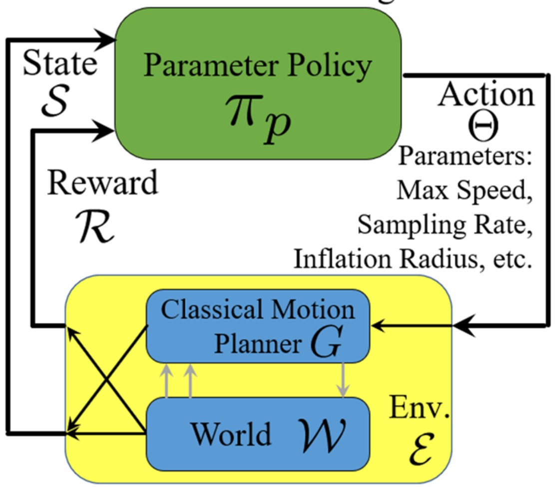


+

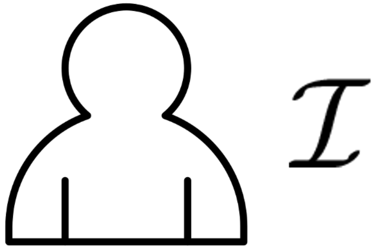


**Algorithm 1** APPL

```
// Training
Input: human interaction 𝓘, space of possible parameters Θ, and navigation stack G.
π = LearnParameterPolicy(𝓘, Θ, G).
// Deployment
Input: navigation stack G, parameter policy π.
for t = 1 : T do
    construct meta-state s_t from x_t and θ_{t−1}.
    θ_t = π(s_t).
    Navigate with G_{θ_t}(x_t).
end for
```

# Adaptive Planner Parameter Learning from Demonstration (APPLD) *[Xiao et al., RA-L20]*

APPLD imposes an internal structure to the general parameter policy.

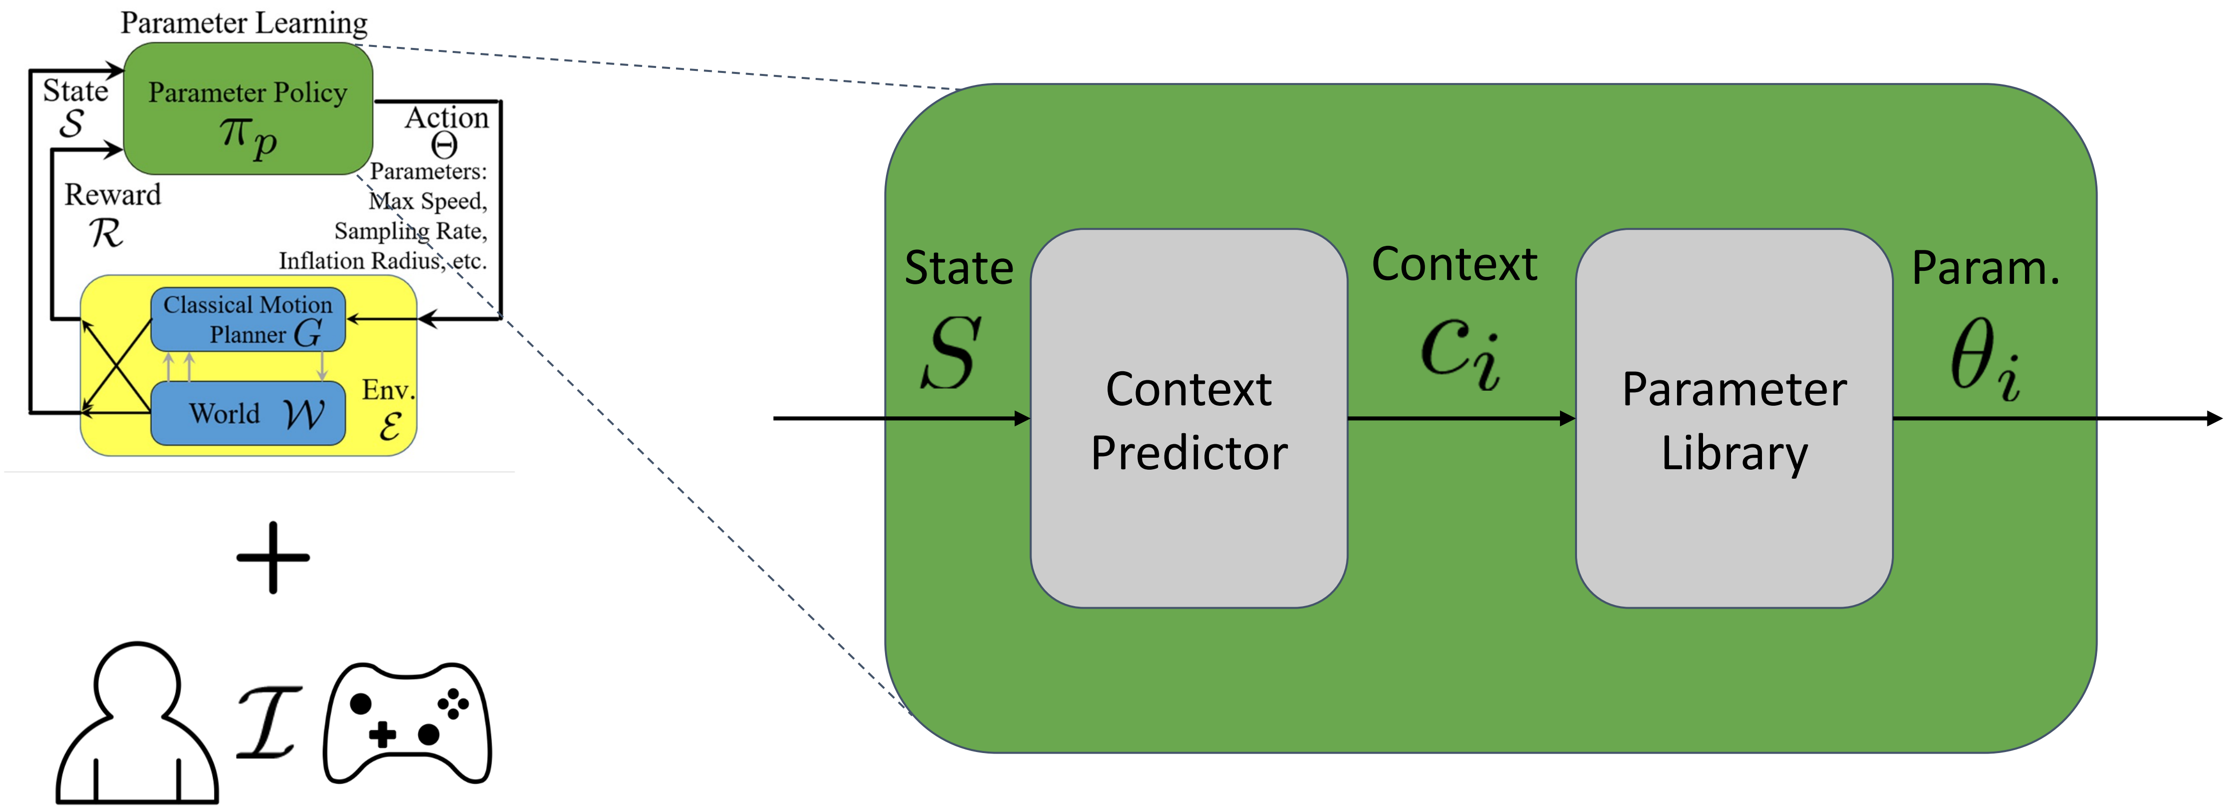

# Adaptive Planner Parameter Learning from Interventions (APPLI) *[W, XX et al., ICRA21]*

Robots do not behave suboptimally everywhere: **Intervention** when necessary.

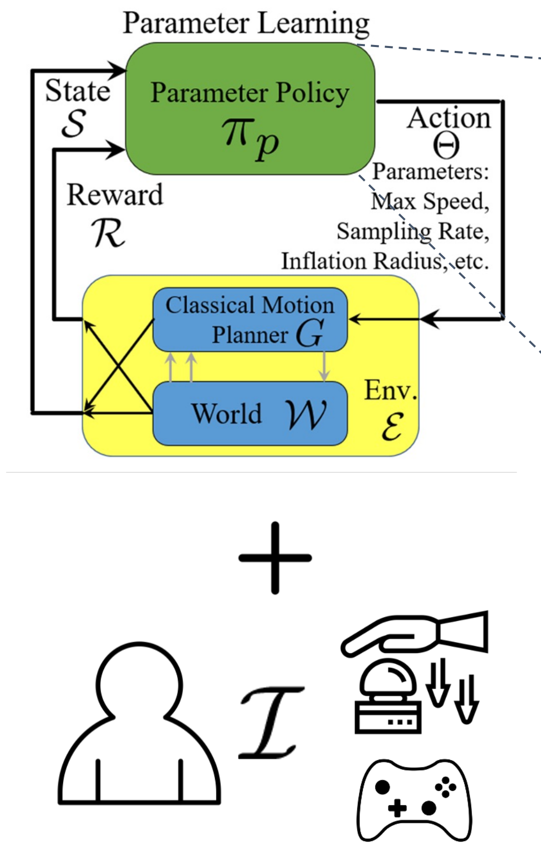


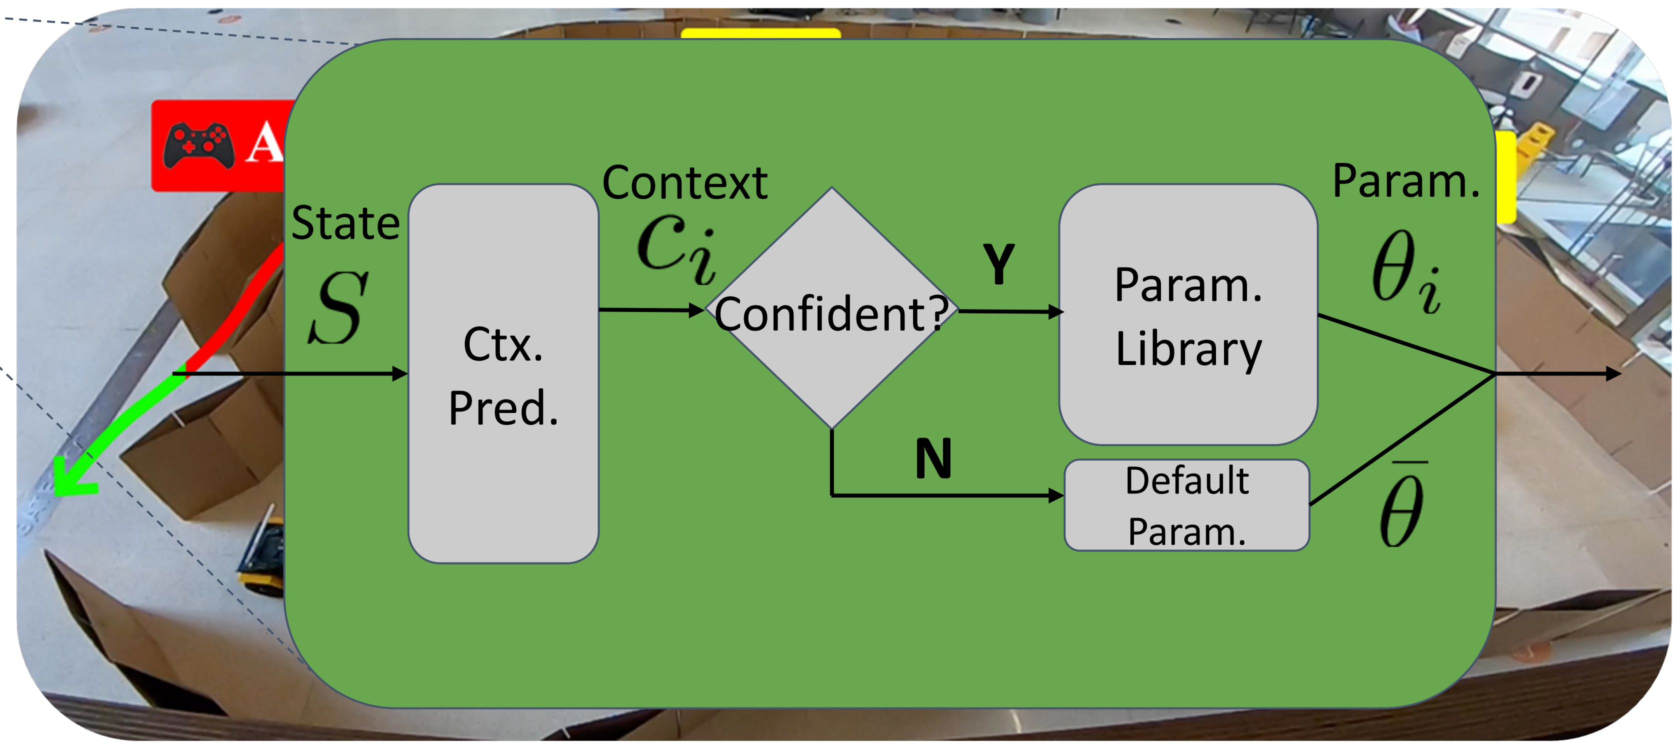

# Adaptive Planner Parameter Learning from Evaluative Feedback (APPLE) *[W, XX et al., RA-L21]*

Non-expert users may not be able to take control of the robot: **Evaluative feedback**.

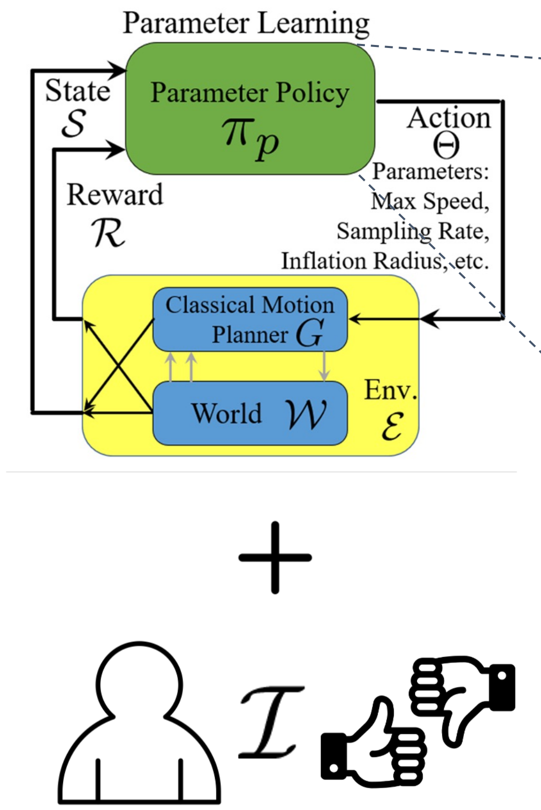


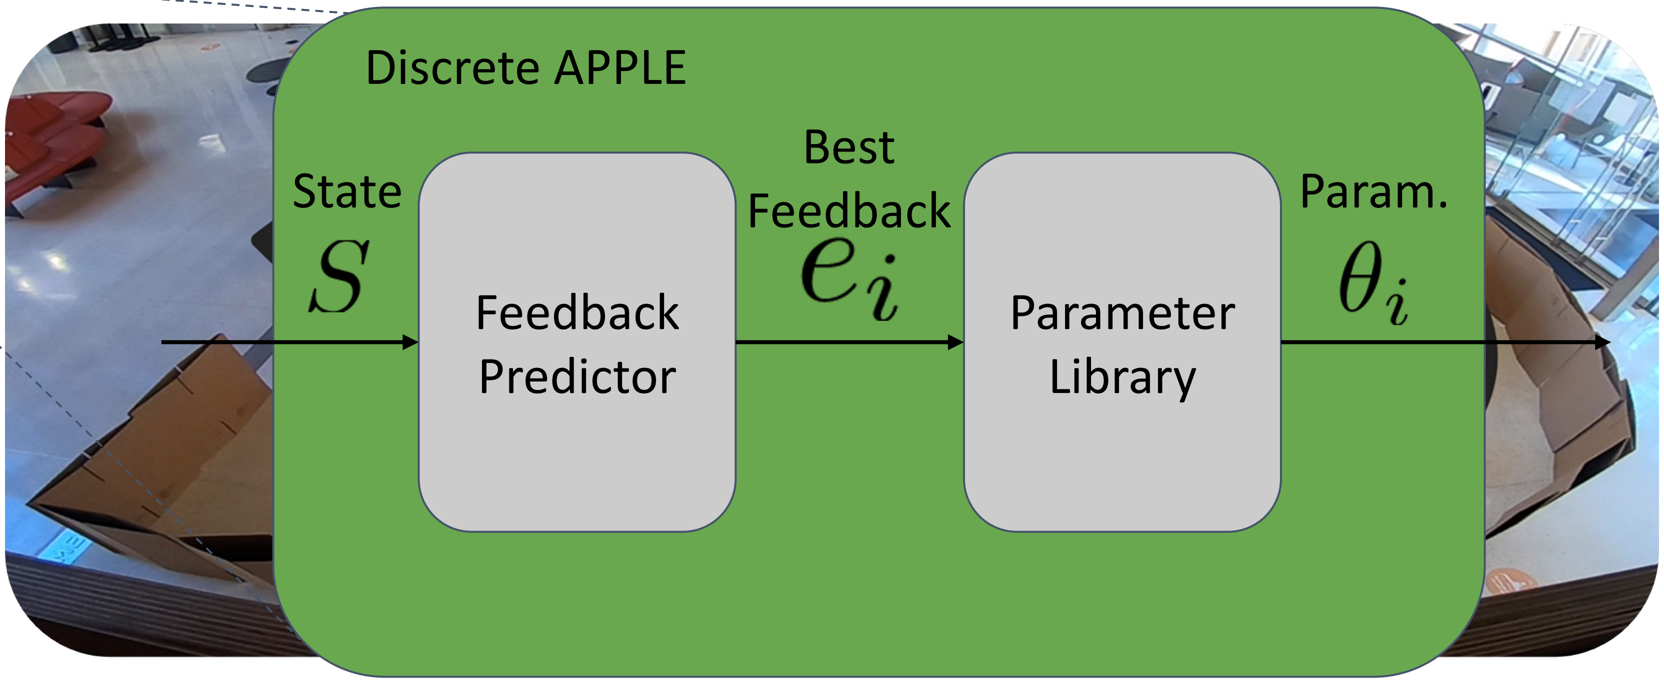

# Adaptive Planner Parameter Learning from Evaluative Feedback (APPLE) *[W, XX et al., RA-L21]*

Non-expert users may not be able to take control of the robot: **Evaluative feedback**.

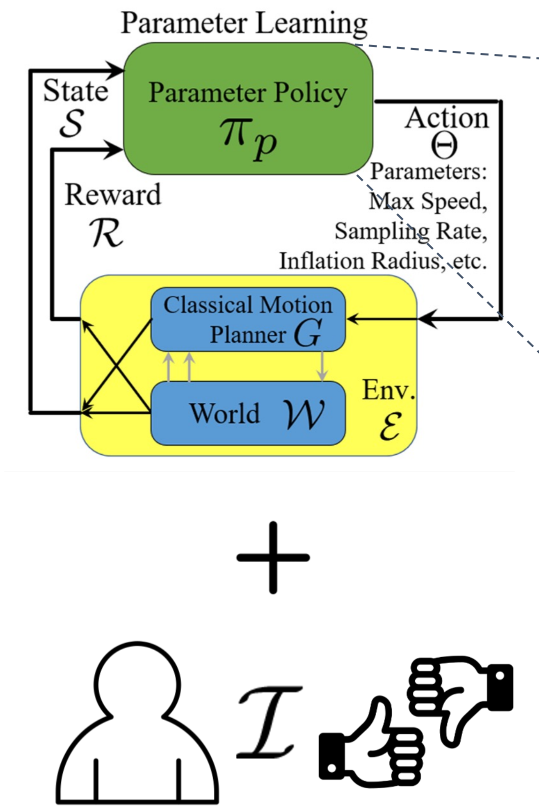


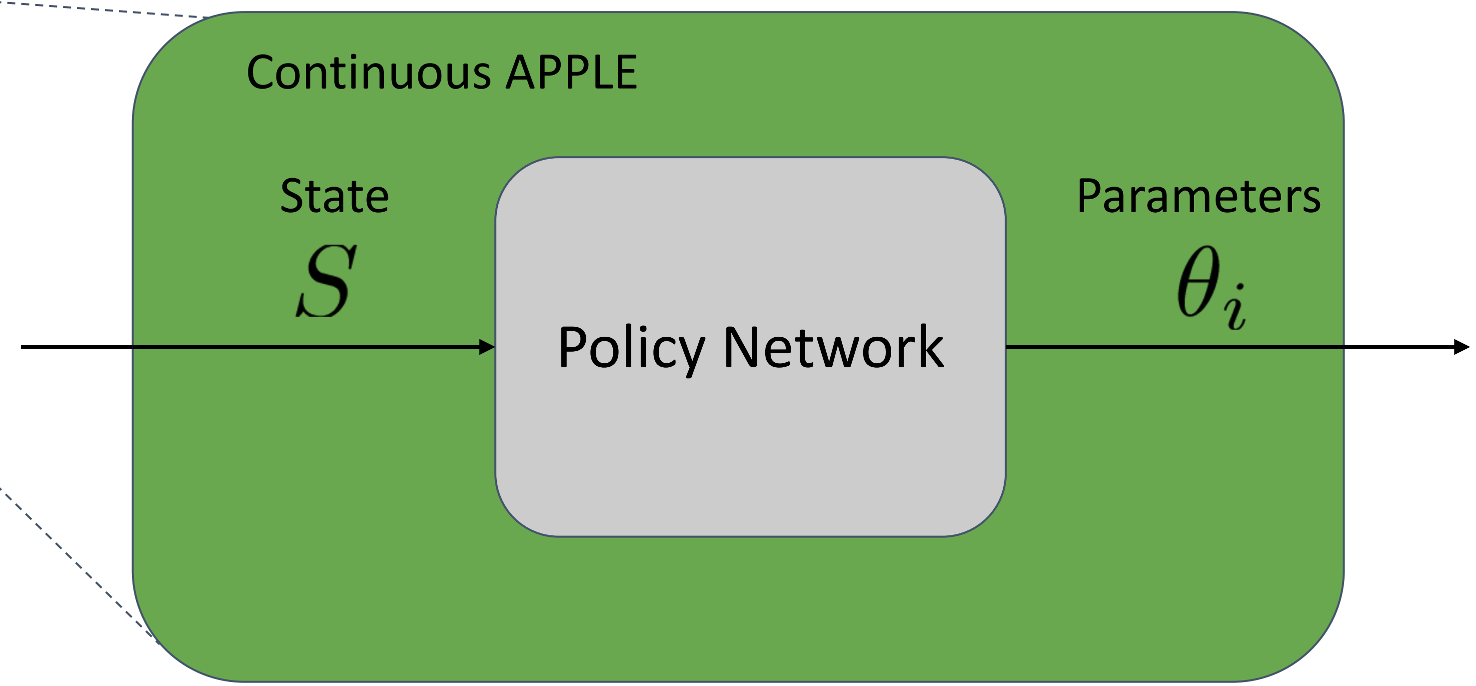

# Adaptive Planner Parameter Learning from Reinforcement (APPLR) *[X, D, N, **XX** et al., ICRA21]*

What about no humans at all? **Reinforcement Learning.**

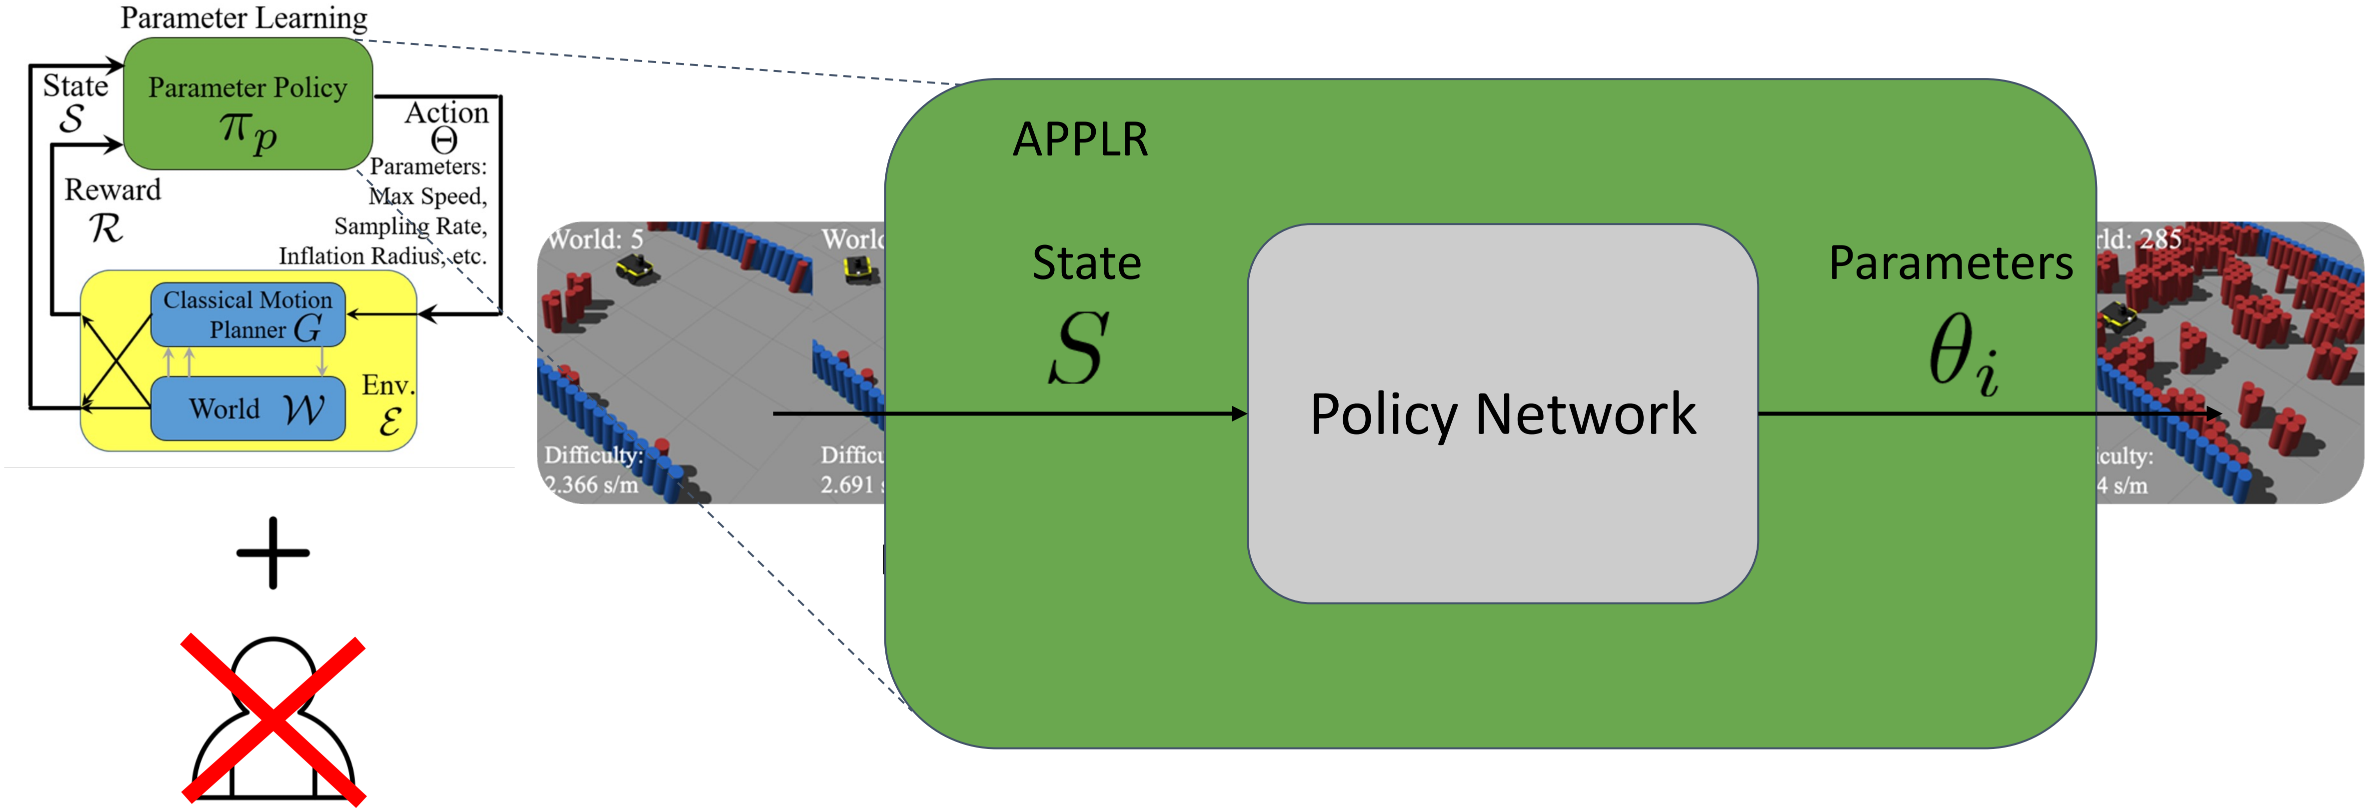

# Cycle-of-Learning from APPL *[XX et al., RAS22]*

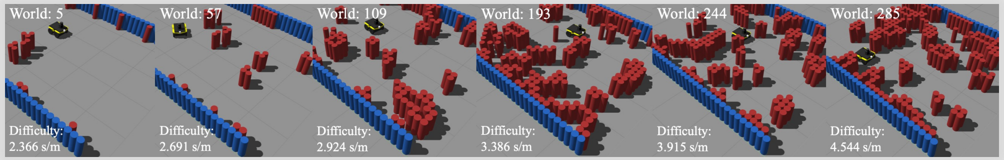


APPLR

*[X, D, N, **XX** et al., ICRA21]*

APPLD

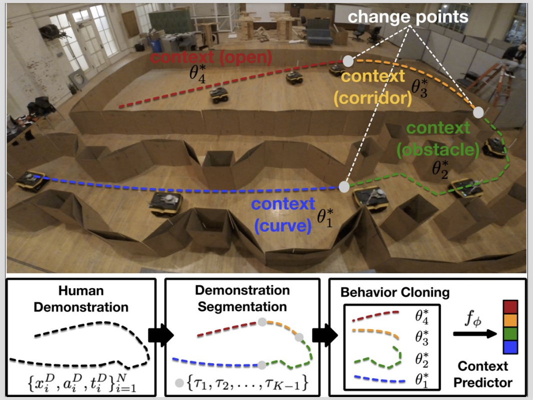


APPLD

*[**XX** et al., RA-L20]*

APPLR

APPLI

APPLE

*[W, **XX** et al., RA-L21]*

APPLE

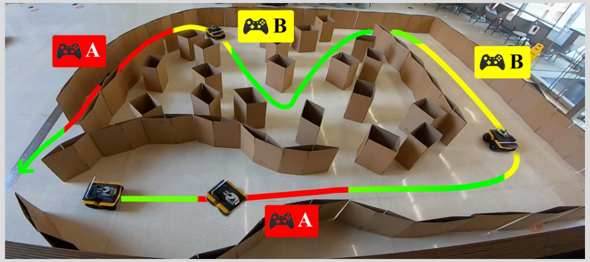


APPLI

*[W, **XX** et al., ICRA21]*

# Cycle-of-Learning from APPL (Future Work)

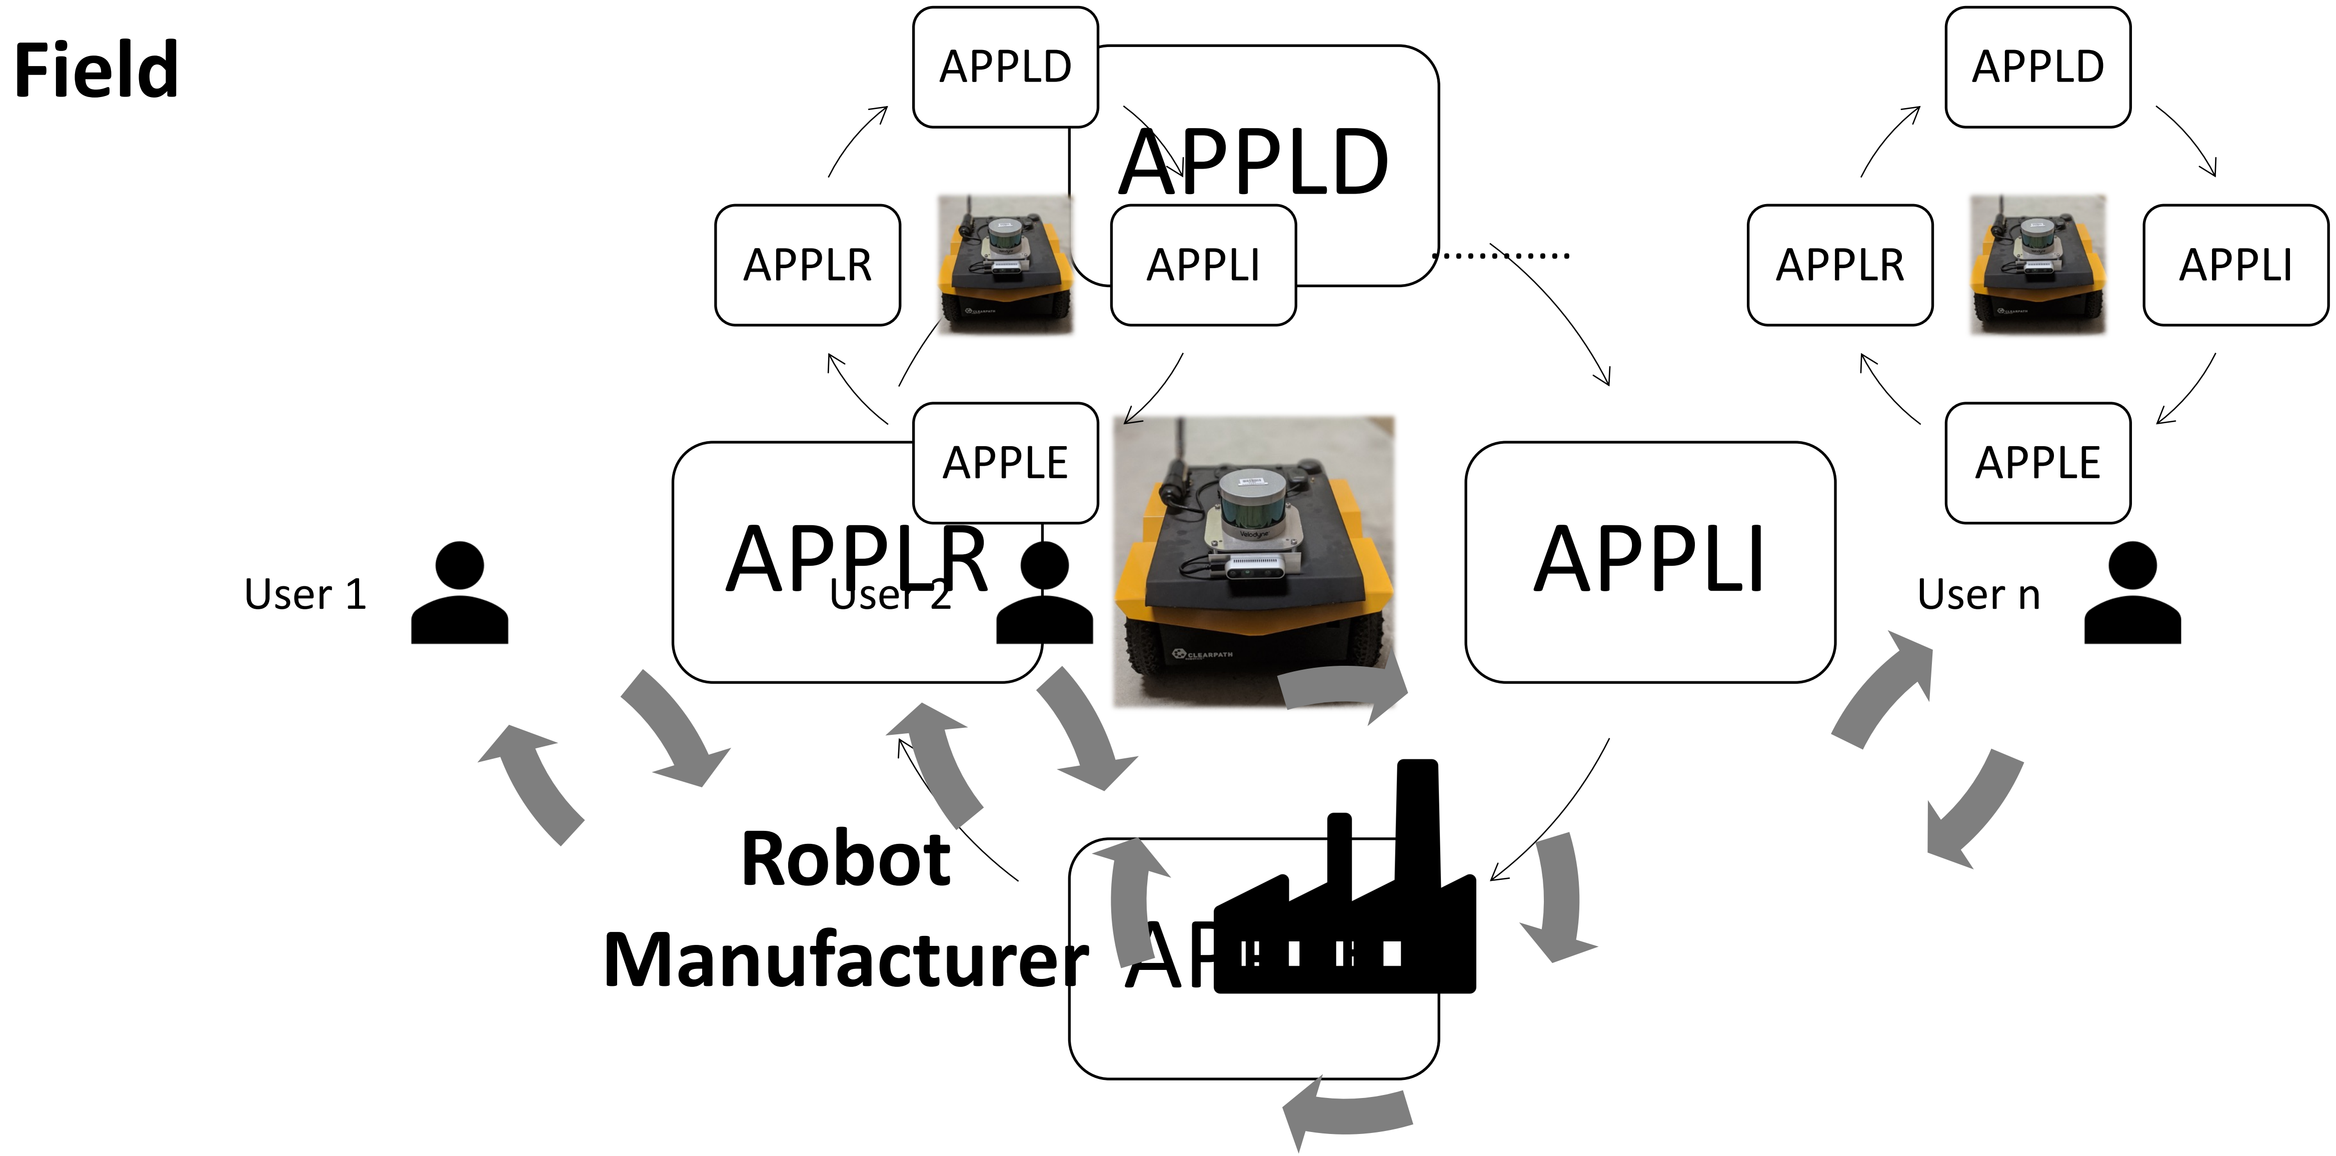

# Human-Interactive Mobile Robots: from Learning to Deployment

- This talk: Human-interactive mobile robots that efficiently learn from and harmoniously deploy among humans
  - Adaptive Planner Parameter Learning (APPL) to utilize the opportunities from easily available non-expert human interactions.
  - Datasets, protocols, principles, guidelines, and learning methods to address social robot navigation challenges.

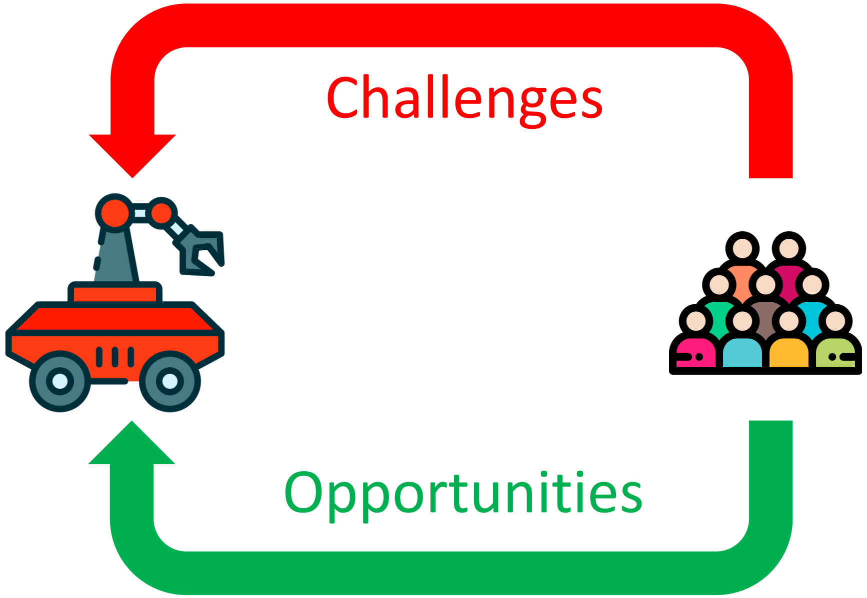

# SCAND: A Large-Scale Dataset of Socially Compliant Navigation Demonstration [K, N, *XX* et al., RA-L22]

- 40km (8.7 hours) of real-world data (~0.5TB)
- 138 trajectories, 15 days
- Data collected on two robots: Jackal and Spot
- Indoor and outdoor environments @ UT Austin
- Four different human demonstrators
- Coarse labels of social interactions

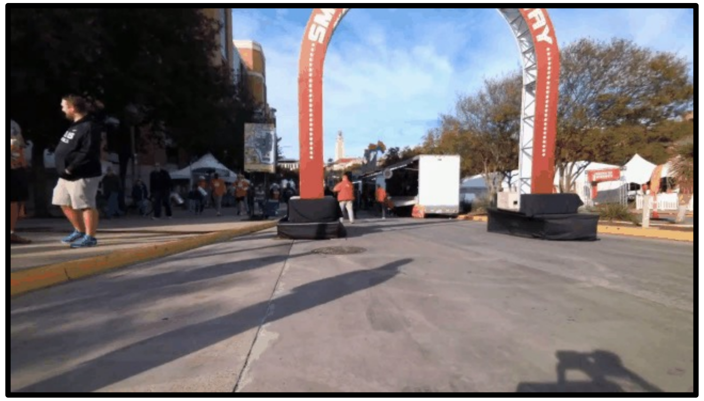

# MuSoHu: Multi-Modal Social Human Navigation Dataset *[N, N, P, D, **XX**, IROS23]*

- 100km, 20 hours, 300 trials, 13 humans, **and counting**!

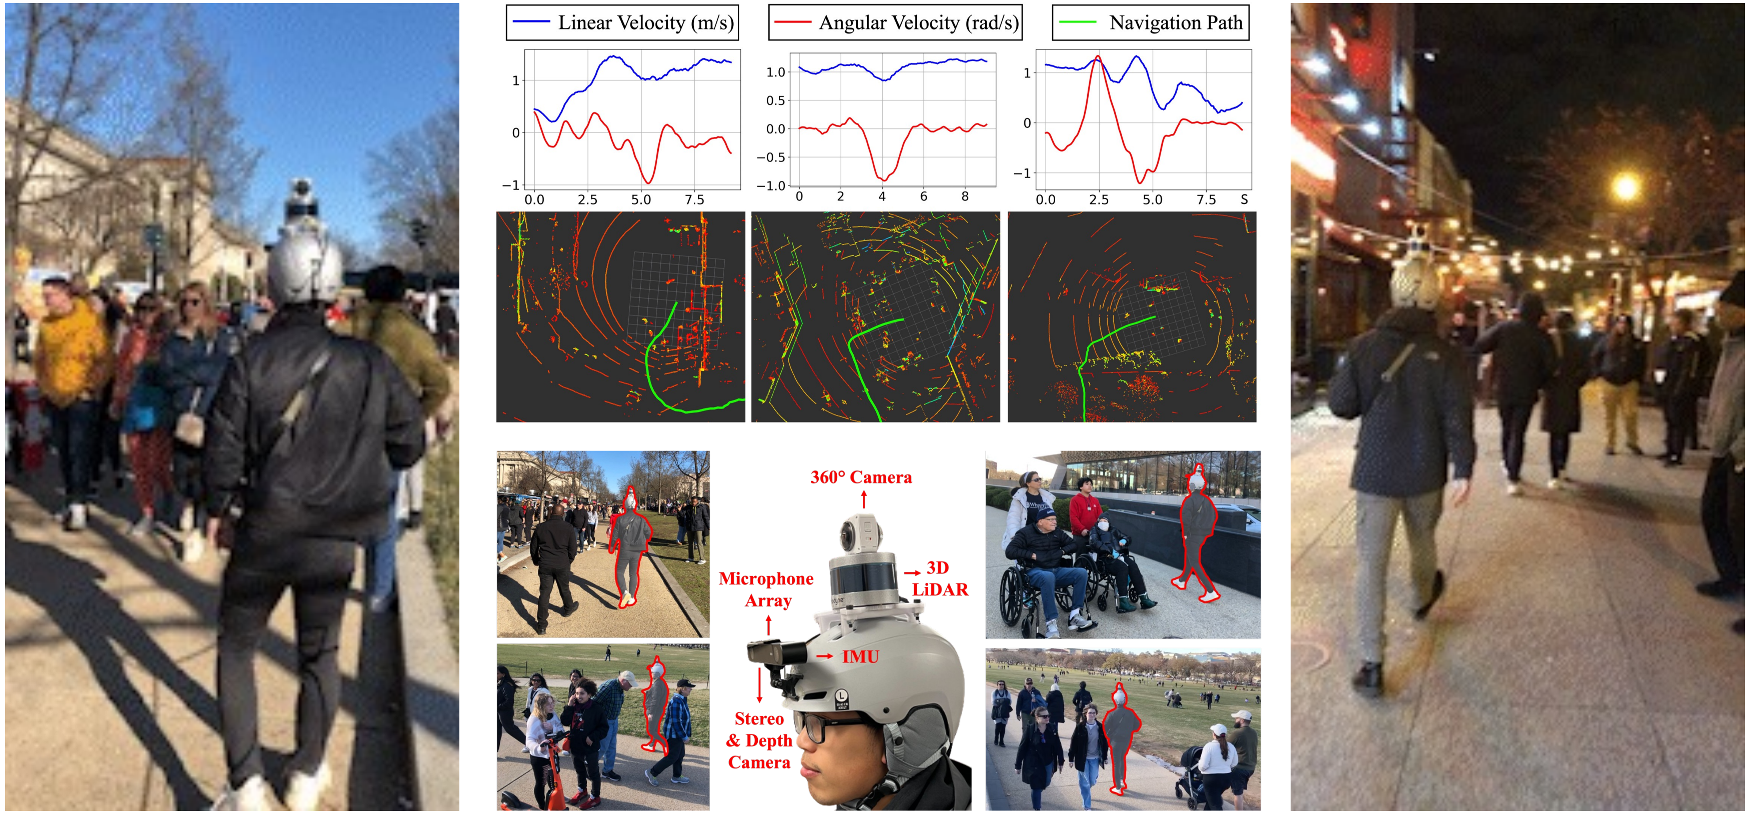

# A Protocol for Validating Social Navigation Policies *[P, L, **XX** et al., ICRA22WS]*

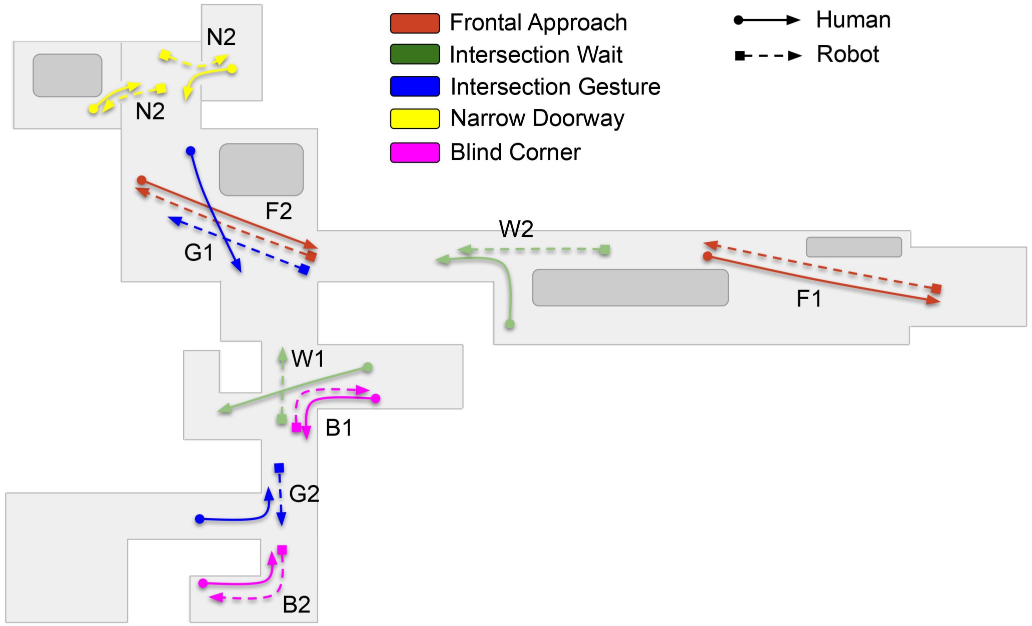


| | |
|---|---|
| **Frontal Approach** | |
| 1 | The robot moved to avoid me. |
| 2* | The robot obstructed my path. |
| 3 | The robot maintained a safe and comfortable distance at all times. |
| 4* | The robot nearly collided with me. |
| 5 | It was clear what the robot wanted to do. |
| **Intersection Wait** | |
| 6 | The robot let me cross the intersection by maintaining a safe and comfortable distance. |
| 7 | The robot changed course to let me pass. |
| 8 | The robot paid attention to what I was doing. |
| 9 | The robot slowed down and stopped to let me pass. |
| **Intersection Gesture** | |
| 10 | The robot maintained a safe and comfortable distance at all times. |
| 11 | The robot slowed down and stopped. |
| 12 | The robot followed my command. |
| 13 | I felt the robot paid attention to what I was doing. |
| **Narrow Doorway** | |
| 14* | The robot got in my way. |
| 15 | The robot moved to avoid me. |
| 16 | The robot made room for me to enter or exit. |
| 17* | It was clear what the robot wanted to do. |
| **Blind Corner** | |
| 18 | The robot moved to avoid me. |
| 19 | The robot stopped to let me pass. |
| 20* | I had to move around the robot. |
| 21* | The robot nearly collided with me head-on. |

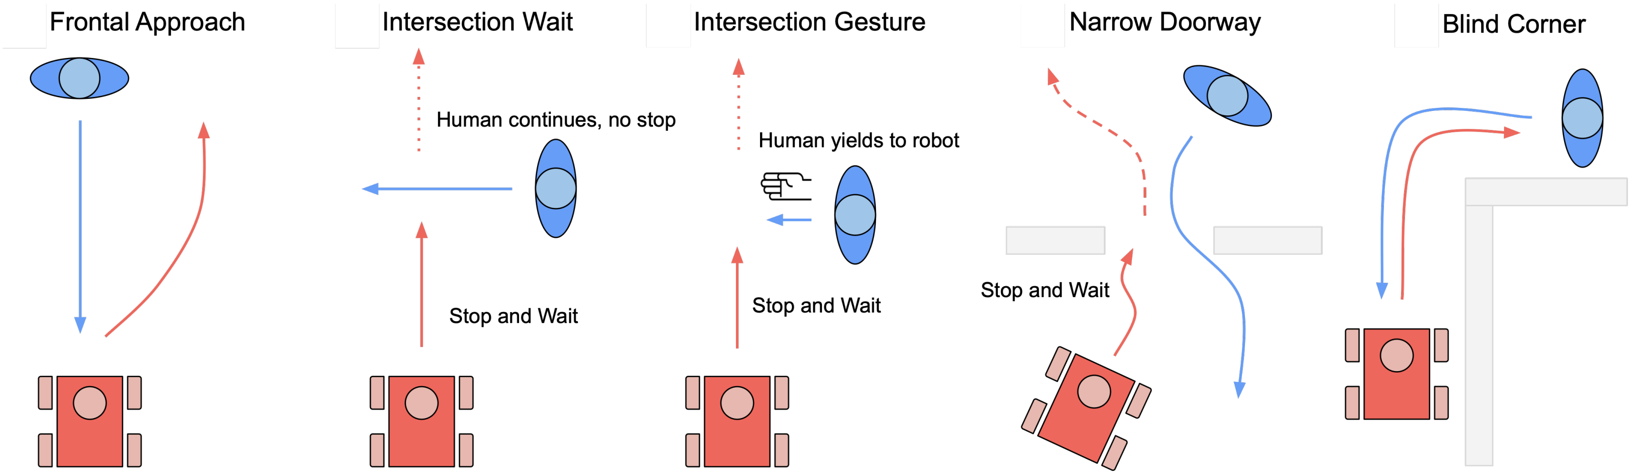

# Principles and Guidelines for Evaluating Social Robot Navigation Algorithms [*F, …,* ***XX*** *et al., under review*]

- Anthony Francis, Claudia Perez-D'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, Hao-Tien Lewis Chiang, Michael Everett, Sehoon Ha, Justin Hart, Jonathan P. How, Haresh Karnan, Tsang-Wei Edward Lee, Luis J. Manso, Reuth Mirksy, Soeren Pirk, Phani Teja Singamaneni, Peter Stone, Ada V. Taylor, Peter Trautman, Nathan Tsoi, Marynel Vazquez, **Xuesu Xiao**, Peng Xu, Naoki Yokoyama, Alexander Toshev, and Roberto Martin-Martin

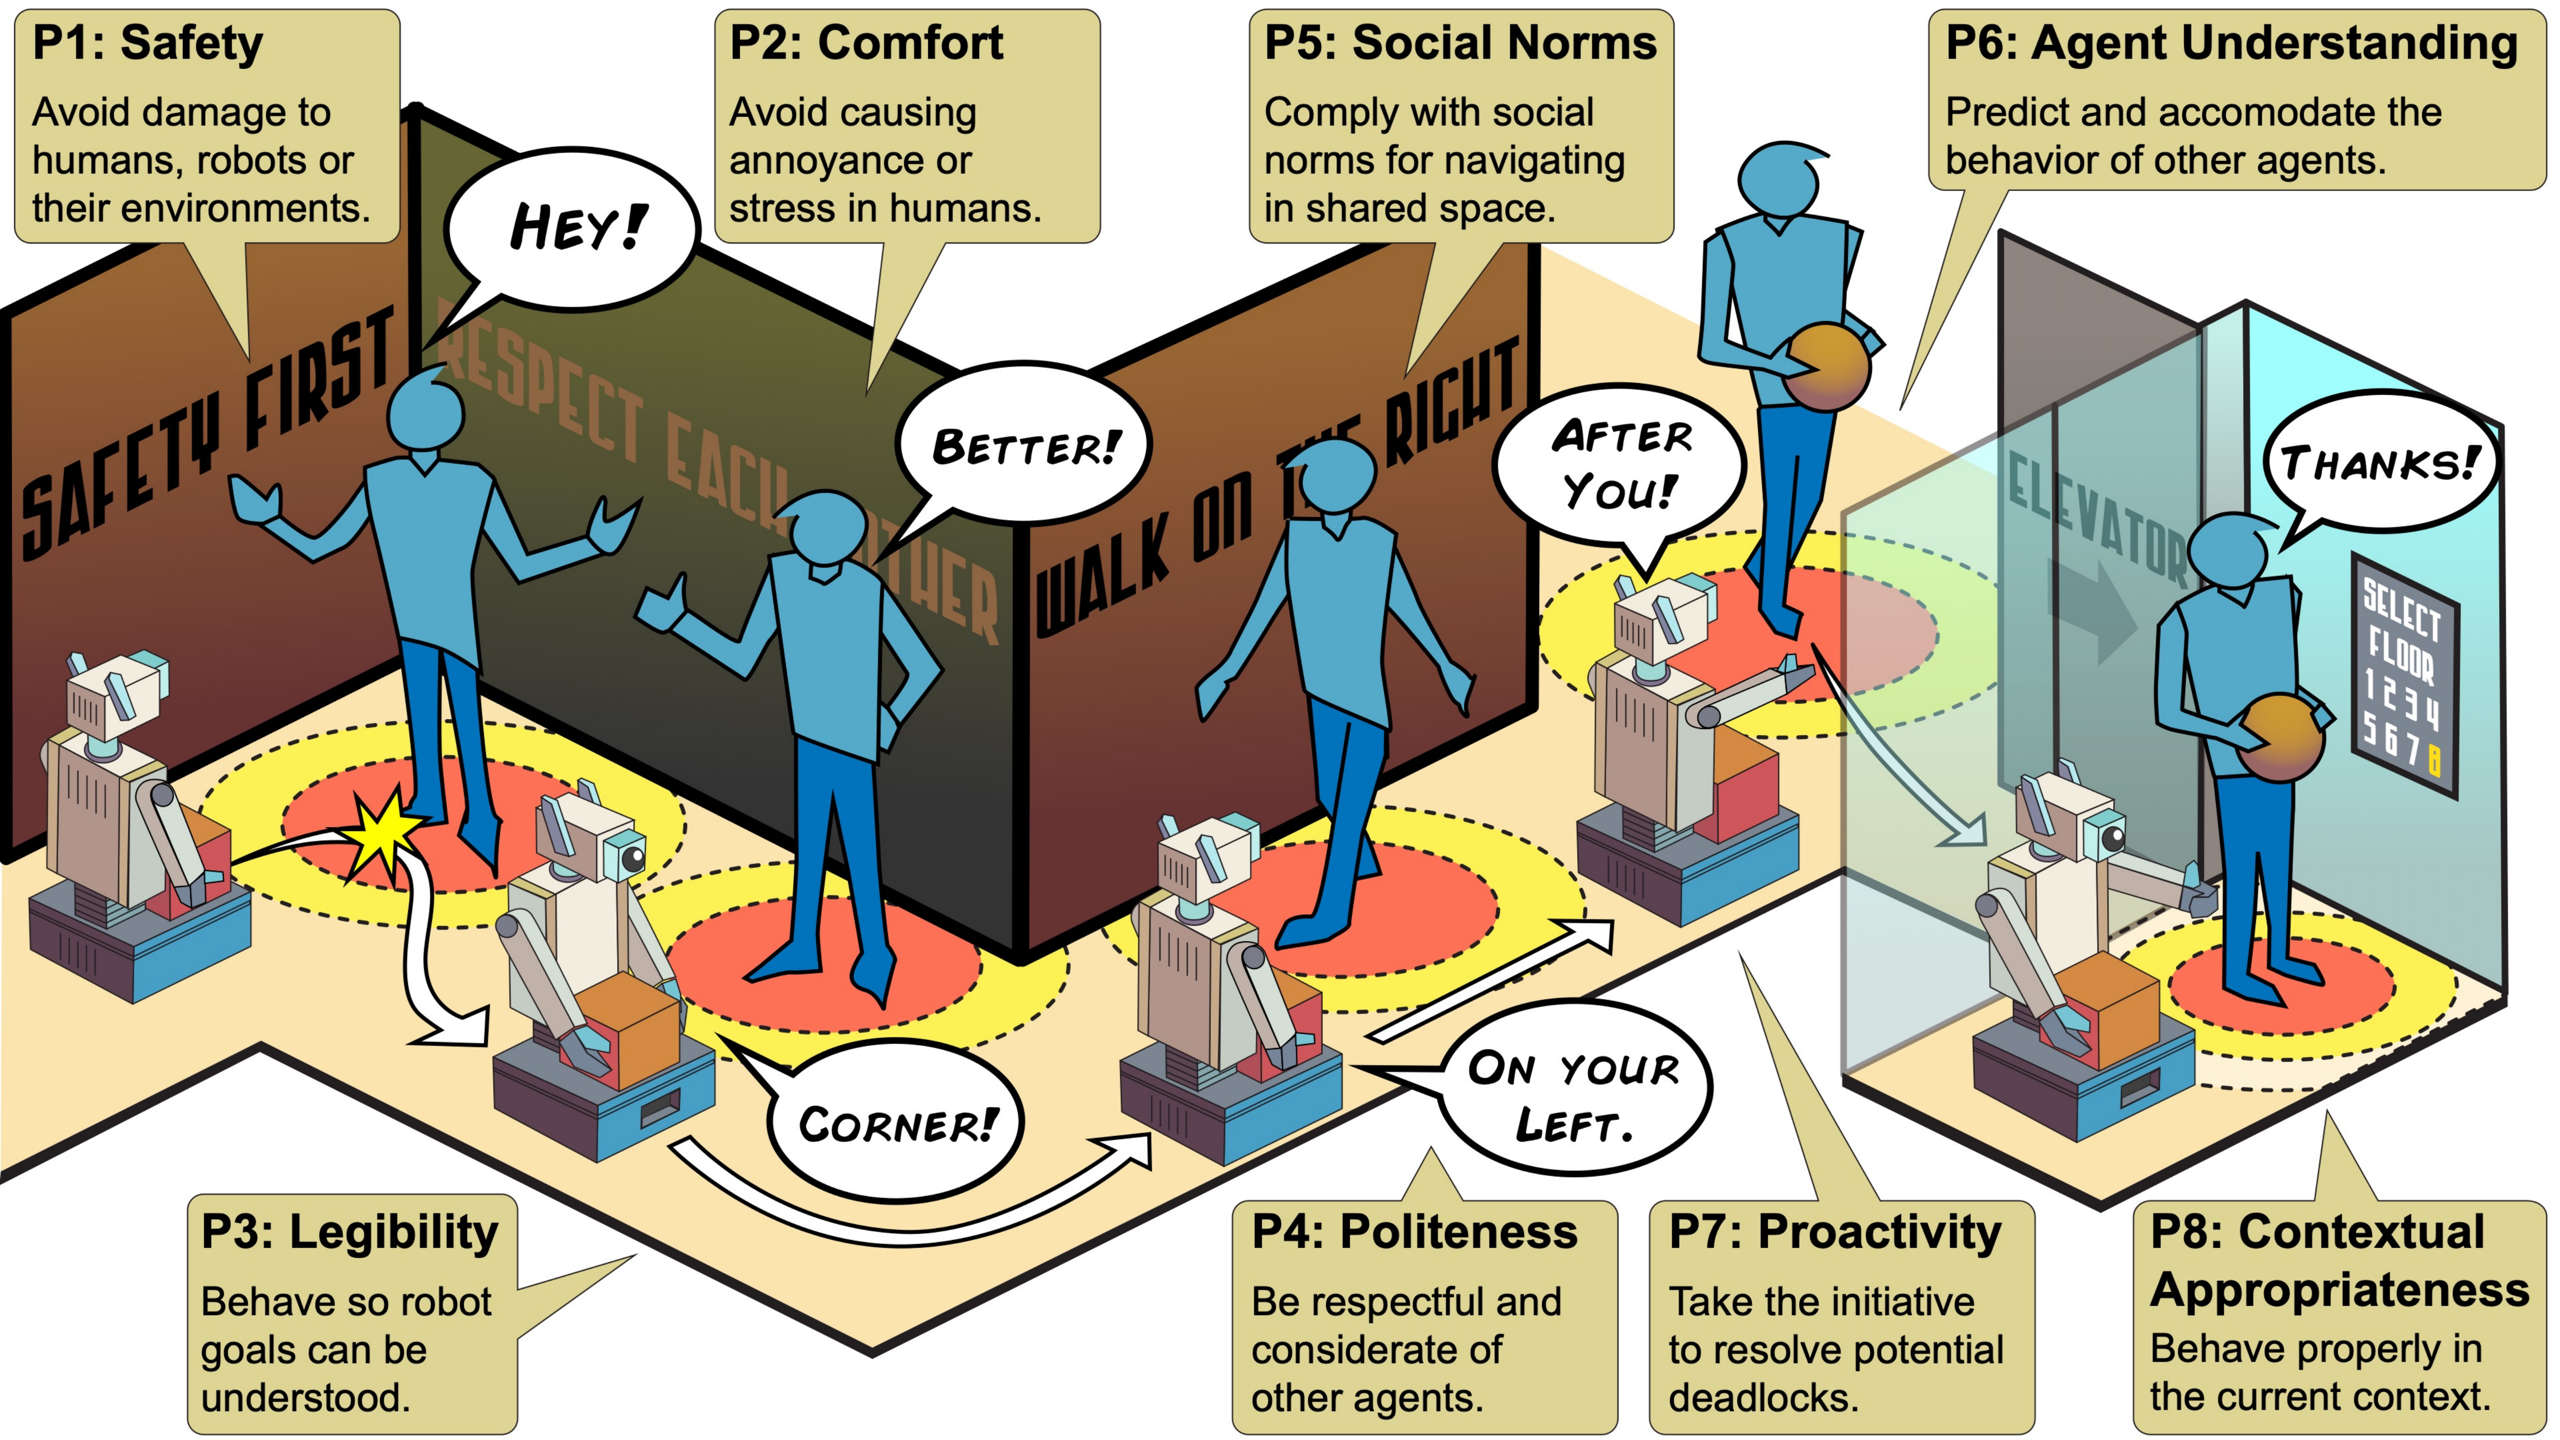
P1: Safety
Avoid damage to humans, robots or their environments.
P2: Comfort
Avoid causing annoyance or stress in humans.
P5: Social Norms
Comply with social norms for navigating in shared space.
P6: Agent Understanding
Predict and accomodate the behavior of other agents.
HEY!
SAFETY FIRST
RESPECT EACH OTHER
BETTER!
WALK ON THE RIGHT
AFTER YOU!
ELEVATOR
THANKS!
SELECT FLOOR
1 2 3 4
5 6 7 8
CORNER!
ON YOUR LEFT.
P3: Legibility
Behave so robot goals can be understood.
P4: Politeness
Be respectful and considerate of other agents.
P7: Proactivity
Take the initiative to resolve potential deadlocks.
P8: Contextual Appropriateness
Behave properly in the current context.

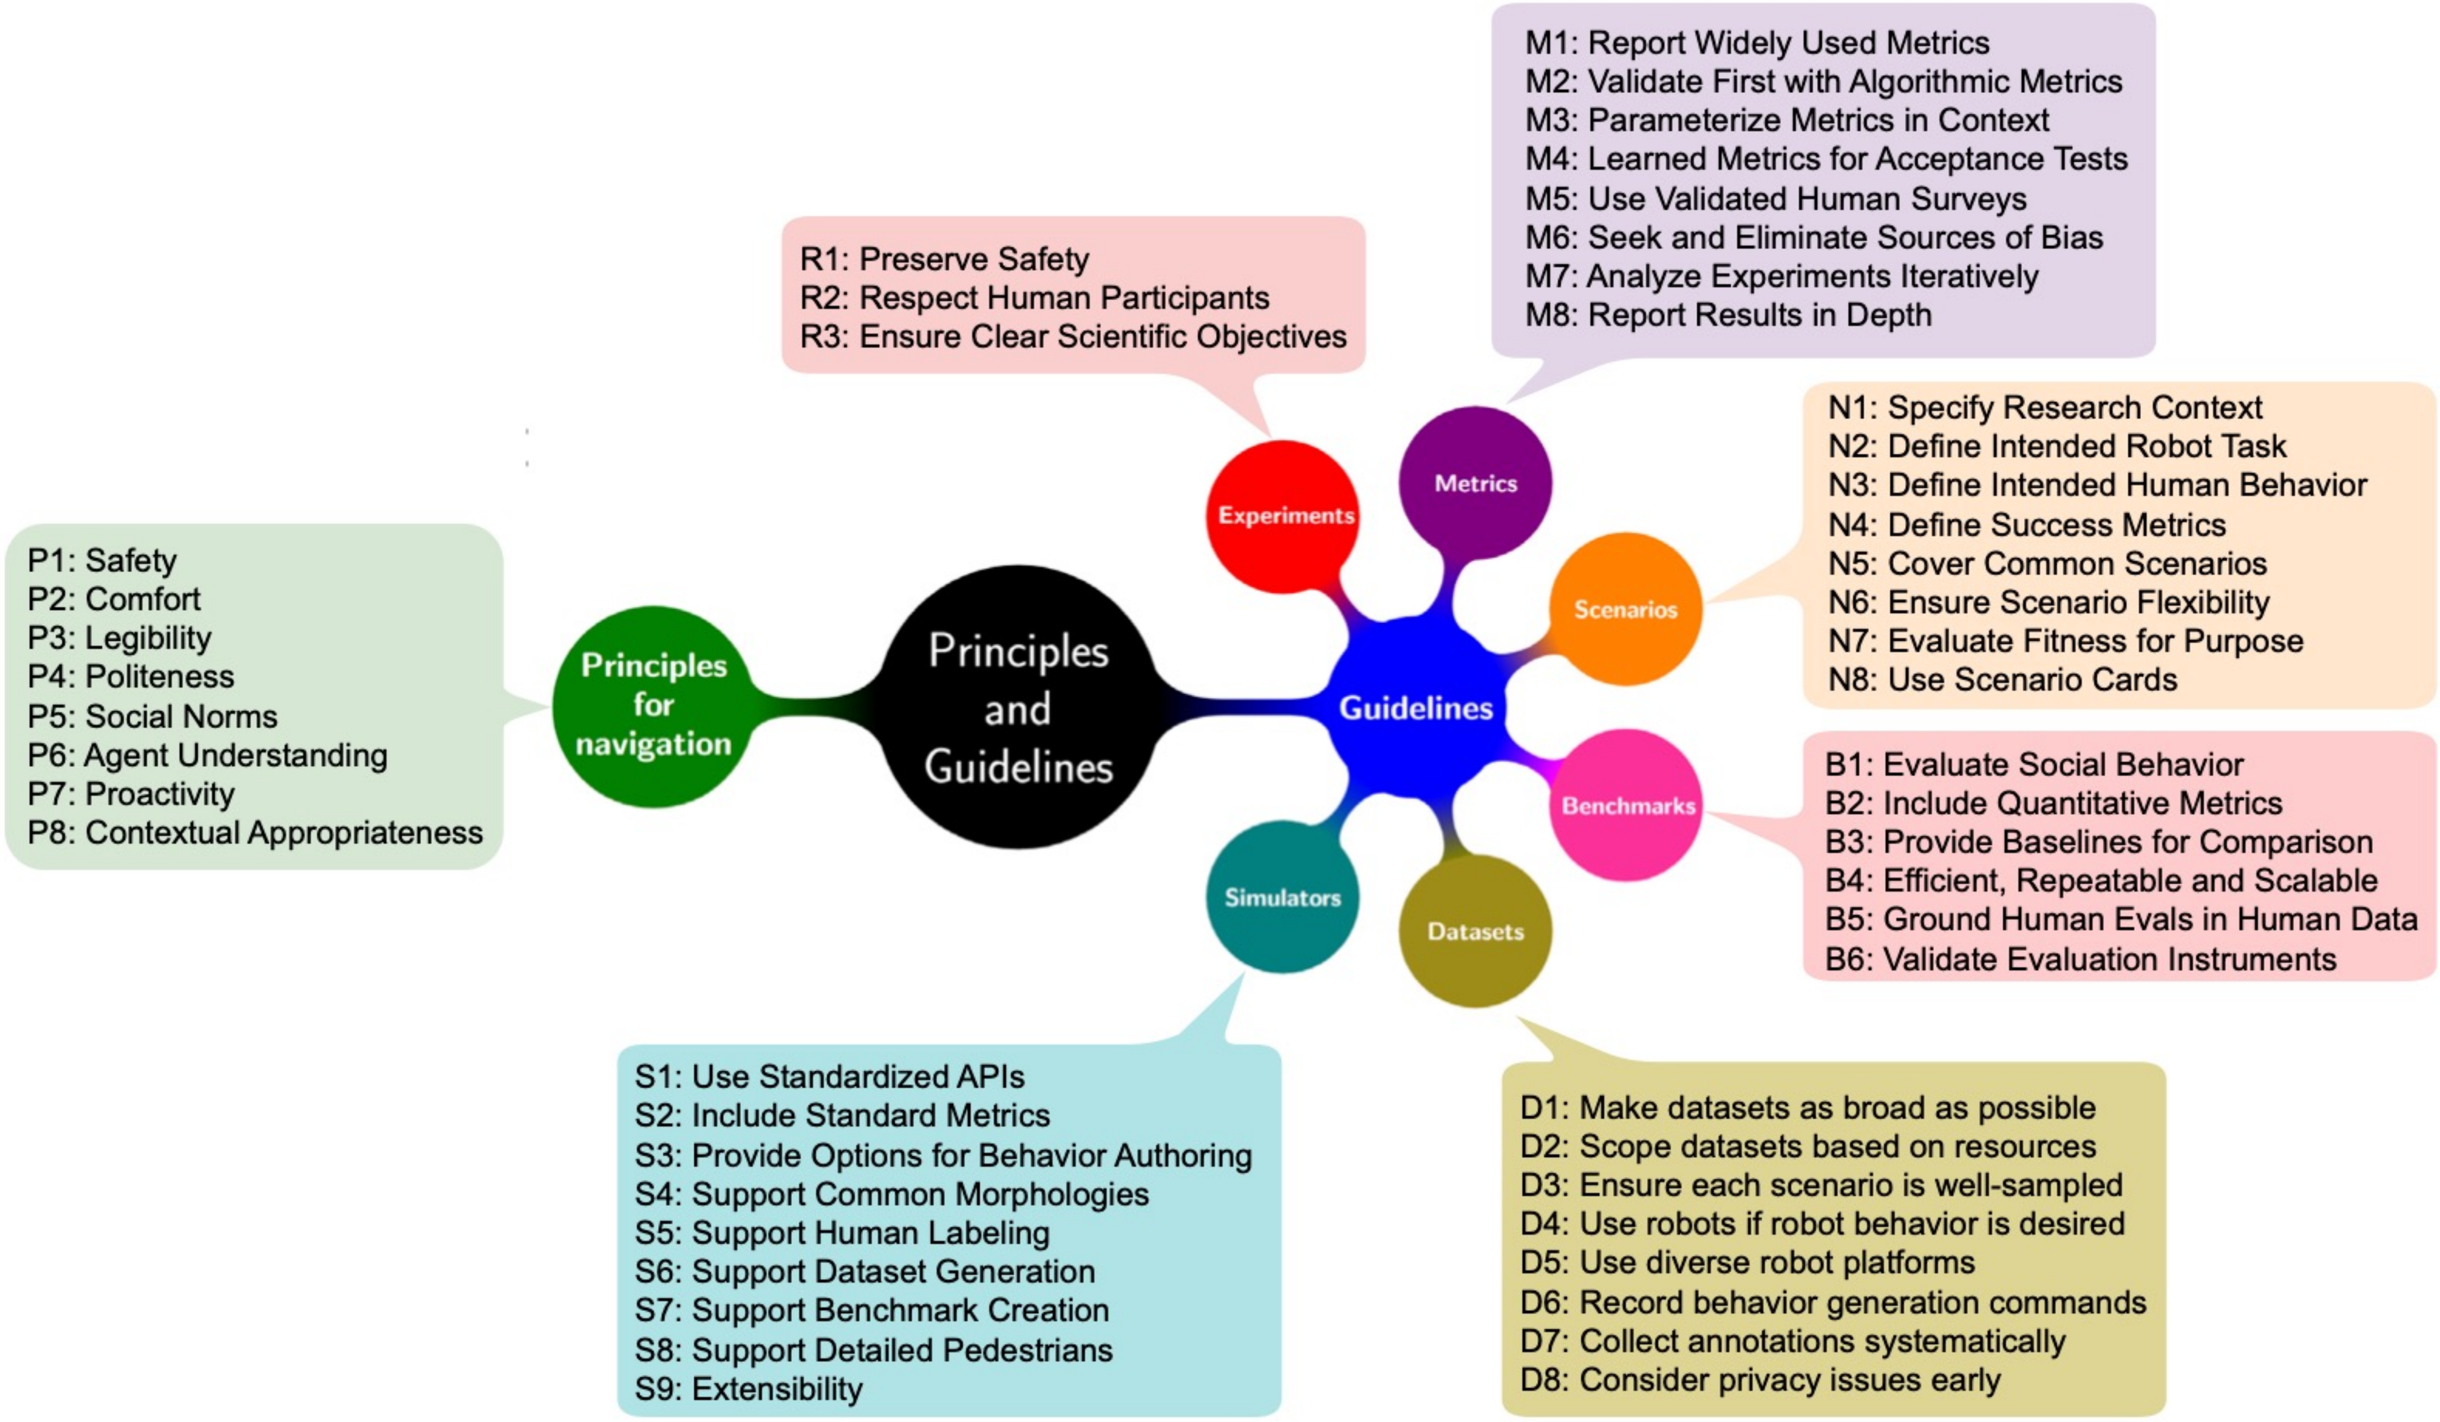
Principles and Guidelines
Principles for navigation
P1: Safety
P2: Comfort
P3: Legibility
P4: Politeness
P5: Social Norms
P6: Agent Understanding
P7: Proactivity
P8: Contextual Appropriateness
Guidelines
Experiments
R1: Preserve Safety
R2: Respect Human Participants
R3: Ensure Clear Scientific Objectives
Metrics
M1: Report Widely Used Metrics
M2: Validate First with Algorithmic Metrics
M3: Parameterize Metrics in Context
M4: Learned Metrics for Acceptance Tests
M5: Use Validated Human Surveys
M6: Seek and Eliminate Sources of Bias
M7: Analyze Experiments Iteratively
M8: Report Results in Depth
Scenarios
N1: Specify Research Context
N2: Define Intended Robot Task
N3: Define Intended Human Behavior
N4: Define Success Metrics
N5: Cover Common Scenarios
N6: Ensure Scenario Flexibility
N7: Evaluate Fitness for Purpose
N8: Use Scenario Cards
Benchmarks
B1: Evaluate Social Behavior
B2: Include Quantitative Metrics
B3: Provide Baselines for Comparison
B4: Efficient, Repeatable and Scalable
B5: Ground Human Evals in Human Data
B6: Validate Evaluation Instruments
Datasets
D1: Make datasets as broad as possible
D2: Scope datasets based on resources
D3: Ensure each scenario is well-sampled
D4: Use robots if robot behavior is desired
D5: Use diverse robot platforms
D6: Record behavior generation commands
D7: Collect annotations systematically
D8: Consider privacy issues early
Simulators
S1: Use Standardized APIs
S2: Include Standard Metrics
S3: Provide Options for Behavior Authoring
S4: Support Common Morphologies
S5: Support Human Labeling
S6: Support Dataset Generation
S7: Support Benchmark Creation
S8: Support Detailed Pedestrians
S9: Extensibility

# Social Robot Navigation is …

- More than Classical Navigation. →**Performer-MPC!** [***XX*** *et al., CoRL22*]

  MPC with Real-Time Transformers.

- Geometric, Semantic, and More. →**Multi-Modal Perception!** [*P, R, N,* ***XX,*** *under review*]

  RGB and Point Cloud for Decision Making.

- Really beyond what Classical Navigation Systems Can Do? →**Targeted Learning!** [*R, …,* ***XX,*** *ICRA24*]

  Learn only When Classical System Fails.

# Performer-MPC: Socially Compliant Navigation Behavior by Real-Time Transformers [*XX* et al., CoRL22]

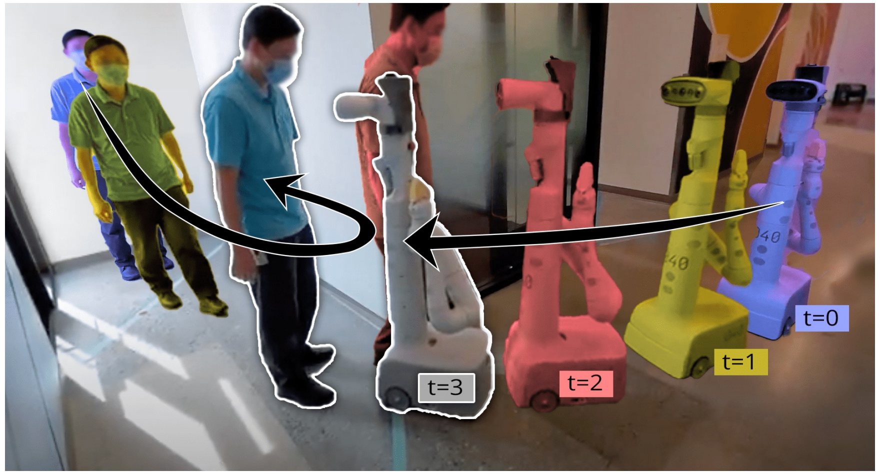


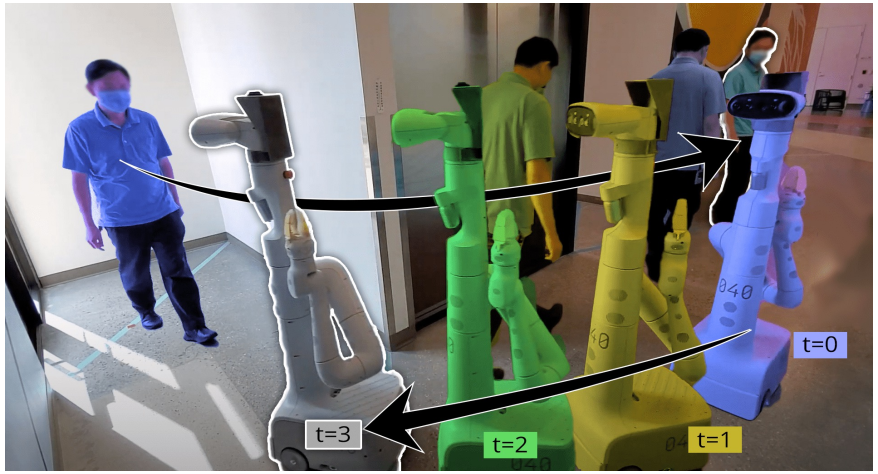


Planning the most efficient, shortest length, minimal time plan?

Social compliance improves motion planning performance!

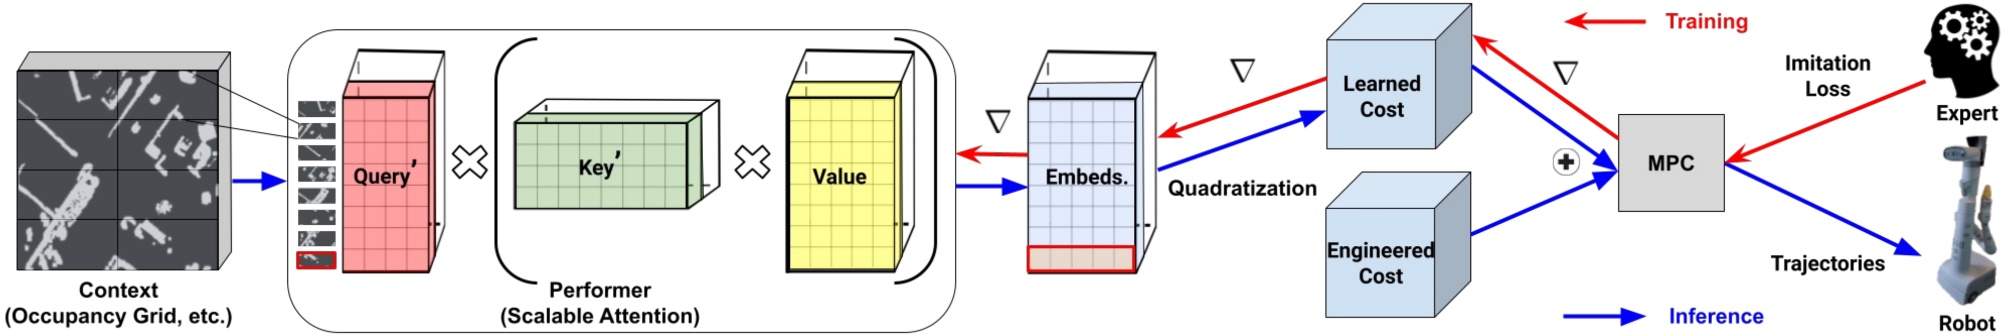

# Performer-MPC: Socially Compliant Navigation Behavior by Real-Time Transformers [*XX et al.*, *CoRL22*]

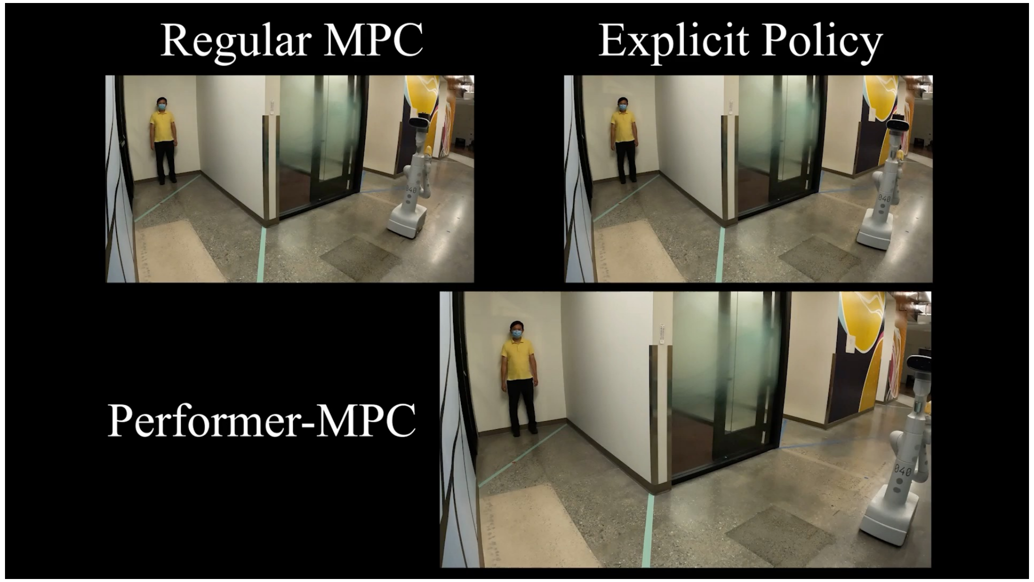


Blind Corner
Learning to anticipate Pedestrians


Pedestrian Obstruction
Learning to respect comfort distance

# Learning Social Robot Navigation with **Multimodal Perception** [P, R, N, ***XX,*** *under review*]

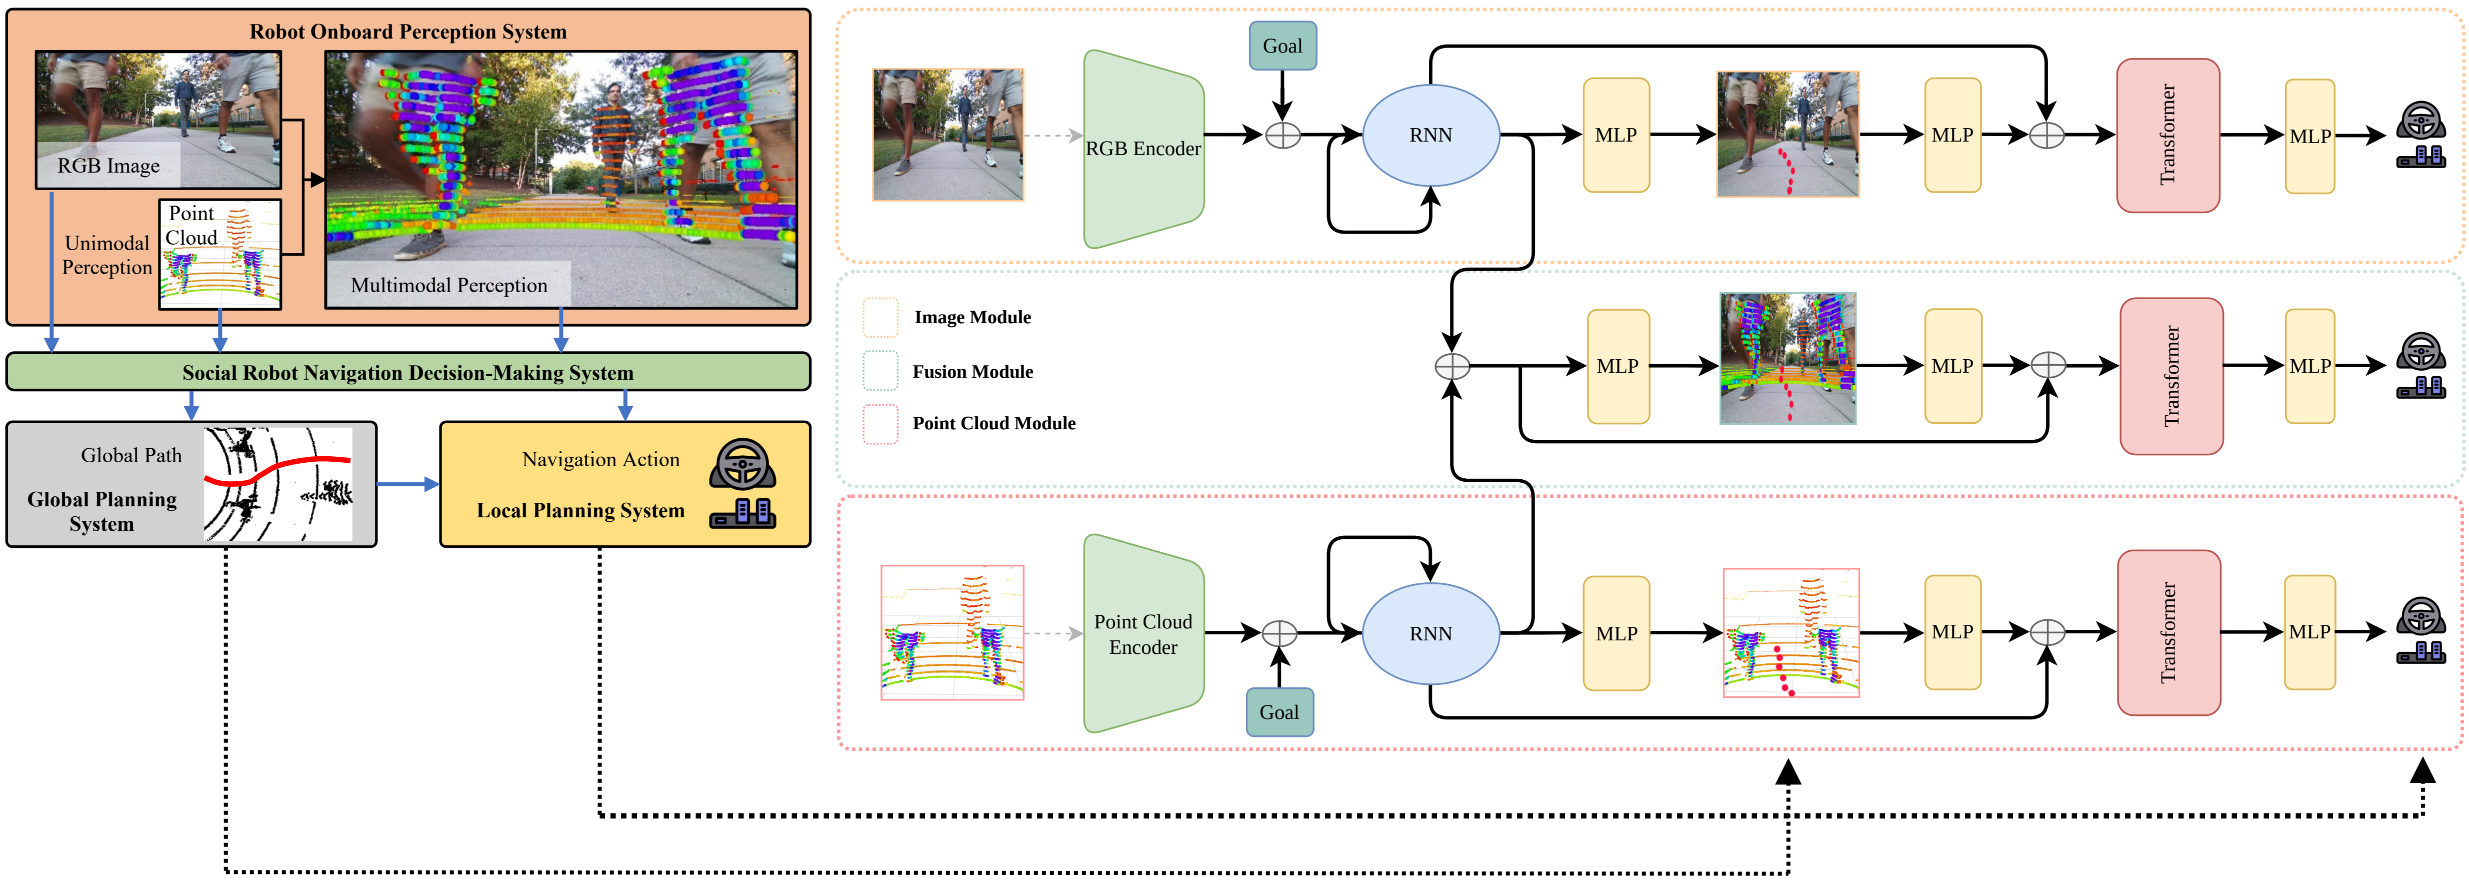

# Learning Social Robot Navigation with **Multimodal Perception** *[P, R, N, **XX,** under review]*

Global Planning

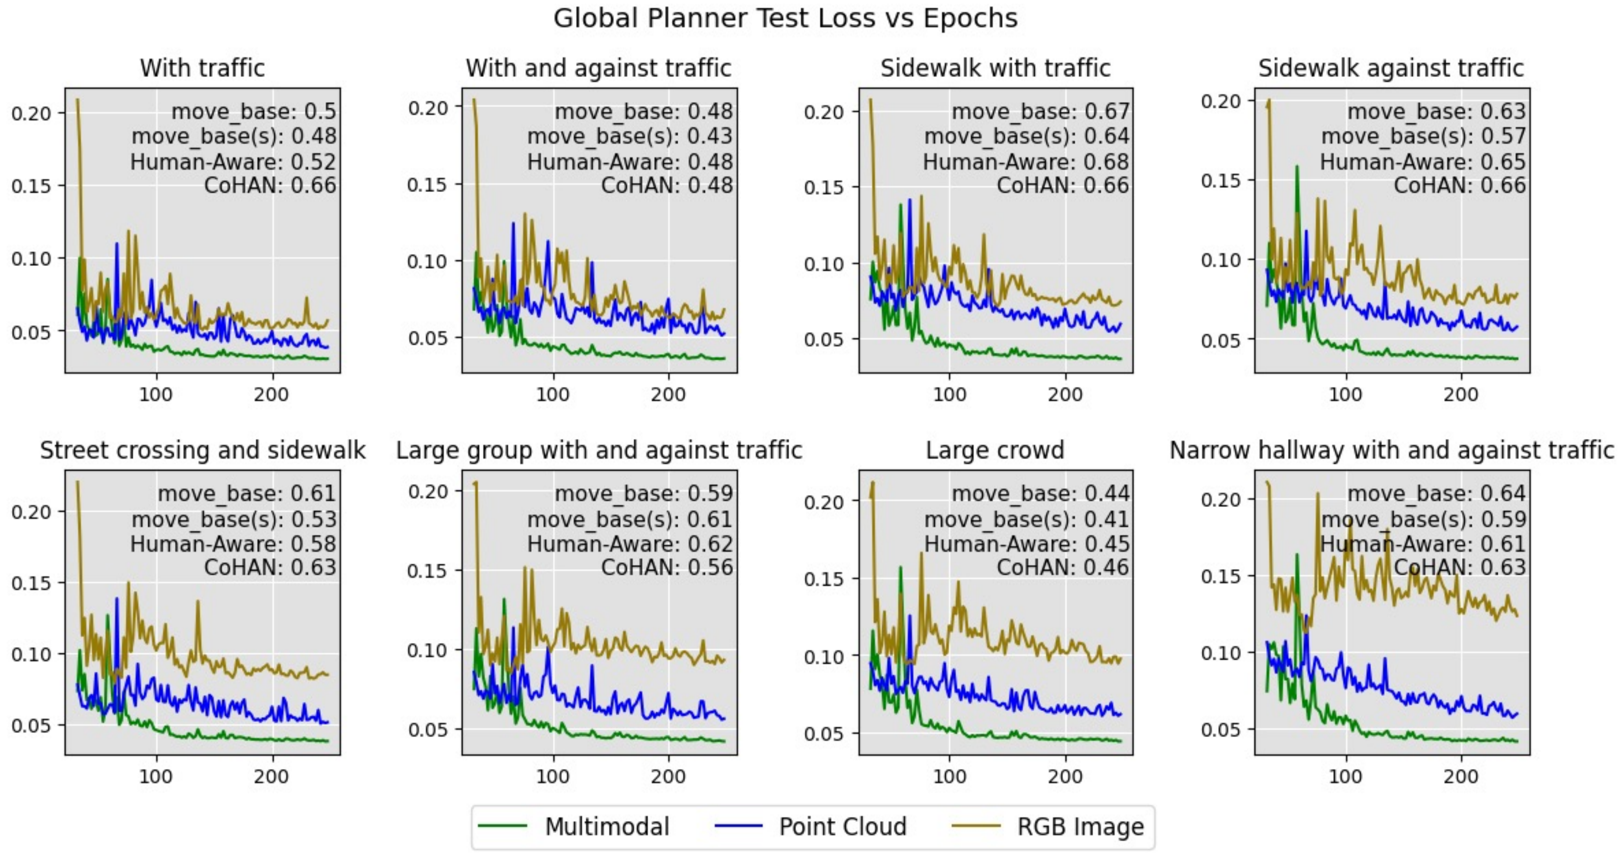


Local Planning

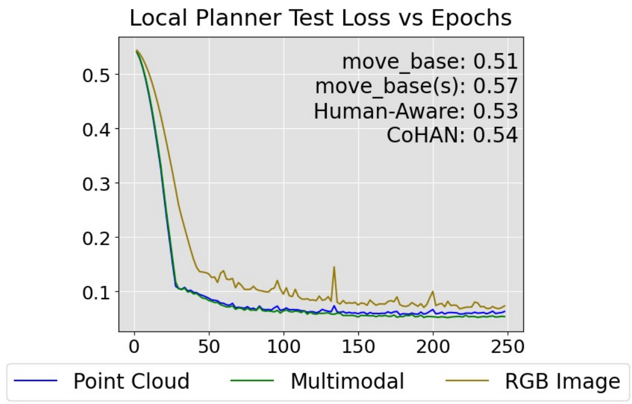


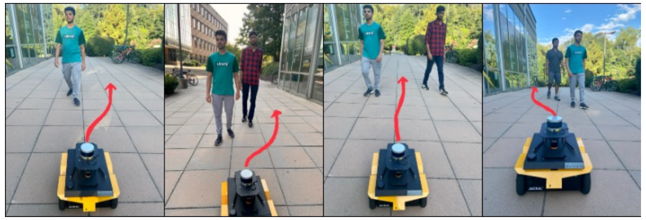

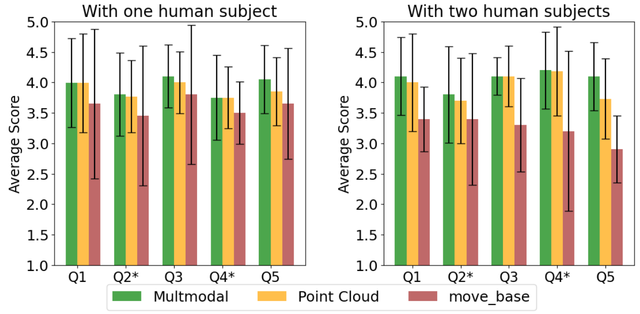


Human Study

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation [*R, …, **XX,** ICRA24*]

- Is social robot navigation really beyond what classical navigation systems can do?
- Probably yes, that’s why we need to study social robot navigation, despite decades of experiences in classical navigation.
- How are they really different?

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …, **XX**, ICRA24]*

- Cutting across a Queue

*[K, N, **XX** et al., RA-L22]*

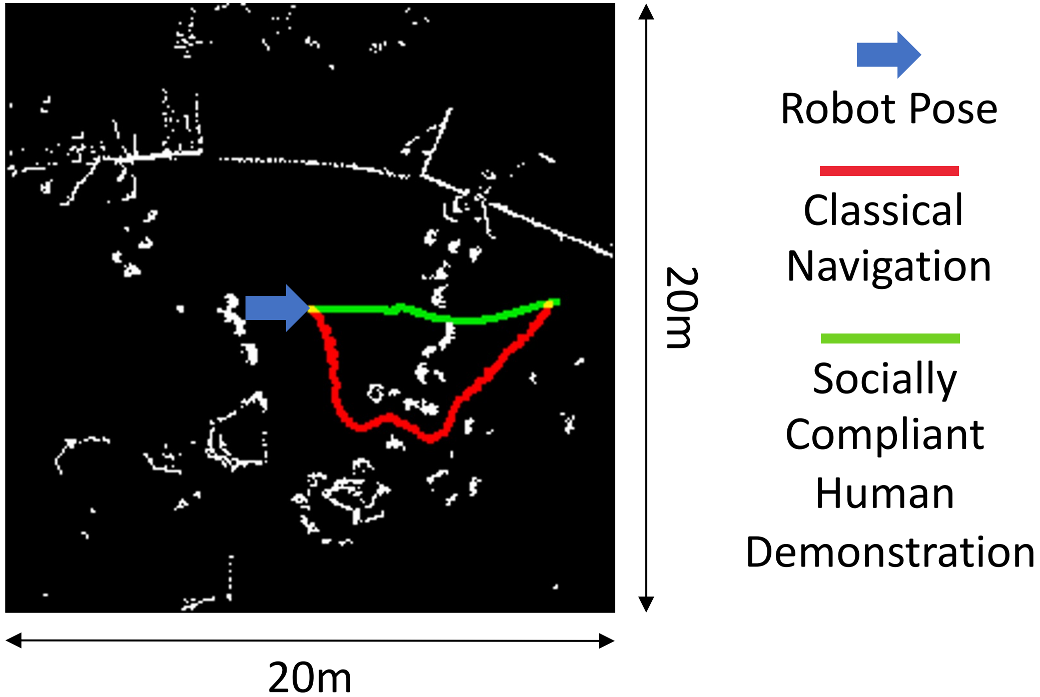

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …, **XX**, ICRA24]*

- Narrow Sidewalk

*[K, N, **XX** et al., RA-L22]*

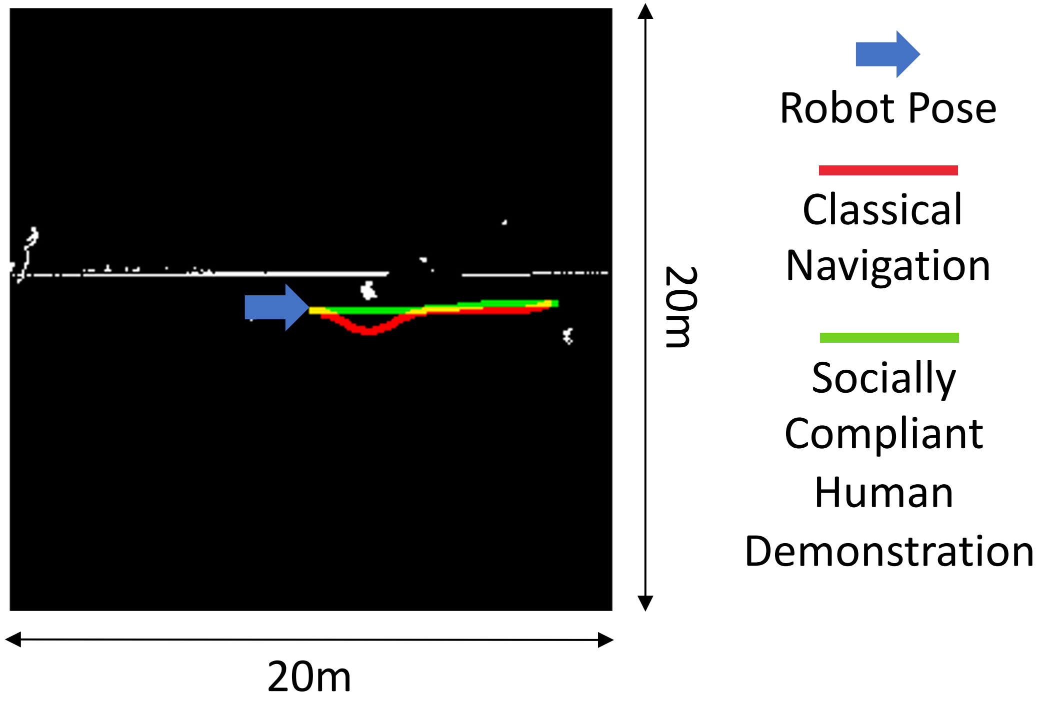

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …,* ***XX,*** *ICRA24]*

- Waiting at a Congested Area

*[K, N,* ***XX*** *et al., RA-L22]*

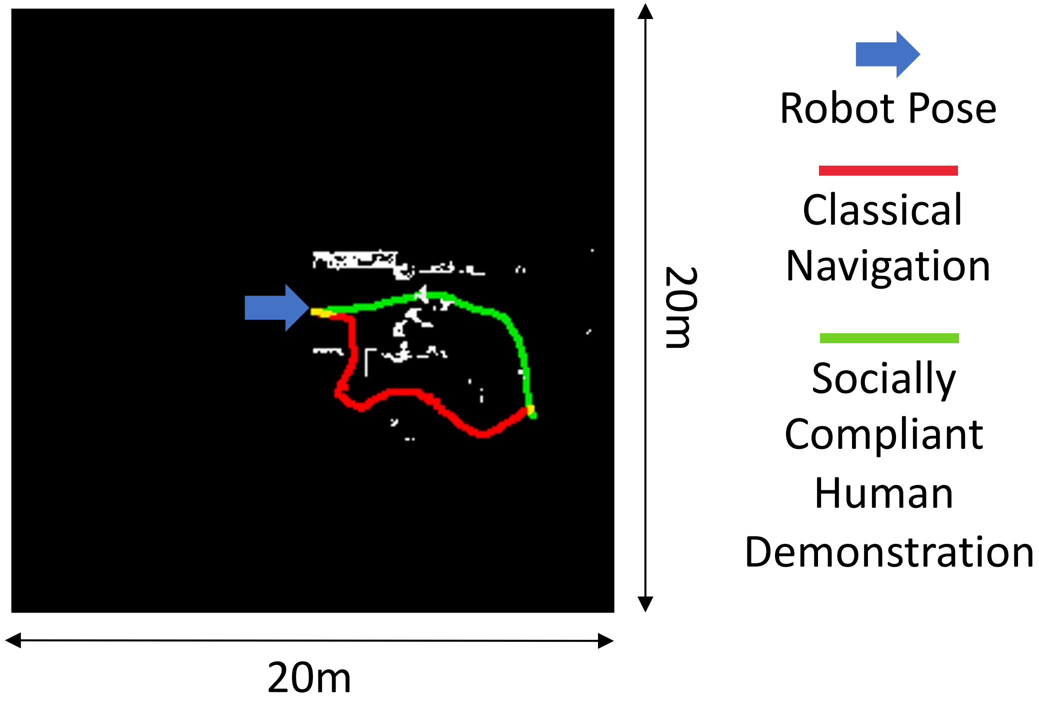

# Targeted Learning: A Hybrid Approach to Social Robot Navigation *[R, …, **XX**, ICRA24]*

- Let's quantify the difference between classical navigation and socially compliant human demonstration!

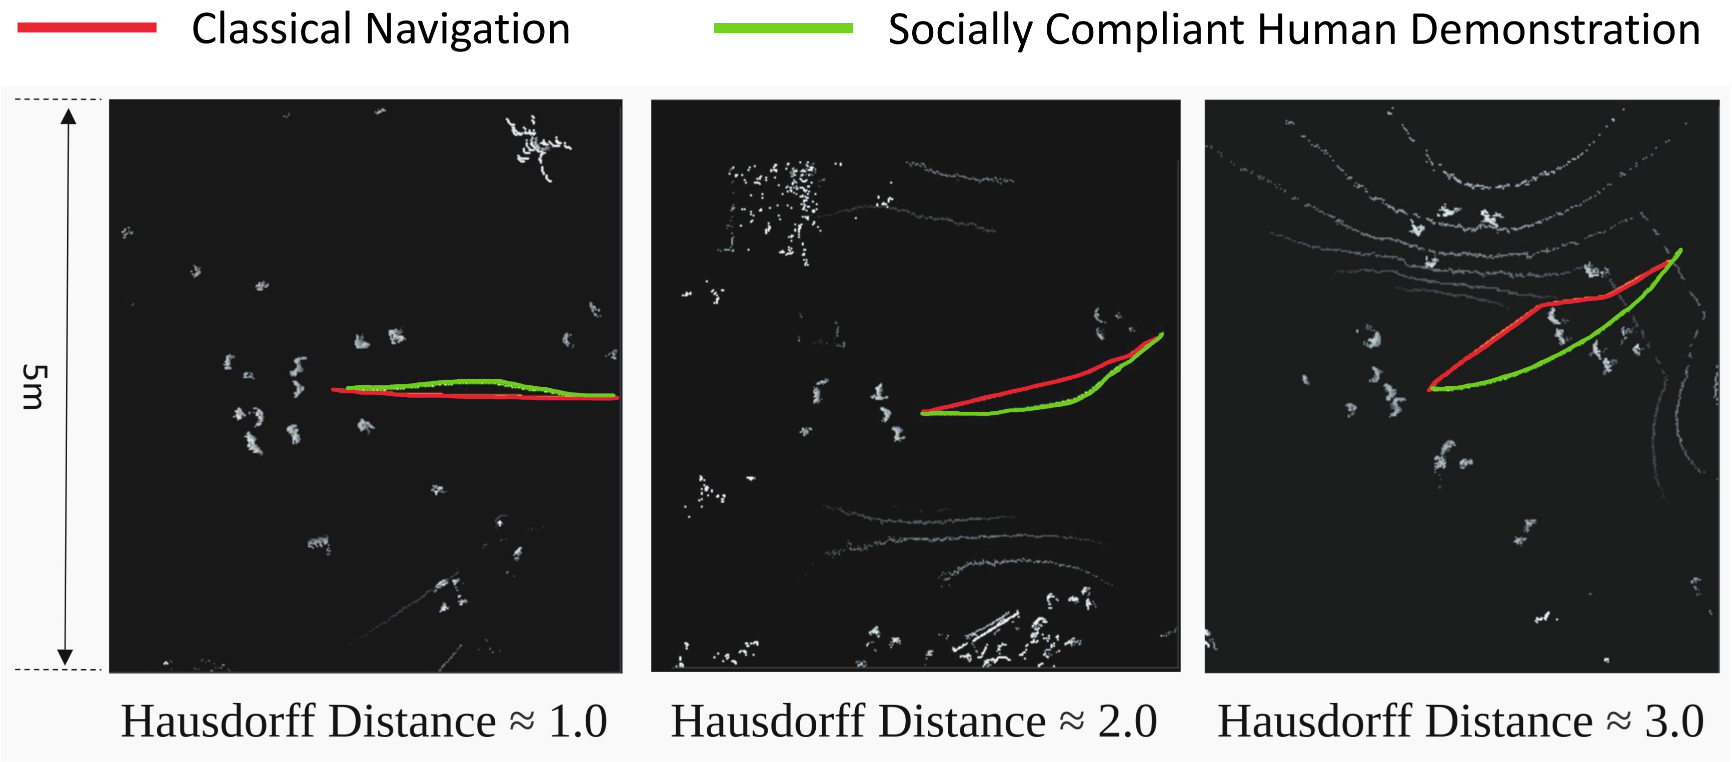

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …,* ***XX,*** *ICRA24]*

- So how does such a difference look like across all SCAND scenarios?

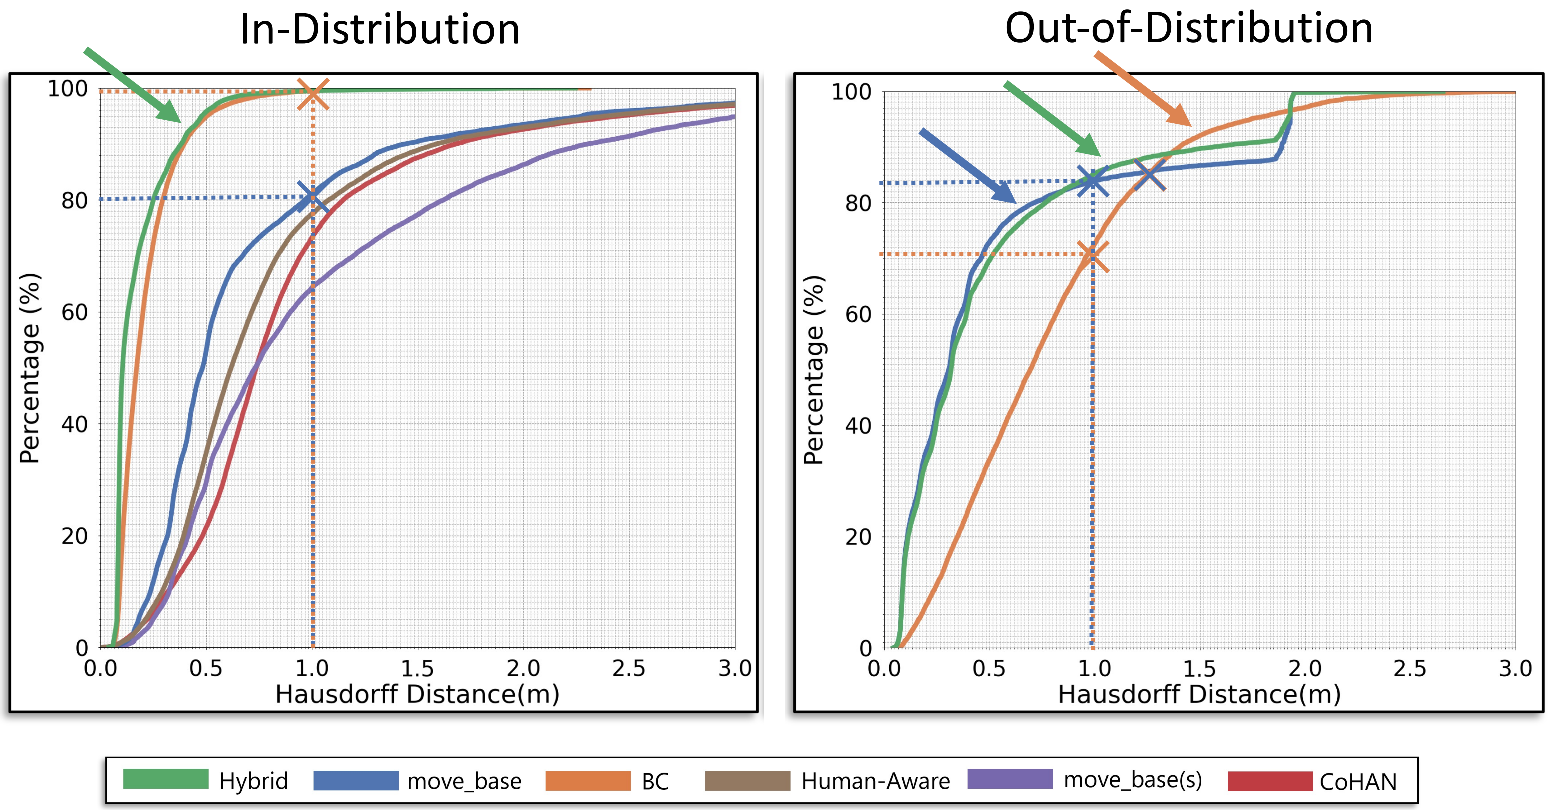

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …,* ***XX,*** *ICRA24]*

- Targeted Learning

```
if (ExpectClassicalGood(s))
    return ClassicalNavigation(s)
else
    return BehaviorCloning(s)
```

Classifier Trained on Labels based on a Hausdorff Distance Threshold

Supervised Learning with only Scenarios where Classical is NOT Good

# **Targeted Learning**: A Hybrid Approach to Social Robot Navigation *[R, …, **XX**, ICRA24]*

**Targeted Learning: A Hybrid Approach to Social Robot Navigation**

Amir Hossain Raj[1*] Zichao Hu[2*] Haresh Karnan[2] Rohan Chandra[2] Amirreza Payandeh[1] Luisa Mao[2]
Peter Stone[2,3] Joydeep Biswas[2] Xuesu Xiao[1]

1 George Mason University 2 Texas The University of Texas at Austin 3 Sony AI

Jackal @ GMU & Spot @ UT Austin
2x Speed

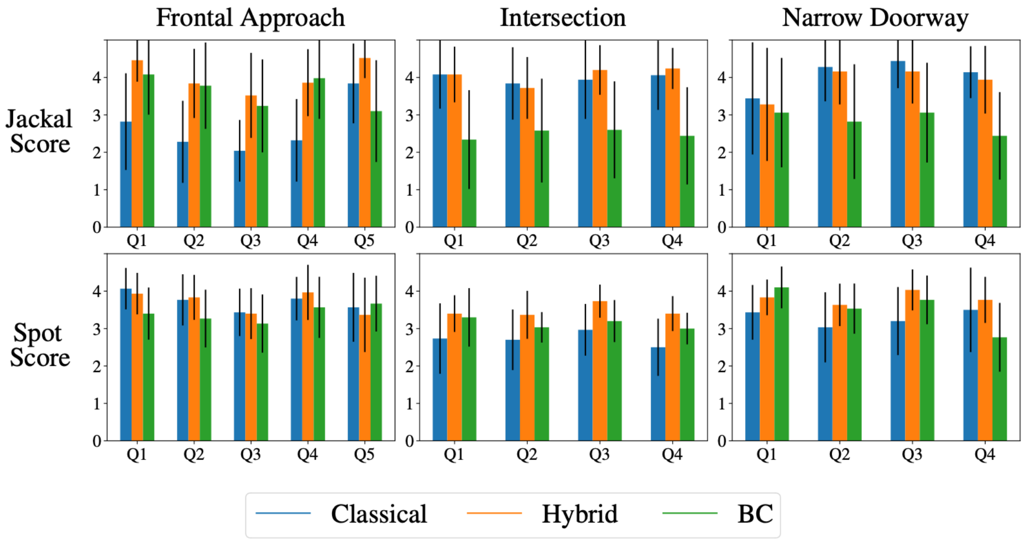


| Jackal | Frontal | Intersection | Doorway |
|---|---|---|---|
| Classical | $2.66 \pm 0.64$ | $3.98 \pm 0.10$ | $\mathbf{4.08 \pm 0.38}$ |
| Hybrid | $\mathbf{4.04 \pm 0.39}$ | $\mathbf{4.06 \pm 0.20}$ | $3.89 \pm 0.36$ |
| BC | $3.63 \pm 0.40$ | $2.49 \pm 0.11$ | $2.84 \pm 0.25$ |
| **Spot** | **Frontal** | **Intersection** | **Doorway** |
| Classical | $\mathbf{3.73 \pm 0.22}$ | $2.72 \pm 0.17$ | $3.29 \pm 0.19$ |
| Hybrid | $3.70 \pm 0.26$ | $\mathbf{3.48 \pm 0.15}$ | $\mathbf{3.82 \pm 0.14}$ |
| BC | $3.41 \pm 0.19$ | $3.13 \pm 0.12$ | $3.54 \pm 0.49$ |

# Social Robot Navigation is …

- More than Classical Navigation. →**Performer-MPC!** [***XX*** *et al., CoRL22*]
  MPC with Real-Time Transformers.
- Geometric, Semantic, and More. →**Multi-Modal Perception!** [*P, R, N,* ***XX,*** *under review*]
  RGB and Point Cloud for Decision Making.
- Really beyond what Classical Navigation Systems Can Do? →**Targeted Learning!** [*R, …,* ***XX,*** *ICRA24*]
  Learn only When Classical System Fails.
- Datasets, Protocols, Principles, and Guidelines!

# Human-Interactive Mobile Robots: from Learning to Deployment

- This talk: Human-interactive mobile robots that efficiently learn from and harmoniously deploy among humans
  - Adaptive Planner Parameter Learning (APPL) to utilize the opportunities from easily available non-expert human interactions.
  - Datasets, protocols, principles, guidelines, and learning methods to address social robot navigation challenges.

## CRASAR

Center for Robot-Assisted Search and Rescue


Robin Murphy


Yiming Fan


M. Suhail


Jan Dufek


T. Woodbury

## Biorobotics Laboratory

THE ROBOTICS INSTITUTE


Howie Choset


E. Cappo


C. Gong


Ke Sun


W. Zhen


Jin Dai


M. Traverse

## LARG

Learning Agents Research Group

The University of Texas at Austin


Peter Stone


G. Warnell

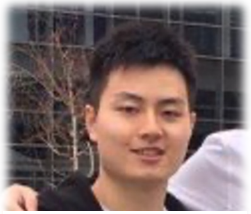
Bo Liu

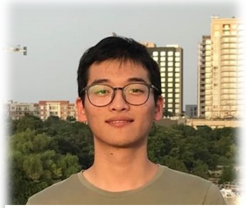
Zifan Xu

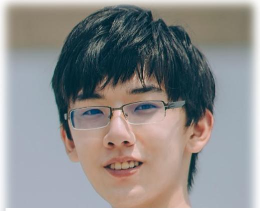
Zizhao Wang


Haresh Karnan


Rohan Chandra

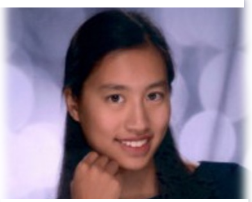
Abigail Truong


Daniel Perille

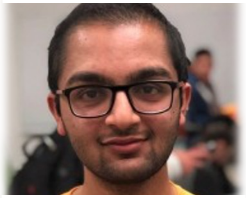
G. Dhamankar


Anirudh Nair


Zichao Hu


Luisa Mao


Joydeep Biswas

AMRL

Everyday Robots

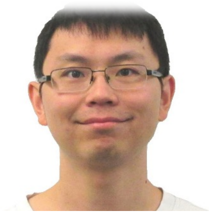
Tingnan Zhang


K. Choromanski


Edward Lee


Anthony Francis


Jake Varley


Stepehn Tu


Sumeet Singh

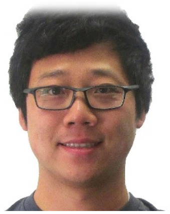
Peng Xu


Fei Xia


Mikael Persson


D. Kalashnikov


Leila Takayama


Roy Frostig


Jie Tan


Carolina Parada


Vikas Sindhwani

Google Research

RobotiX

GEORGE MASON UNIVERSITY | College of Engineering and Computing


Aniket Datar


Chenhui Pan


Aaron Nguyen


Mohammad Nazeri


Amir Payandeh


Amir Raj


Bhabaranjan Panigrahi

# Human-Interactive Mobile Robots: from Learning to Deployment

- This talk: Human-interactive mobile robots that efficiently learn from and harmoniously deploy among humans
  - Adaptive Planner Parameter Learning (APPL) to utilize the opportunities from easily available non-expert human interactions.
  - Datasets, protocols, principles, guidelines, and learning methods to address social robot navigation challenges.