APPLD: Adaptive Planner Parameter Learning From Demonstration

Xuesu Xiao[1*], Bo Liu[1*], Garrett Warnell[2], Jonathan Fink[2], and Peter Stone[1]

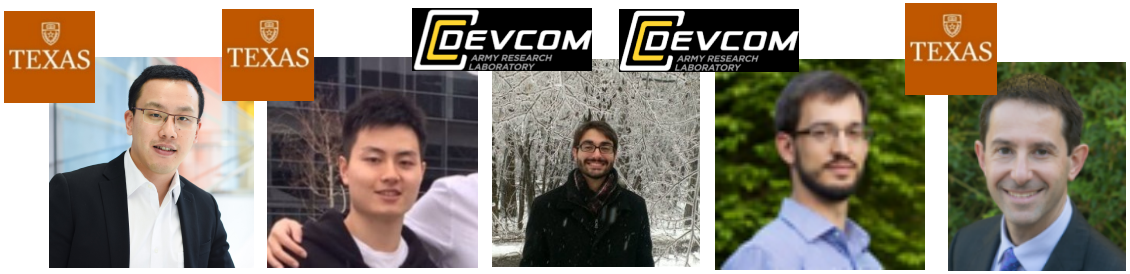

[1]The University of Texas at Austin

[2]Army Research Laboratory

[*]Equally Contributing Authors

LfD for Improved Robot Navigation

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

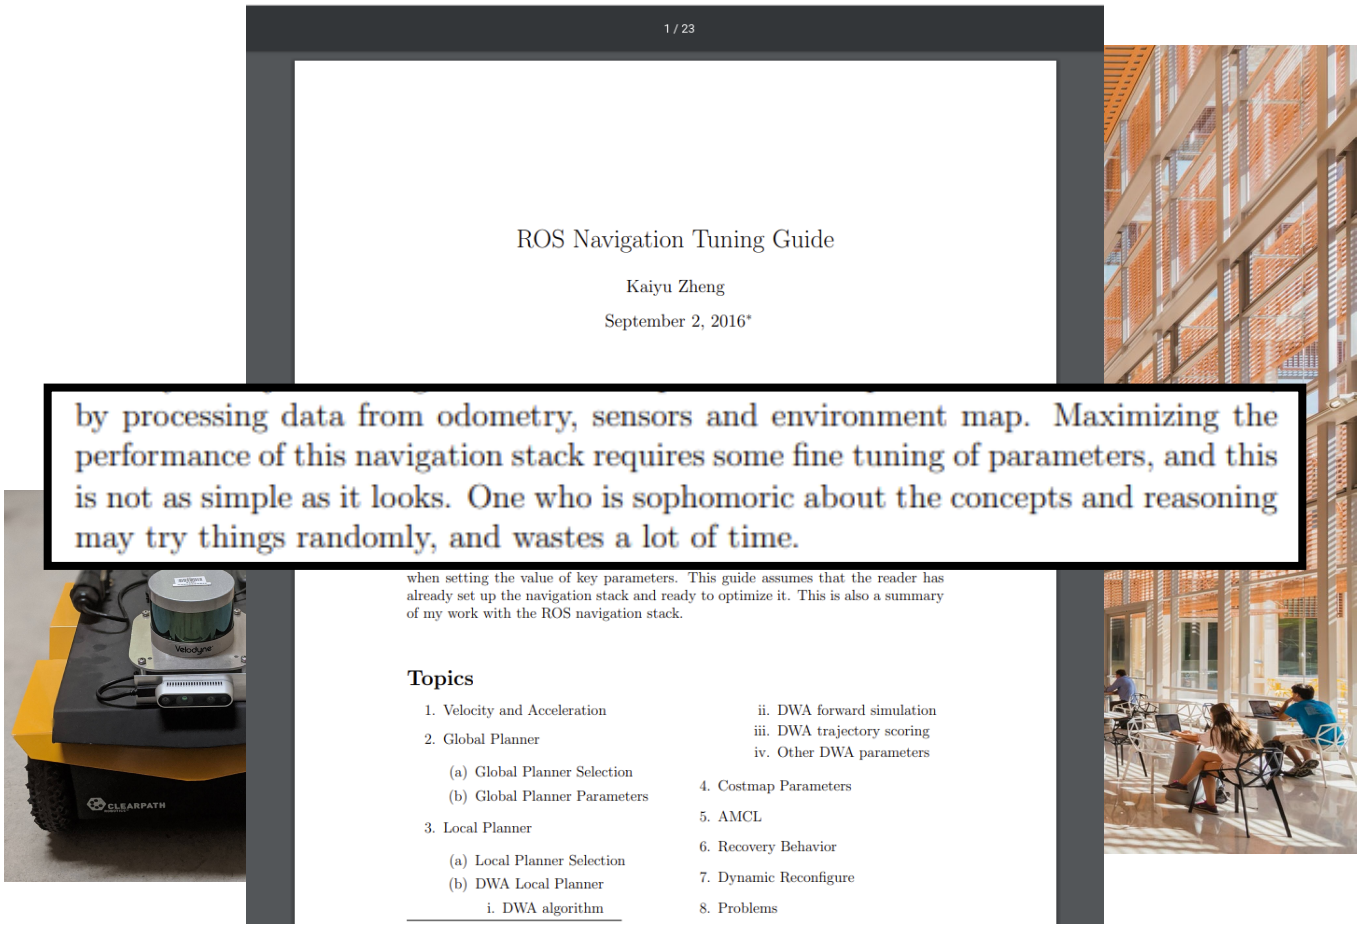

LfD for Improved Robot Navigation

Inspiration: Humans do this effortlessly.

LfD for Improved Robot Navigation

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

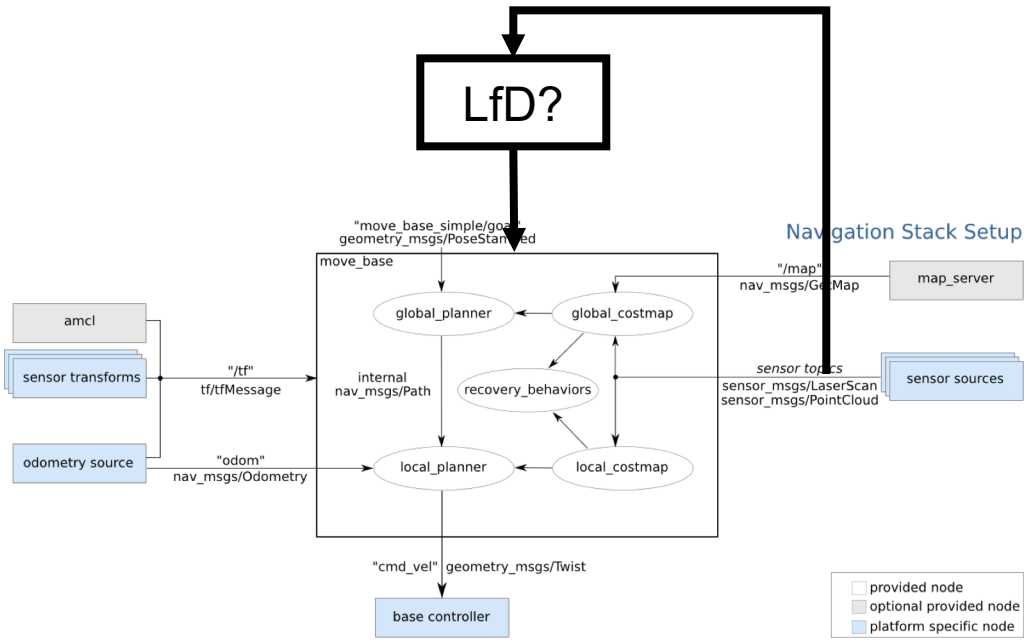

ROS **`move_base`** navigation stack

LfD for Improved Robot Navigation

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

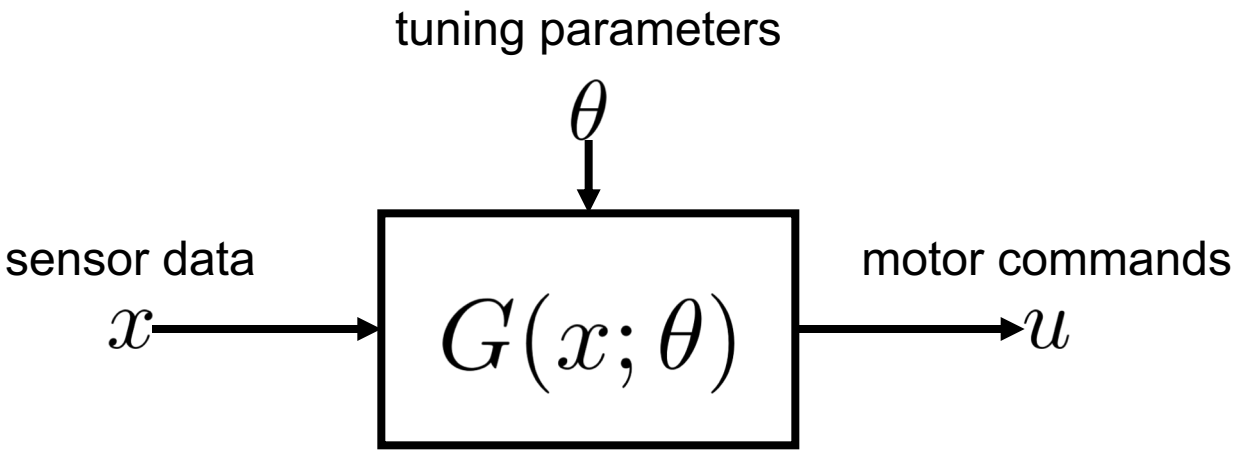

Behavioral Cloning: Learn parameters from a demonstration using supervised learning.

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

LfD for Improved Robot Navigation

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

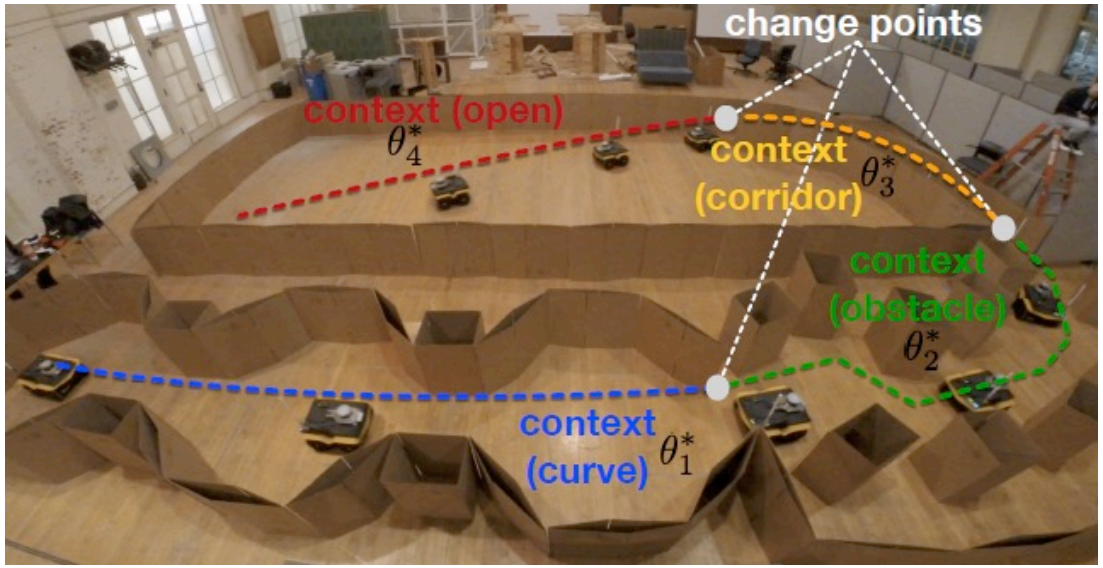

$$\{\theta_1^*, \ldots, \theta_K^*\}$$

LfD for Improved Robot Navigation

APPLD: Adaptive planner parameter learning from demonstration

1. Collect demonstration.

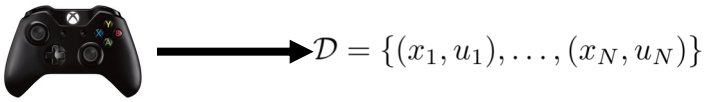

2. Perform automatic demonstration segmentation.

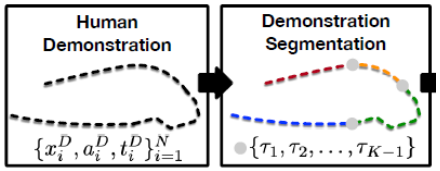

3. Use black-box optimization to find set of optimal parameters.

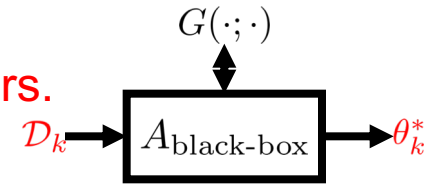

4. Use supervised learning to train a context predictor.

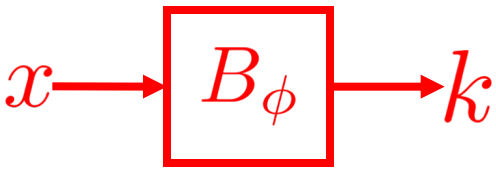

LfD for Improved Robot Navigation

APPLD Deployment

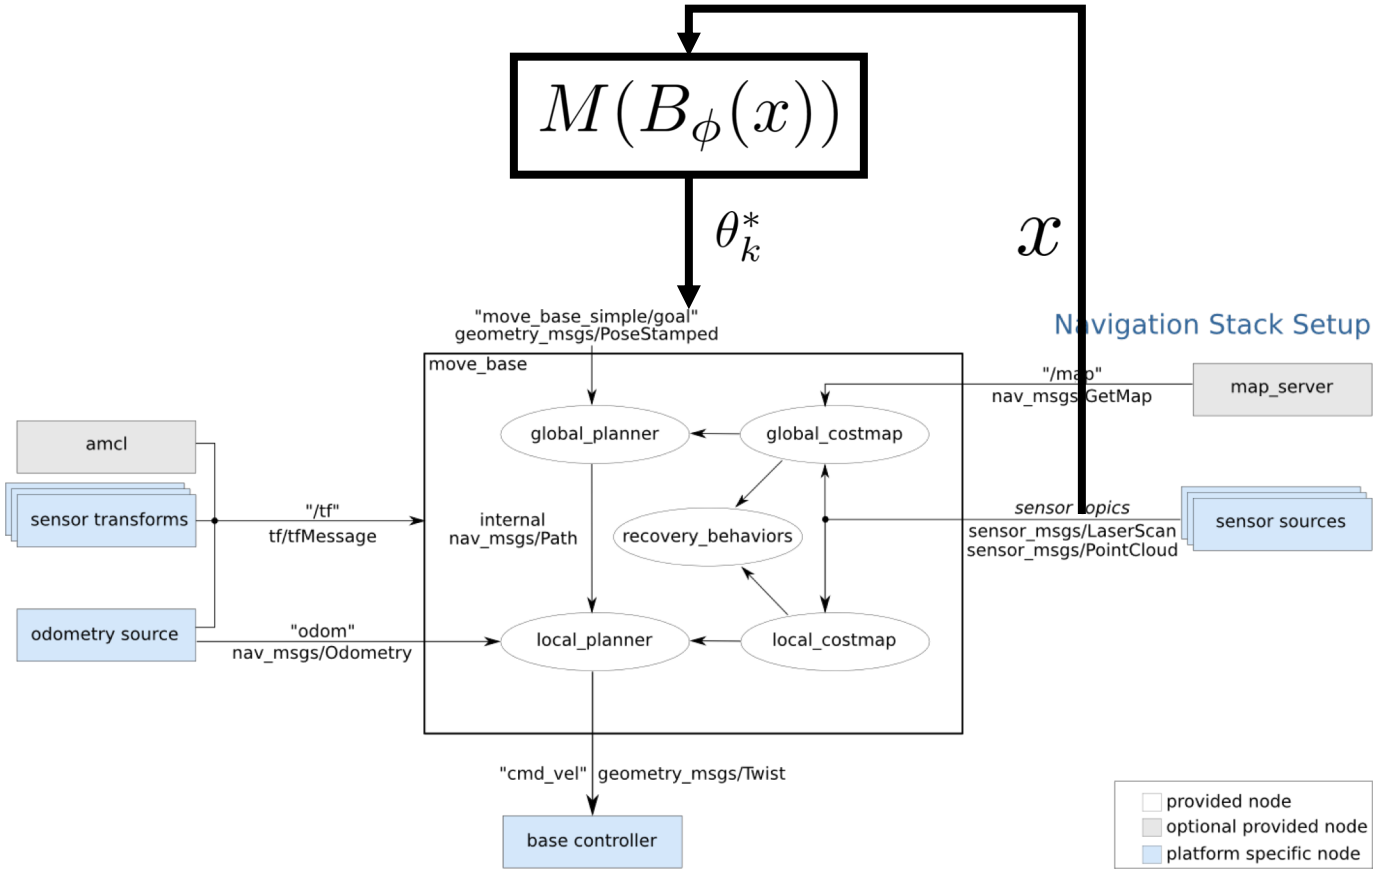

LfD for Improved Robot Navigation

Experiments

Robot: Clearpath Jackal (Velodyne Puck lidar)

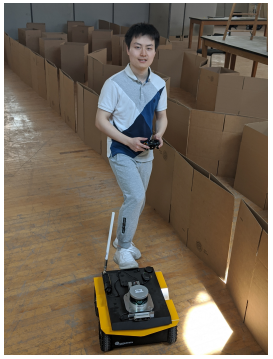

Human: An author (Xbox wireless controller)

Environment: Challenging obstacle course

Results

APPLD: Deployment in demonstration environment

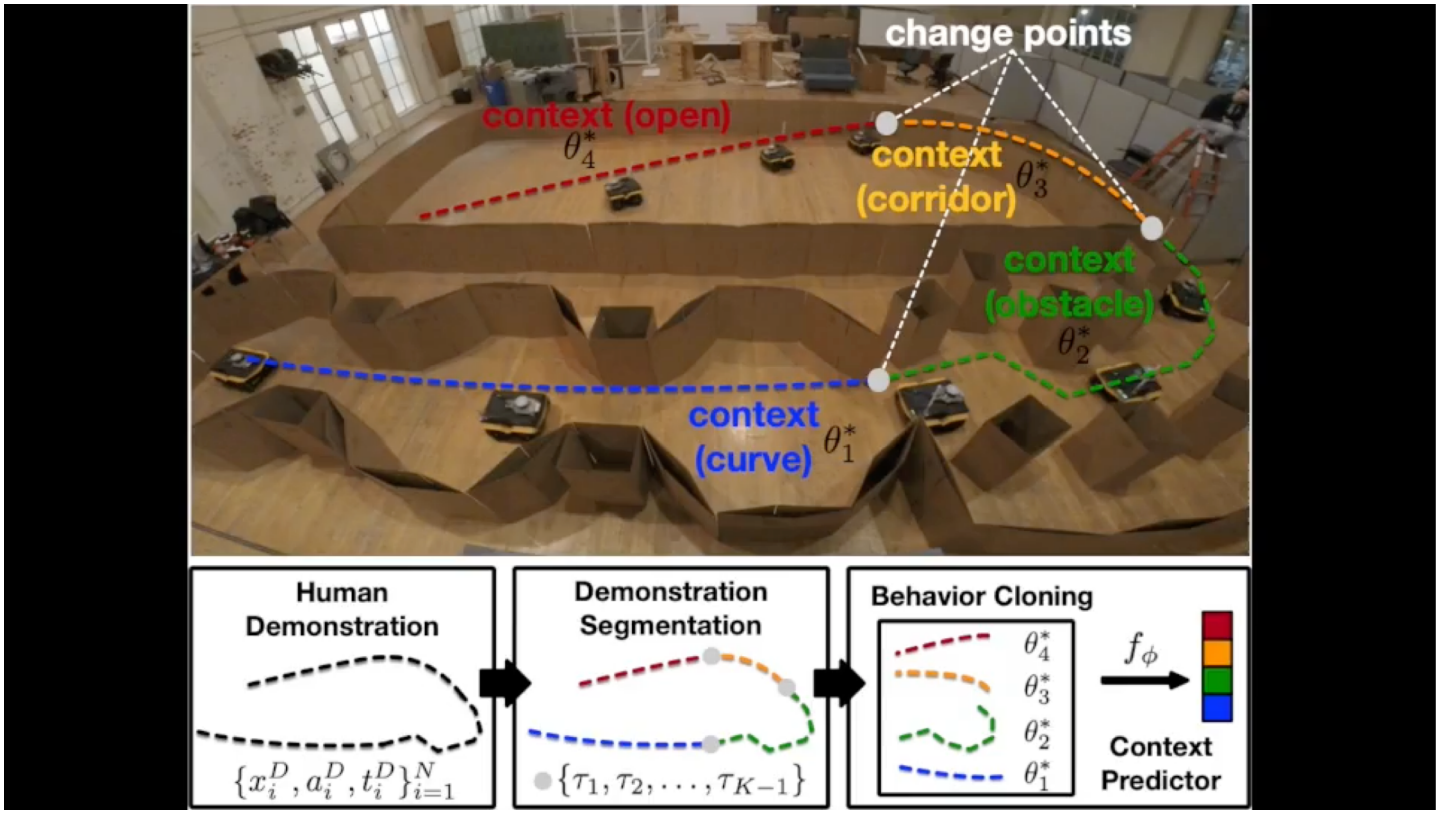

Results

Xuesu Xiao
xiao@cs.utexas.edu

Different robot, navigation stack, and environment