

# Navigation

Saurabh Gupta

# Tasks

- Locomotion
  - Wheeled ground vehicle
  - Drone
  - Quadruped
  - Off-road terrain



# Tasks

- Locomotion
  - Platforms:
    - Wheeled ground vehicle, Drone, Quadrupeds, Off-road terrain
  - Move safely, and gracefully. Avoid collisions.
  - Follow high-level commands
  - Visual Servoing, follow an object
- Goal-directed behavior
  - Goals known / unknown
  - Environment seen / novel

# Axes of variation

## Supervision Signal



**Expert Trajectories**

**Reward Signals**

**Ecological Supervision**

**No supervision (SIFT + Geometry)**

## Environment Modeling



**No Modeling (Reactive Policies)**

**Implicit Modeling (via LSTMs)**

**Explicit modeling of space (metric / topological maps)**

## Task



**Locomotion (move safely)**

**Path following**

**Long range goals (points / semantics / visual)**

**Navigation for actions**

## Generalization



**Extensive training on test environments**

**Drop in novel environment**

# Axes of variation

**Abstraction**



State space  
discretization  
(eg. grid world)

Discrete  
Actions

Continuous  
actions

**Environment Properties**



Random Mazes  
(limited semantics)

Real world  
environments

**Experimental setup**



Simulation

Simulation (based of real  
world environments)

Real world

**Training**

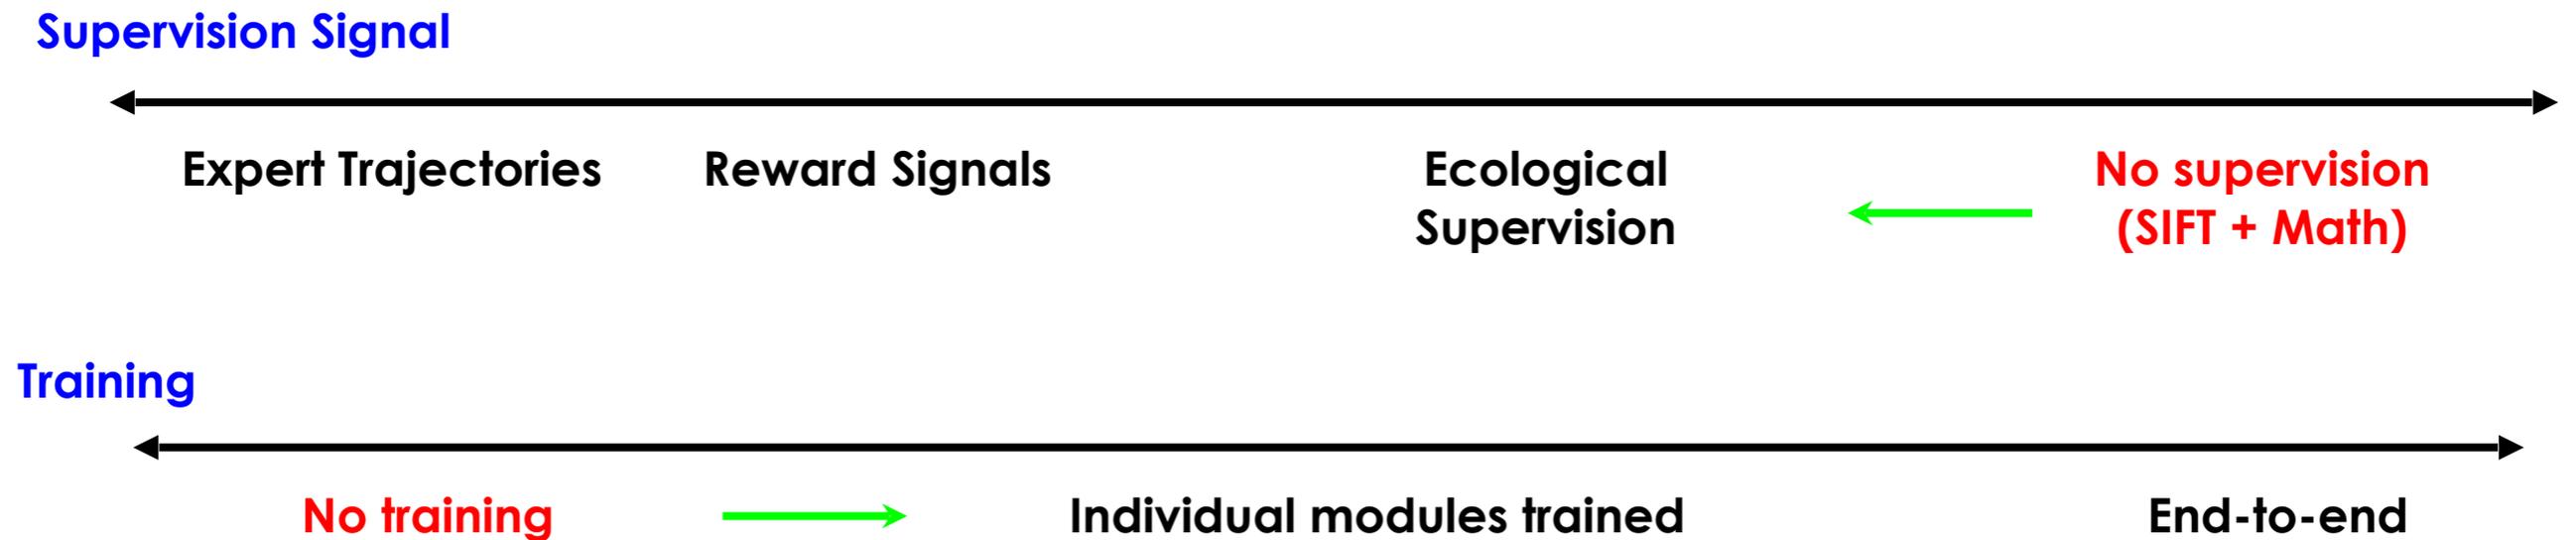


No training

Individual modules trained

End-to-end

# CNNs for SLAM

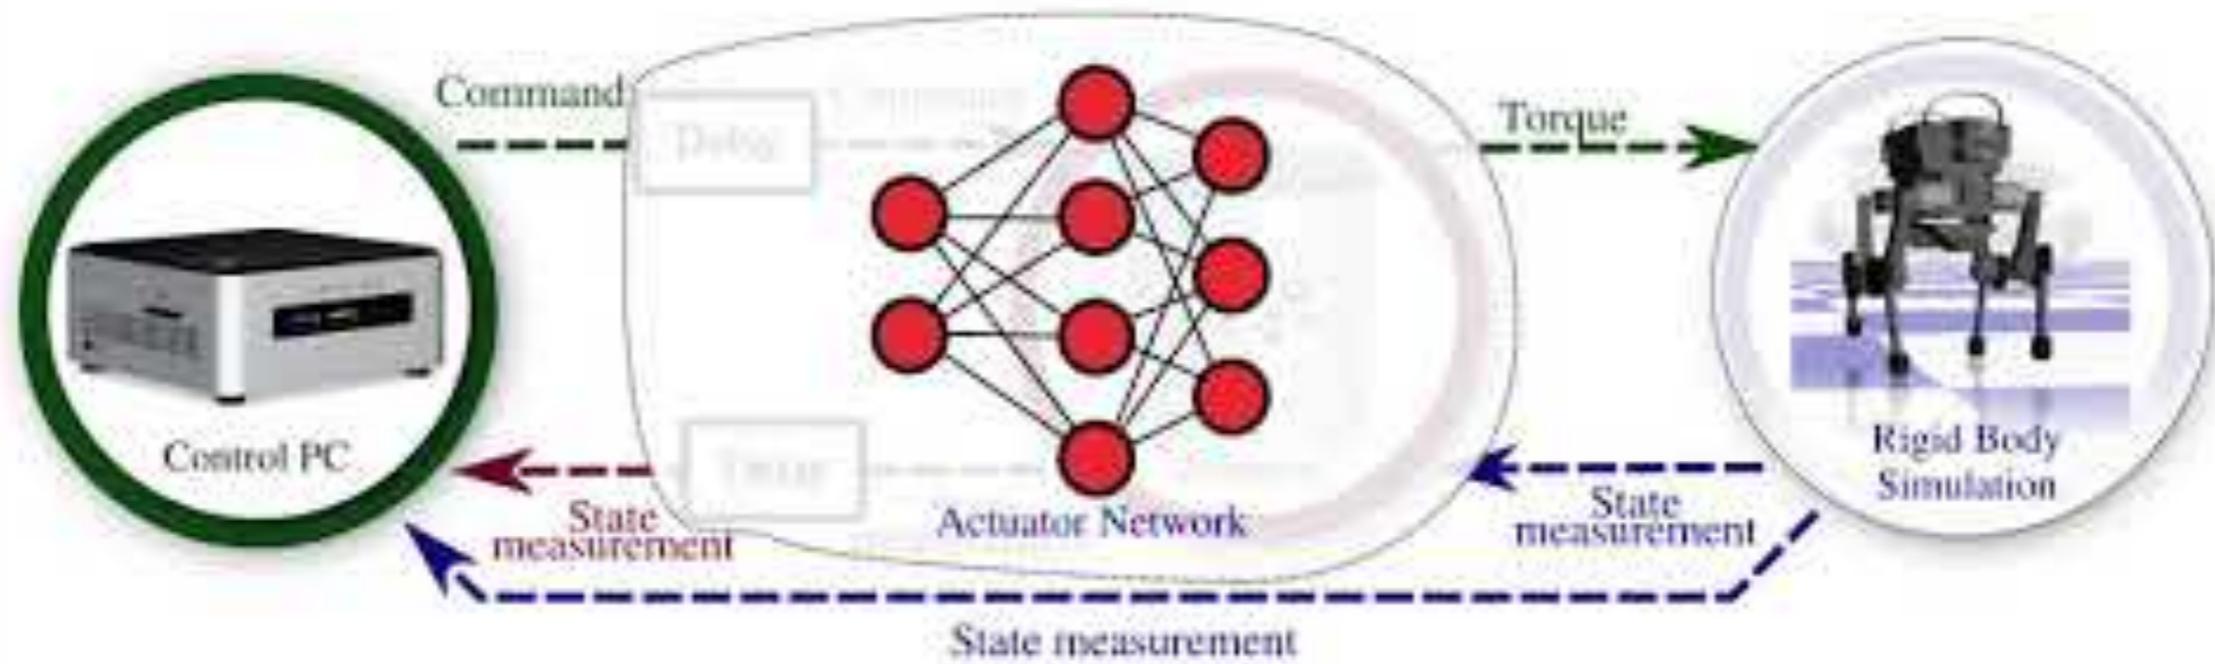


## Classical SLAM

- **Localization:**
  - **Semantic Visual Localization** J. Schonberger, M. Pollefeys, A. Geiger, T. Sattler arXiv 17 [pdf](#)
  - **Self-Supervised Place Recognition in Mobile Robots** S. Pillai, J. Leonard. IROS17 [pdf](#)
  - **Active neural localization** D. Chaplot arXiv17 [pdf](#)
- **Visual Odometry**
  - **Towards Visual Ego-motion Learning in Robots.** S. Pillai and J. Leonard IROS17 [pdf](#)
  - **Unsupervised Learning of Depth and Ego-motion from Video** Zhou et al. CVPR17 [pdf](#)
- **Reconstruction**
  - **Learned Stereo Machines.** Kar et al. NIPS17 [pdf](#)
  - **Differentiable ray consistency.** Tulsiani et al. CVPR17 [pdf](#)



- Learning Agile and Dynamic Motor Skills for Legged Robots J. Hwango et al. Science Robotics 18 [pdf](#), [video](#)



To this end, we train a neural network representing this complex dynamics with data from the real robot.

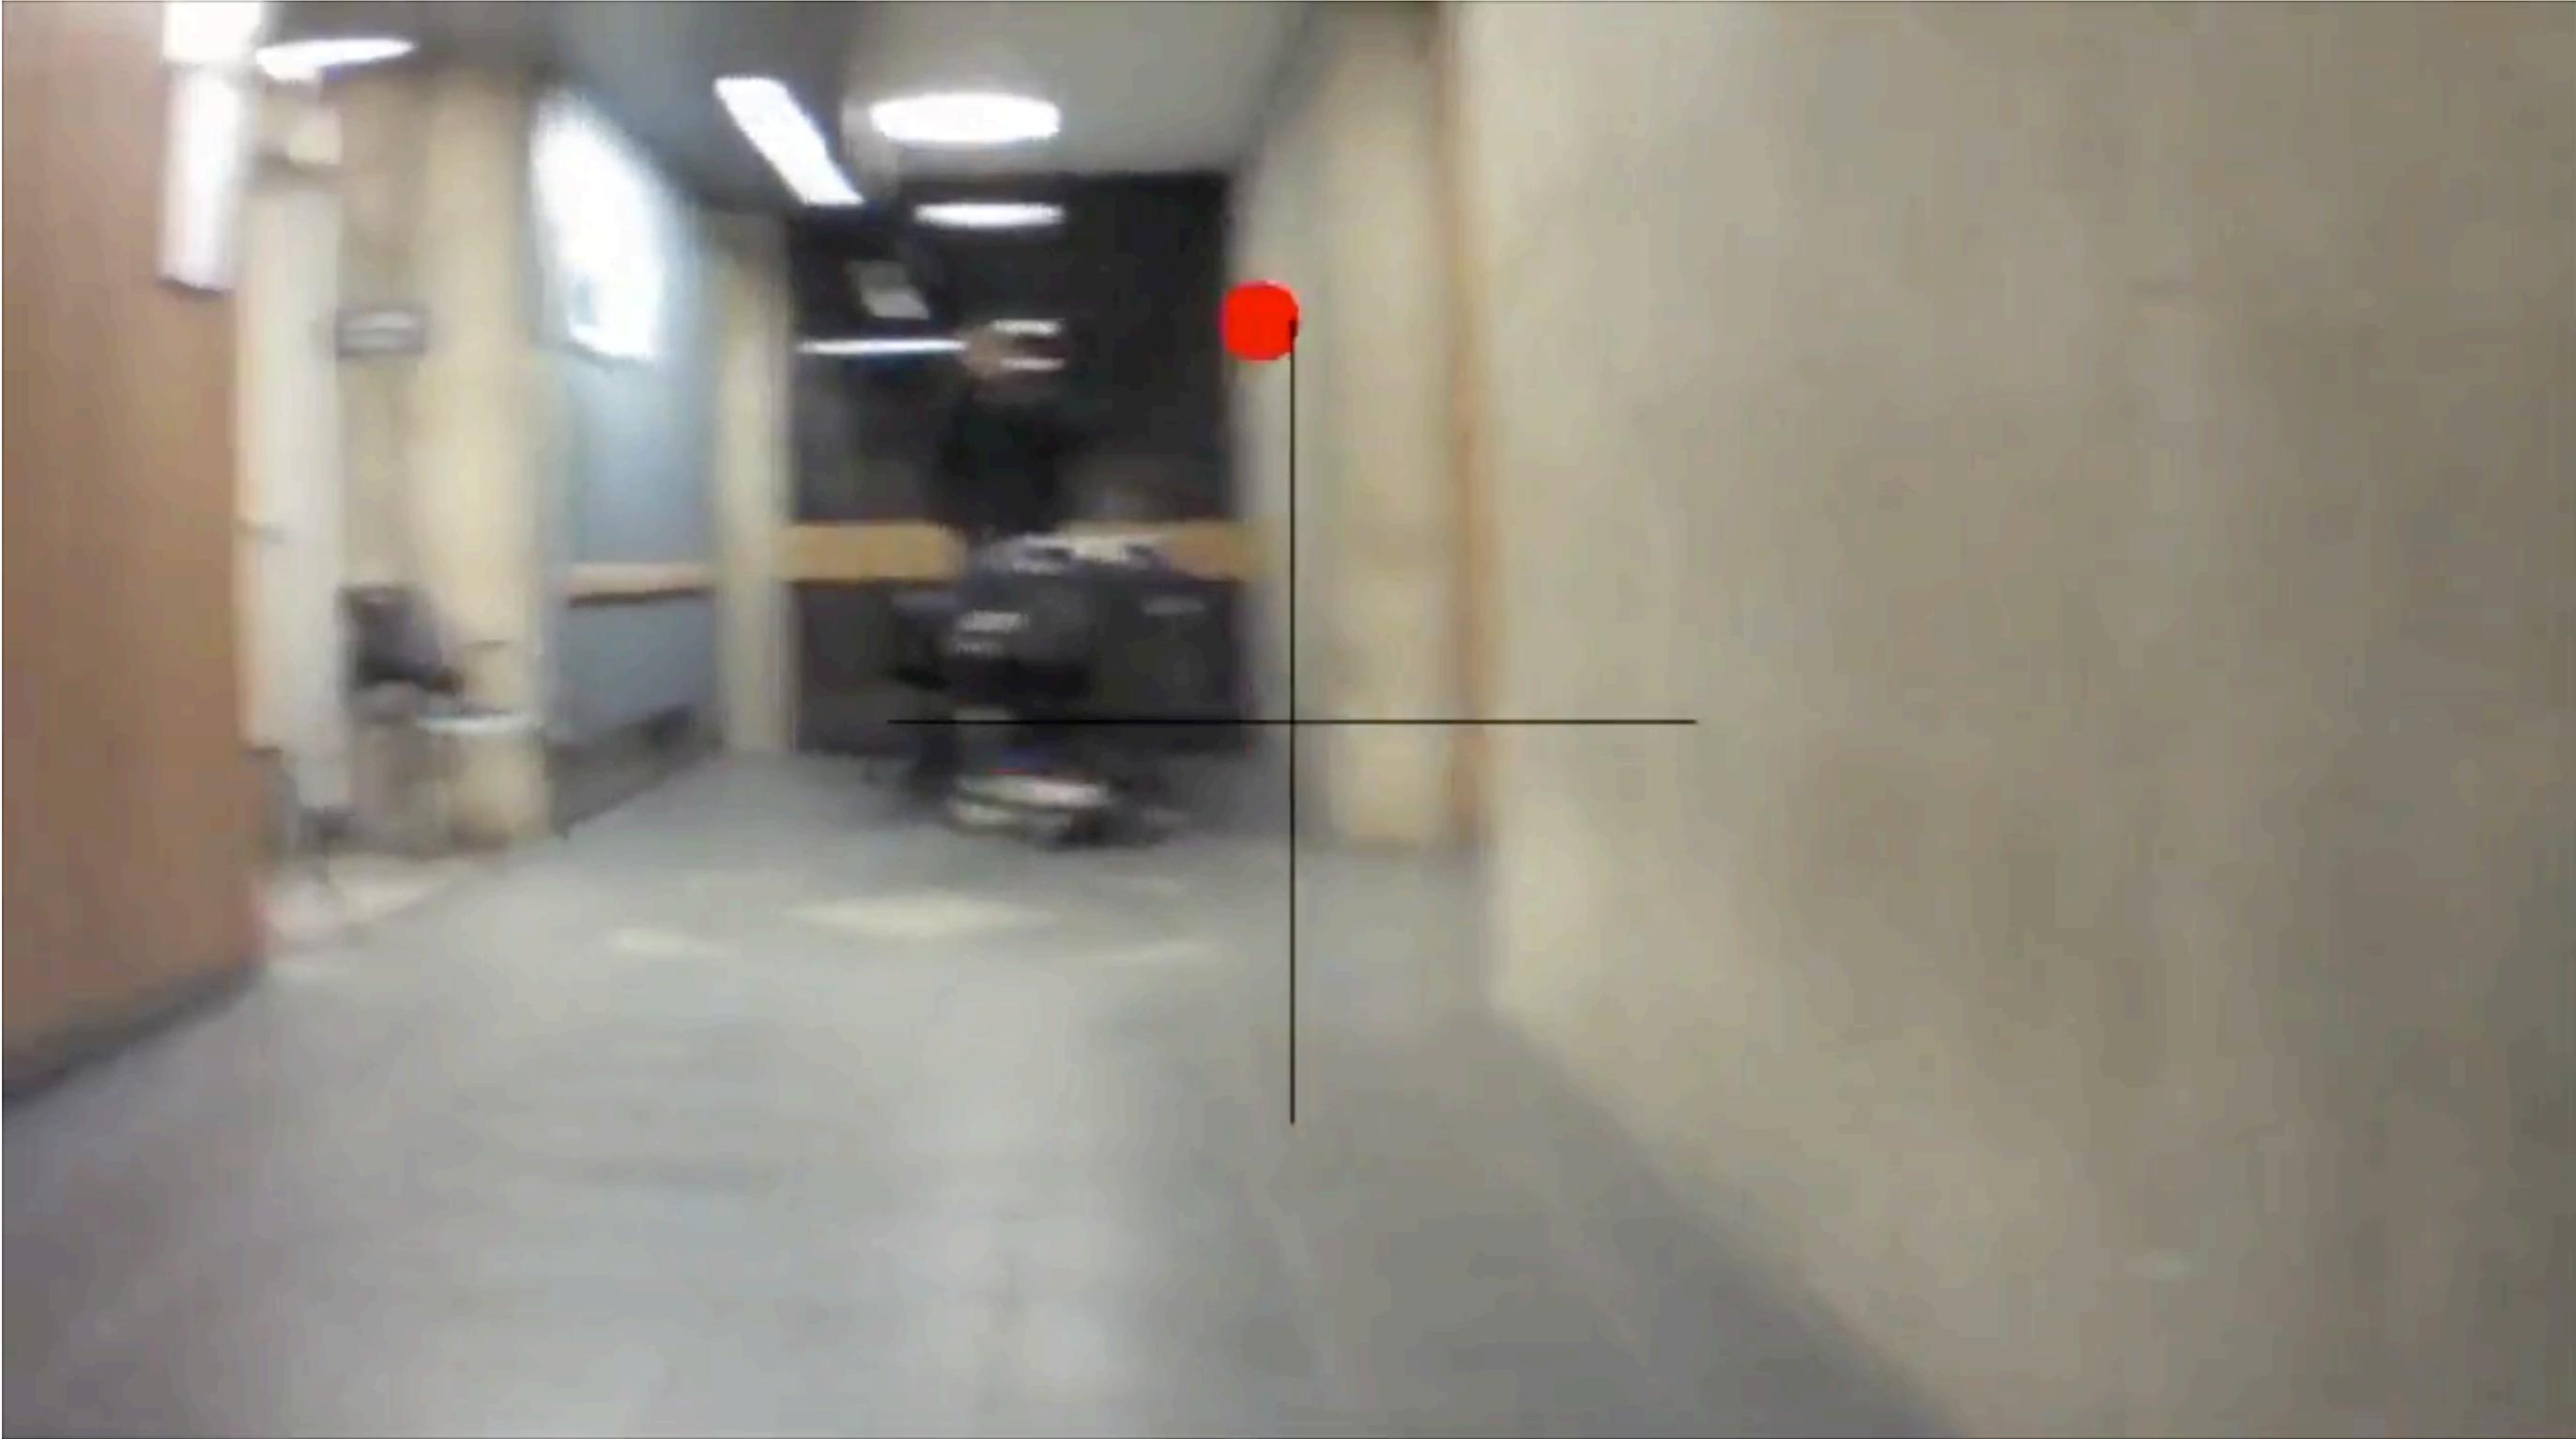
- **Agile Autonomous Driving** Y. Pan et al. R:SS18 [pdf](#), [video](#)



- [Beauty and the Beast: Optimal Methods Meet Learning for Drone Racing](#) E. Kaufmann et al. ICRA19 [pdf](#), [video](#)



**Learning to Fly by Crashing** D. Gandhi, L. Pinto, A. Gupta. IROS17. [pdf](#), [video](#)



**Self-supervised Deep Reinforcement Learning with Generalized Computation Graphs for Robot Navigation** G. Kahn S. Levine. ICRA18. [pdf](#), [video](#)

## Self-supervised Deep Reinforcement Learning with Generalized Computation Graphs for Robot Navigation



Gregory Kahn, Adam Villaflor, Bosen Ding, Pieter Abbeel, Sergey Levine



BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

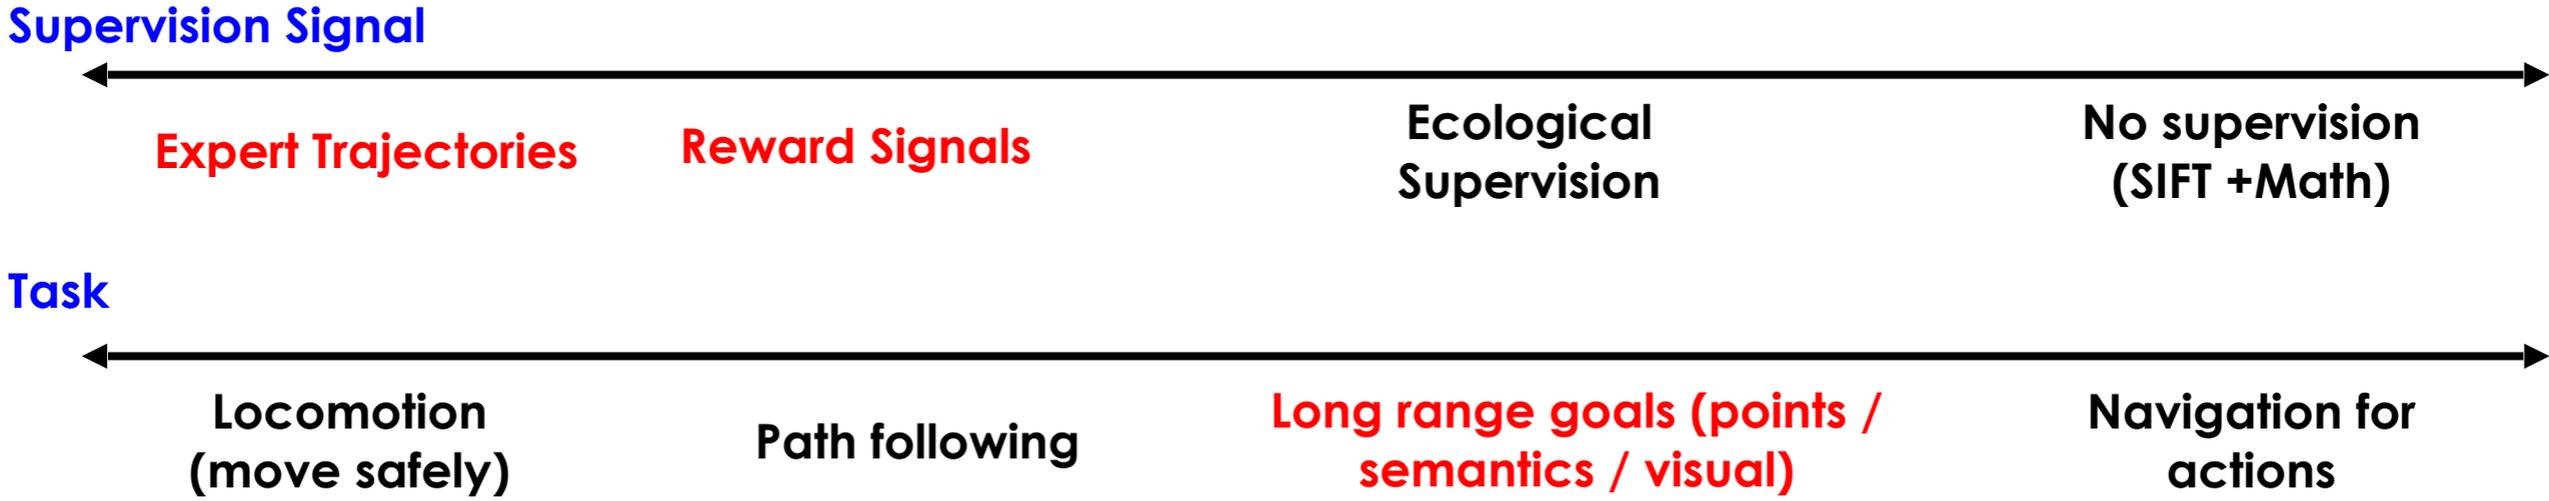
# Goal directed navigation

- Goal-directed behavior
  - Goals known / unknown
  - Environment seen / novel
- Central question is that of how to represent space?
  - Reason about what you have seen and where / what you have not seen
  - Spatial reasoning
  - Robustness to actuation noise
- For unknown environments, we may additionally need:
  - Semantic reasoning
  - Easy to build

# Representations for Space

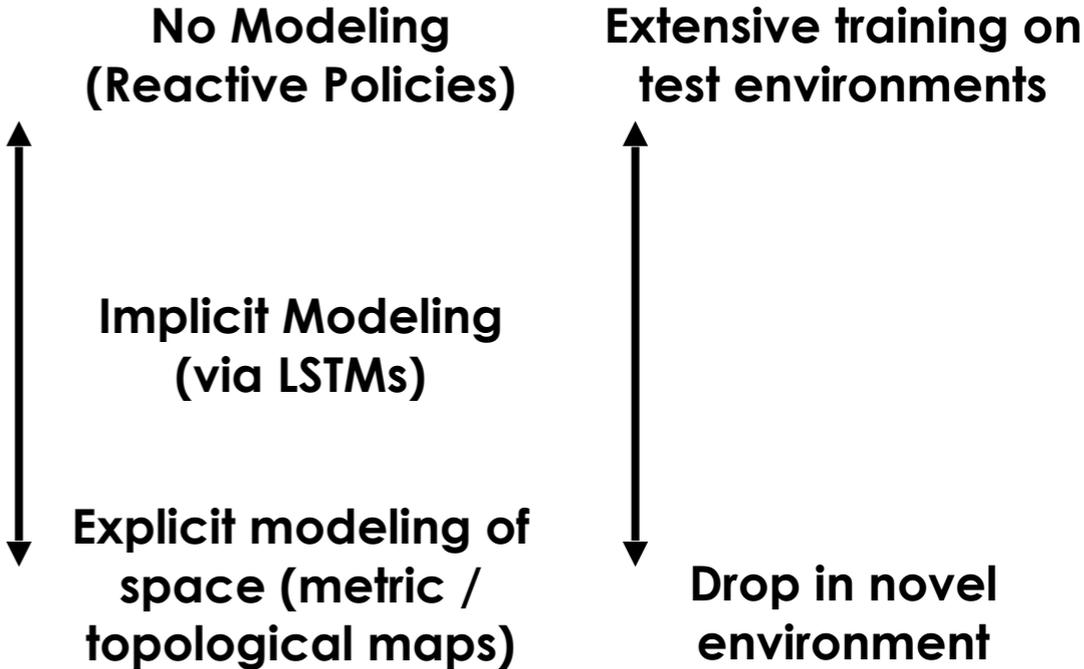
- Metric maps
  - Occupancy maps
  - Occupancy maps flavored with semantic information
- Topological maps
  - Collection of images, and connectivity information

# Learned Task driven End-to-end navigation



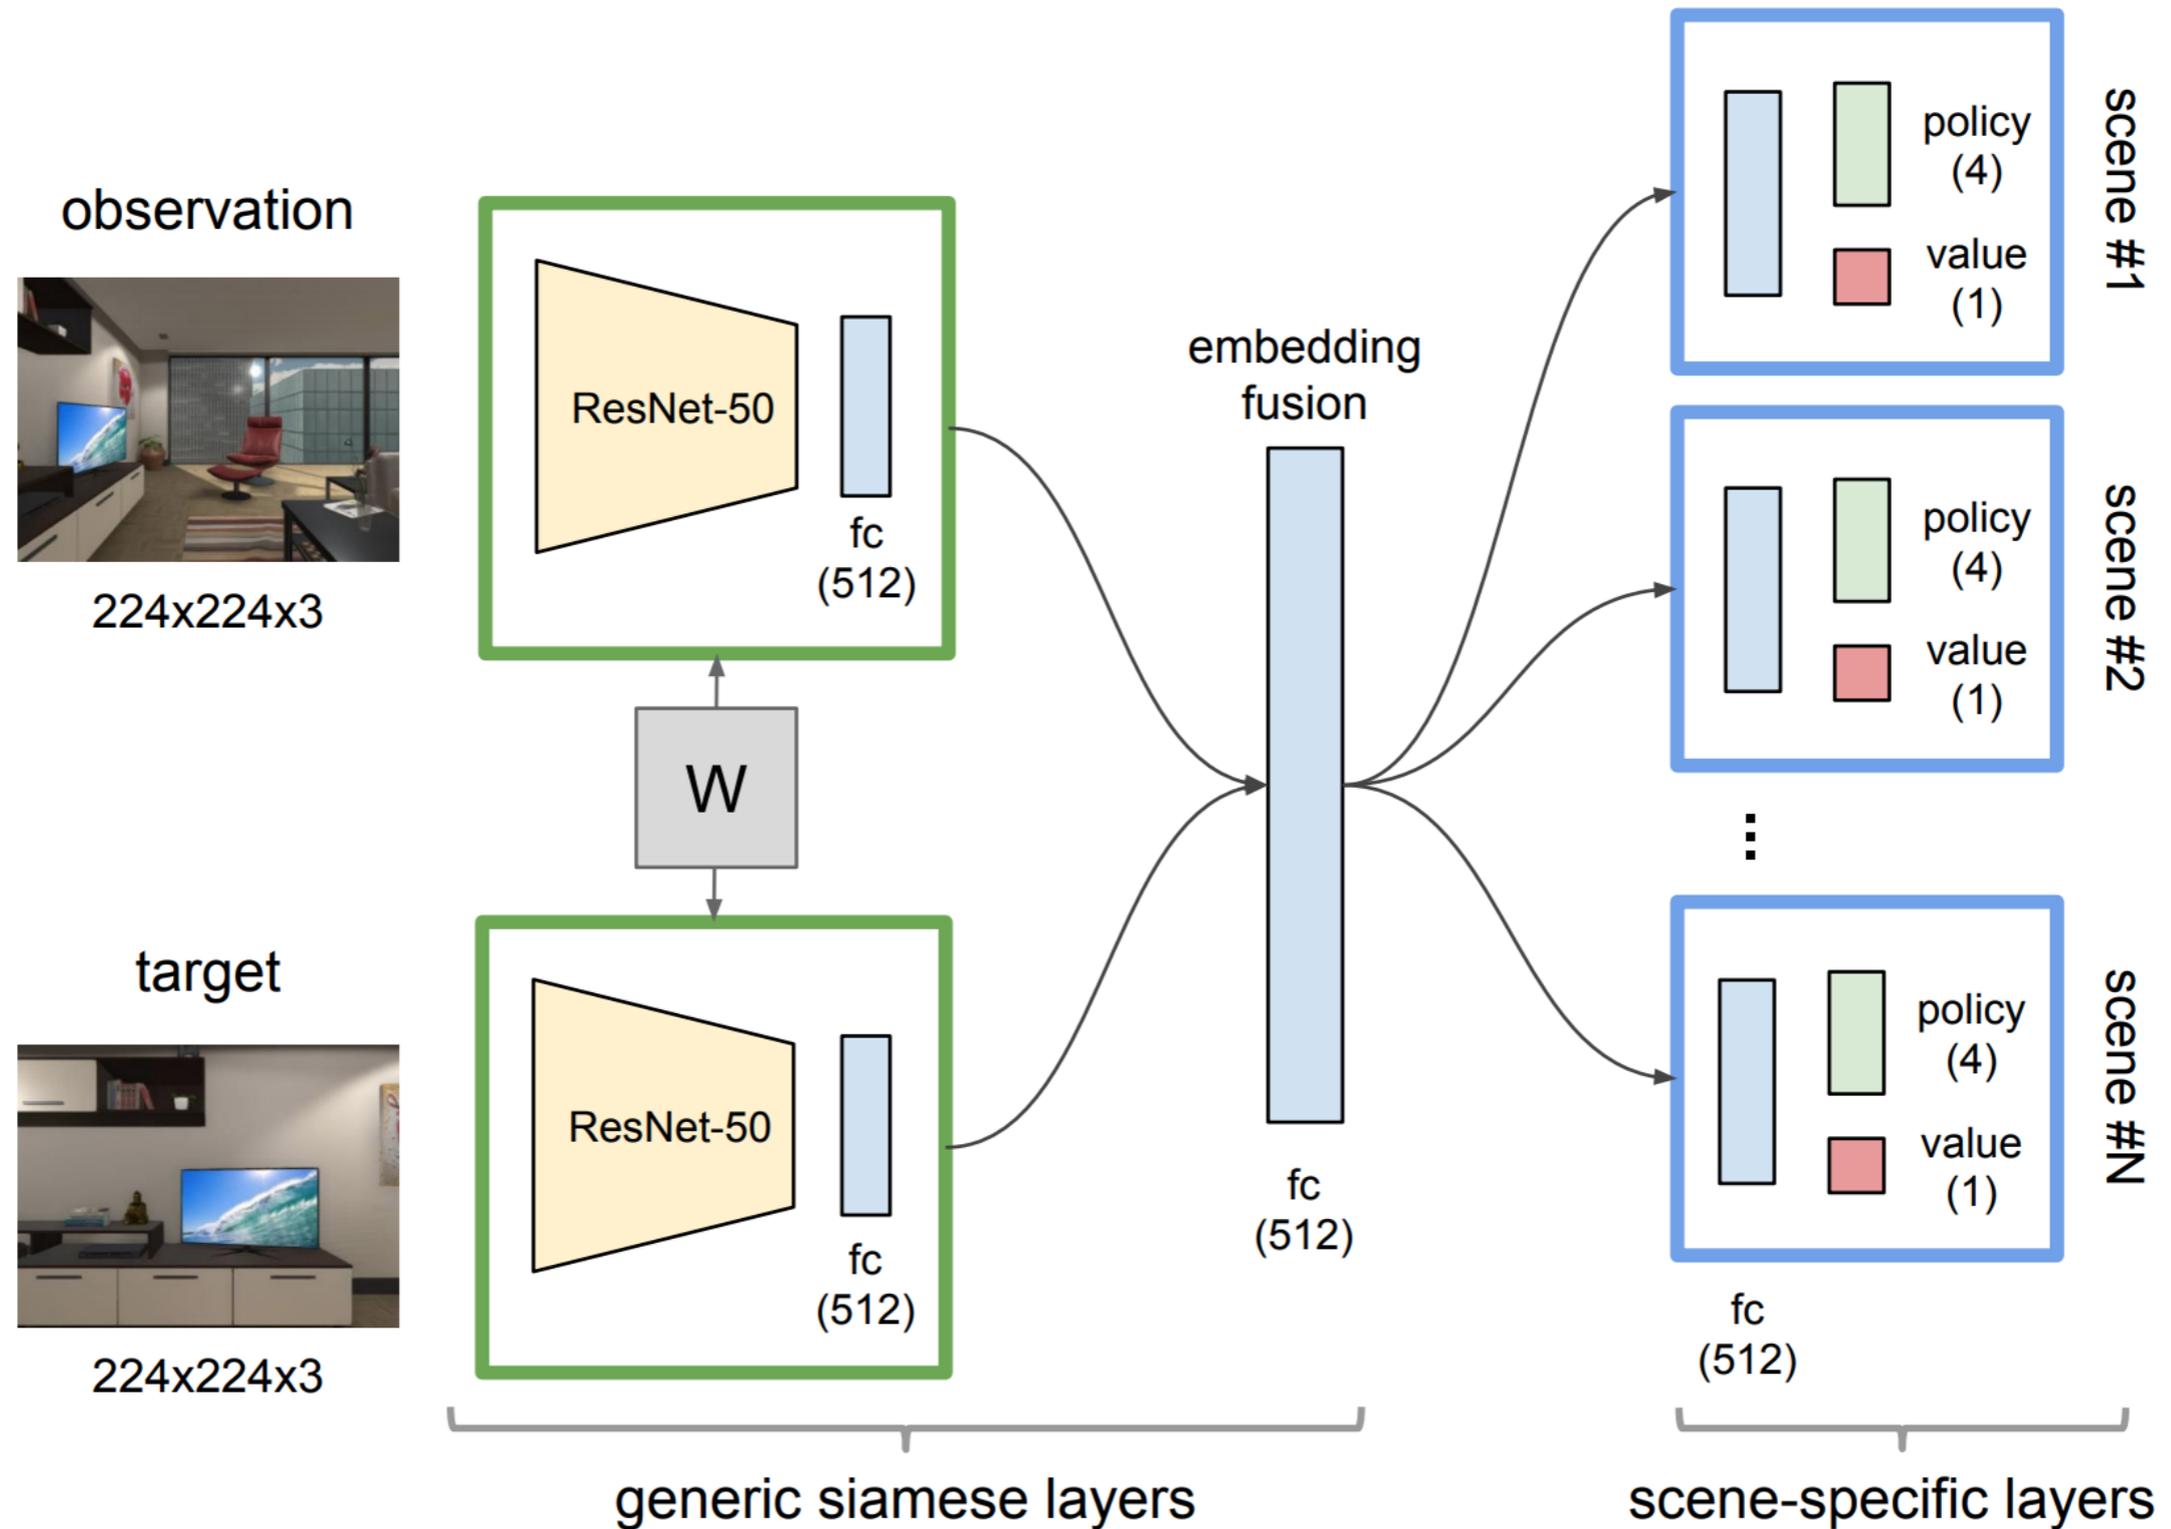
- Target driven visual navigation in indoor scenes using deep reinforcement learning. Zhu, ... Farhadi ICRA17 [pdf](#), [video](#)
- Learning to navigate in complex environments P. Mirowski, ... R. Hadsell ICLR17 [pdf](#), [video](#)
- Cognitive Mapping and Planning CVPR17 [pdf](#), [video](#)
- Neural Topological SLAM CVPR 20 [pdf](#), [video](#)

**Environment Modeling**                      **Generalization**



# Target driven visual navigation in indoor scenes using deep reinforcement learning.

Y. Zhu, ..., A. Gupta, ..., A. Farhadi. ICRA17. [pdf](#), [video](#)



# Simulation Environments

**Abstraction**



**Environment Properties**

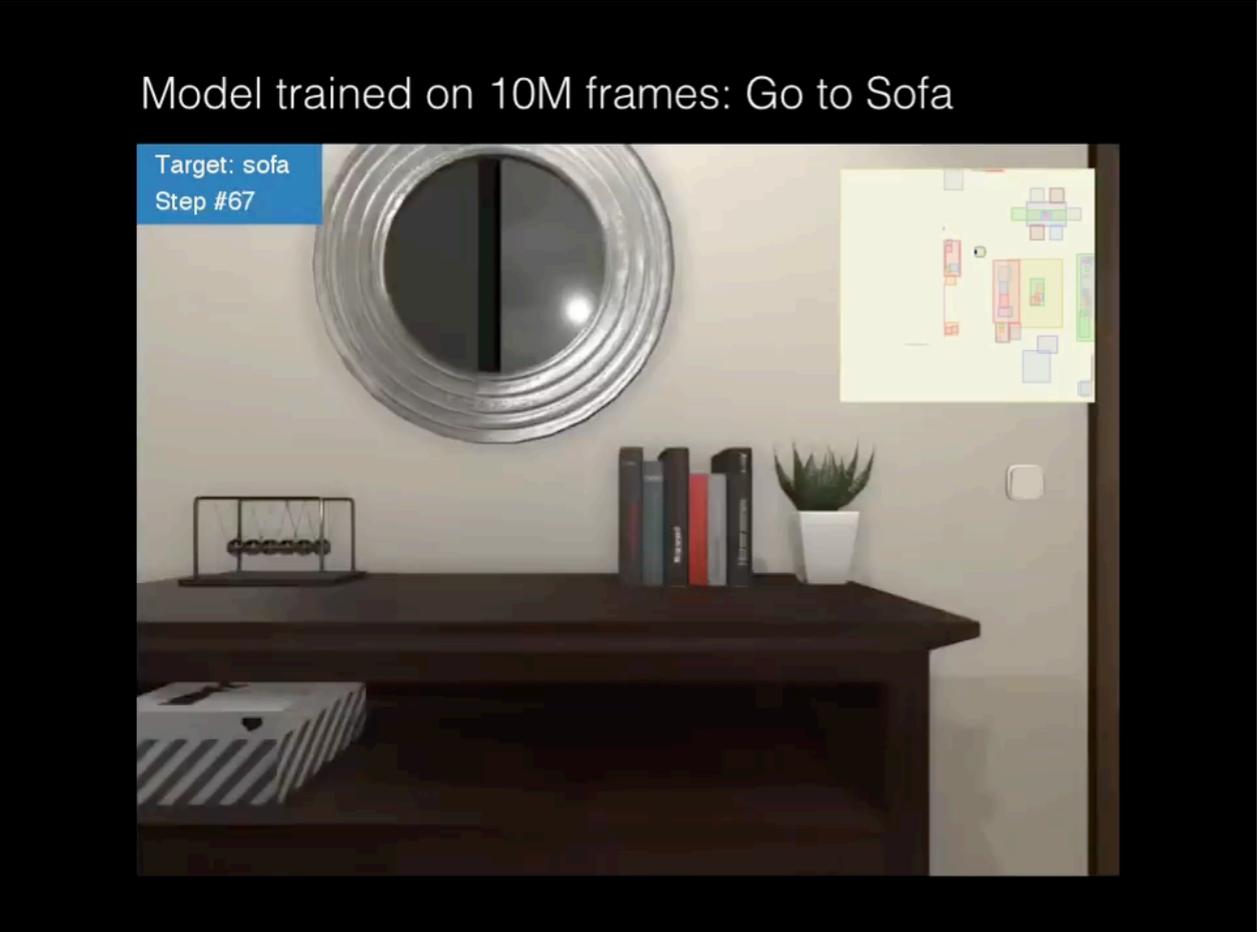


**Experimental setup**



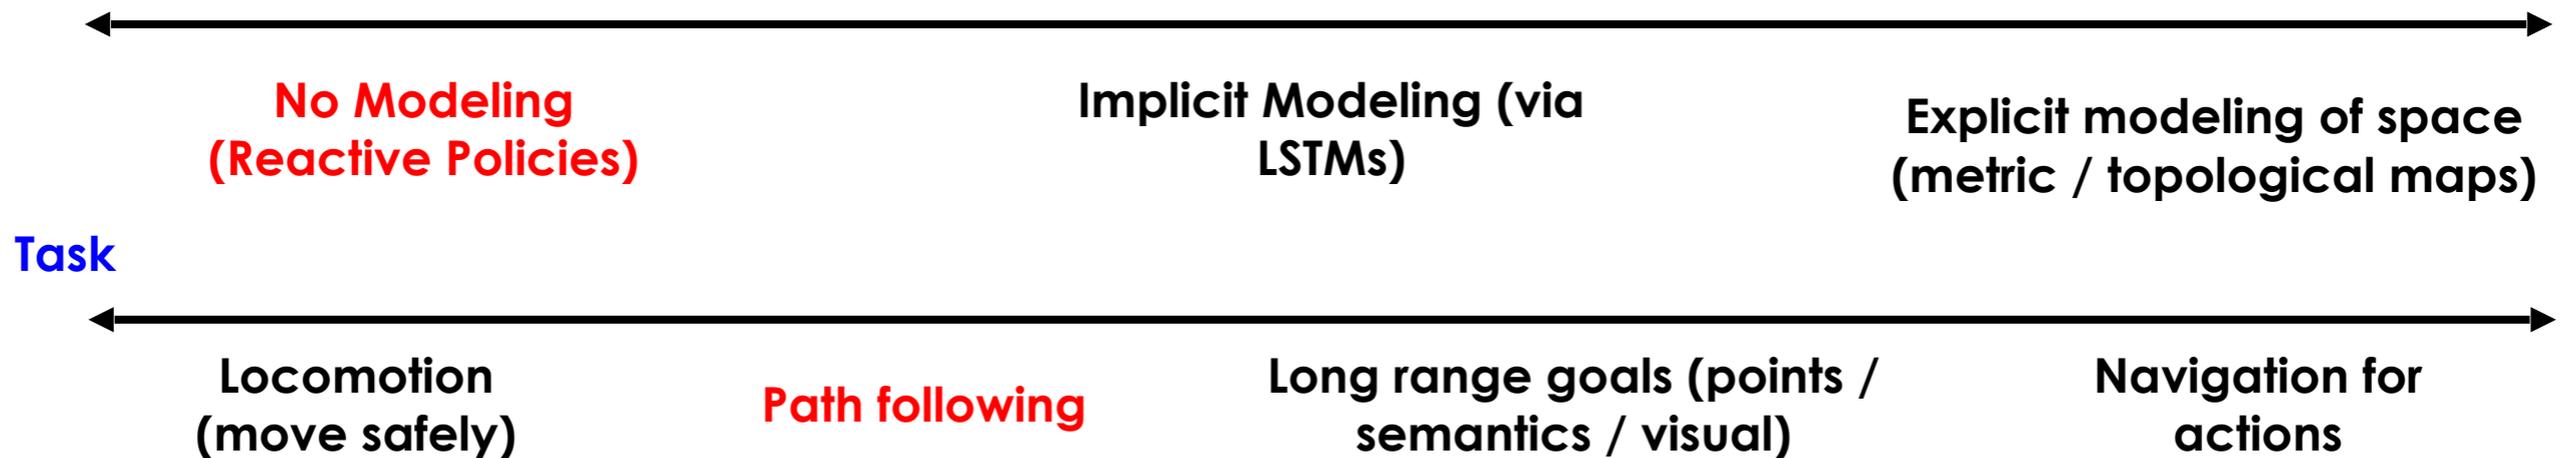
Simulators based off games / CG (eg: AI2Thor, VizDoom, CARLA)

Simulators based off real world scans (eg: CMPSim, Habitat, Gibson)



# Visual Servoing

## Environment Modeling



- **Visual Memory for Robust Path Following**  
Kumar, Gupta, ... Malik
- **End-to-end Driving via Conditional Imitation Learning** F. Codevilla, ... V. Koltun ICRA18? [pdf](#), [video](#)
- **Learning Visual Servoing with Deep Features and Fitted Q-Iteration**, Alex X. Lee, Sergey Levine, Pieter Abbeel. ICLR17. [pdf](#)
- **Zero-Shot Visual Imitation** D. Pathak, ..., T. Darrell pdf, video ICLR18? [pdf](#), [video](#)
- **Classical Visual Servoing**

## Supervision Signal

Expert Trajectories

Reward Signals

Ecological Supervision

No supervision (SIFT + Geometry)

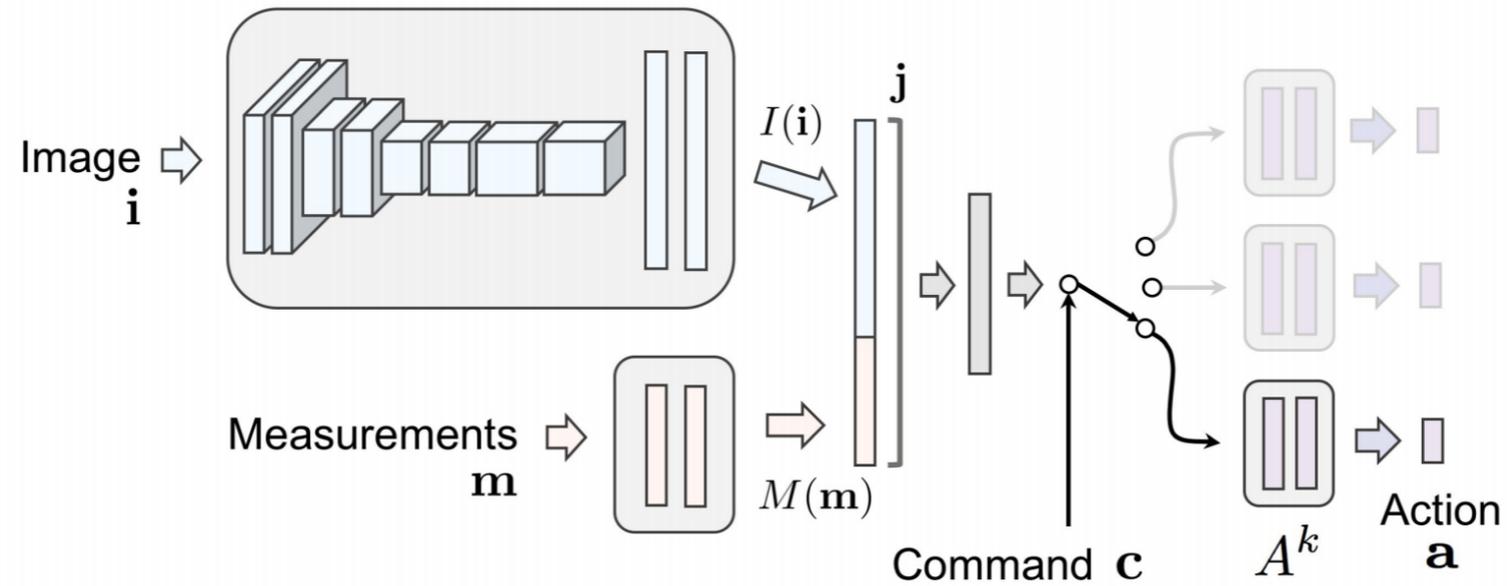
# Learning Visual Servoing with Deep Features and Fitted Q-Iteration

Alex X. Lee, Sergey Levine, Pieter Abbeel. ICLR17. [pdf](#)



# End-to-end Driving via Conditional Imitation Learning

F. Codevilla, ... V. Koltun ICRA18? [pdf](#), [video](#)



# End-to-end Driving via Conditional Imitation Learning

Felipe Codevilla, Matthias Mueller, Alexey Dosovitskiy, Antonio Lopez, Vladlen Koltun

Submitted to ICRA 2018

# Other Unexplored Axes?

## Environment Scale



Small mazes

Houses / Offices

City Scale Environments?

## Task



Locomotion (move safely)

Path following

Long range goals (points / semantics / visual)

Navigation for actions?

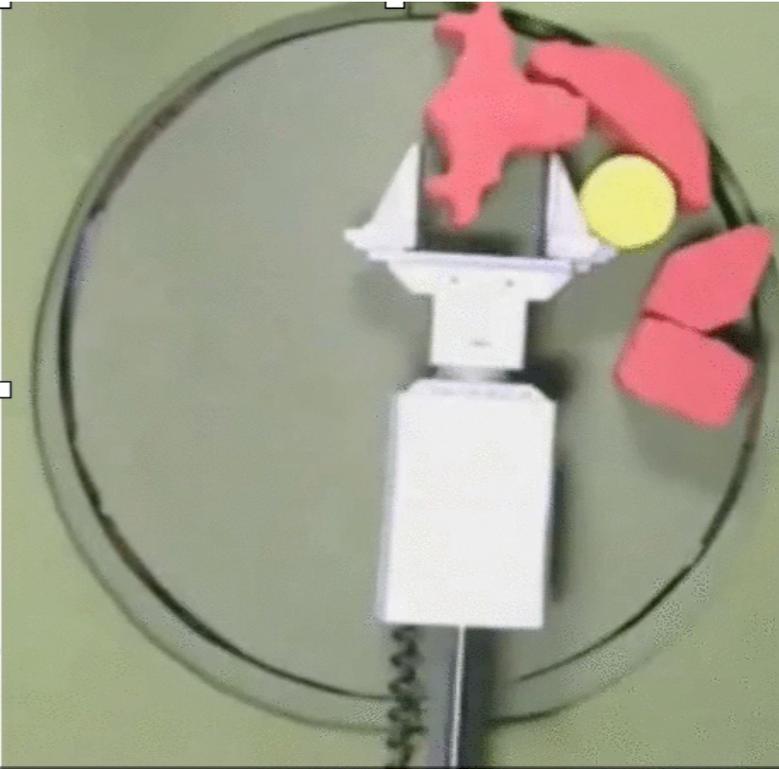
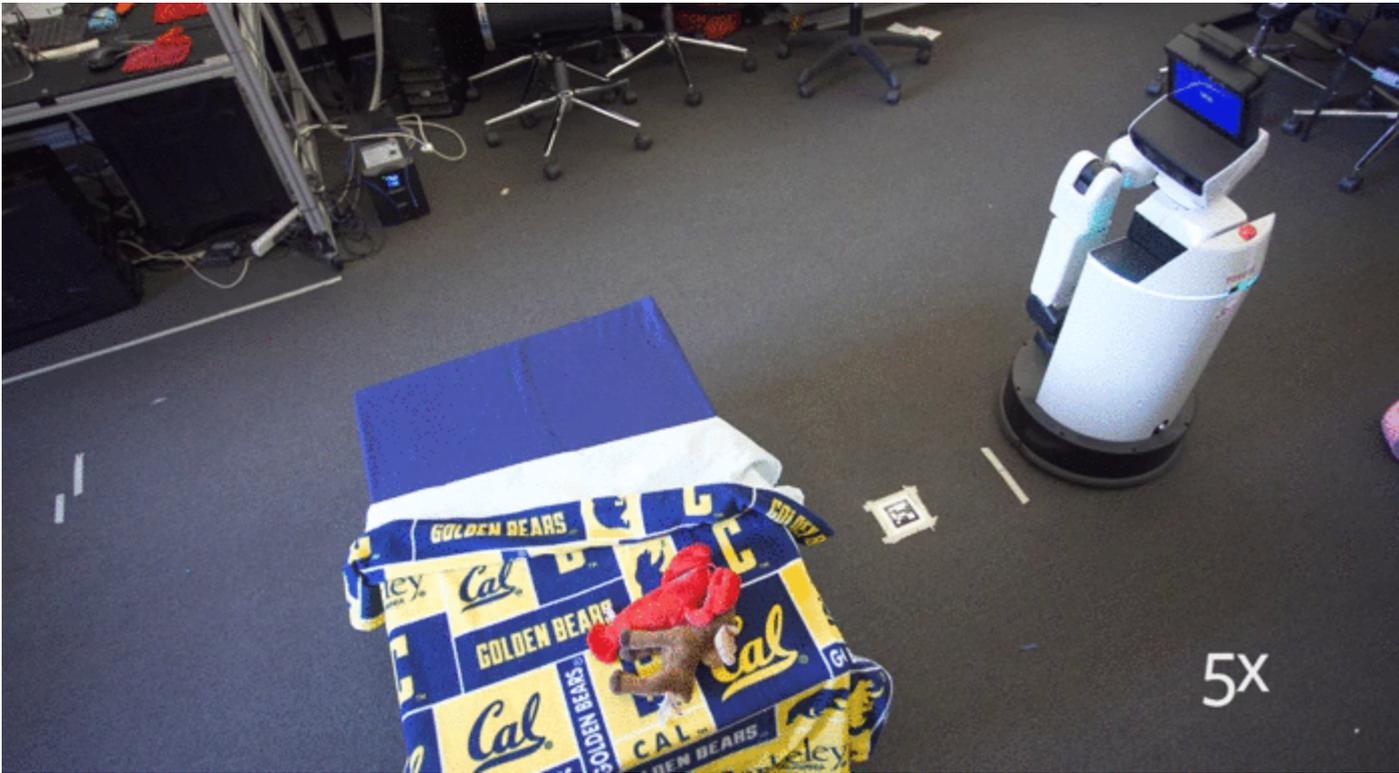
## Dynamic Environments



World is static

World has objects that need to be moved around?

Reason about other agents?



# Discussion

- *What if any is the role of learning for solving these different navigation tasks?*

# Agile Autonomous Driving using End-to-End Deep Imitation Learning

Yunpeng Pan<sup>\*‡</sup>, Ching-An Cheng<sup>\*</sup>, Kamil Saigol<sup>\*</sup>, Keuntaek Lee<sup>†</sup>, Xinyan Yan<sup>\*</sup>,  
Evangelos A. Theodorou<sup>\*</sup>, and Byron Boots<sup>\*</sup>

<sup>\*</sup>Institute for Robotics and Intelligent Machines, <sup>†</sup>School of Electrical and Computer Engineering  
Georgia Institute of Technology, Atlanta, Georgia 30332–0250

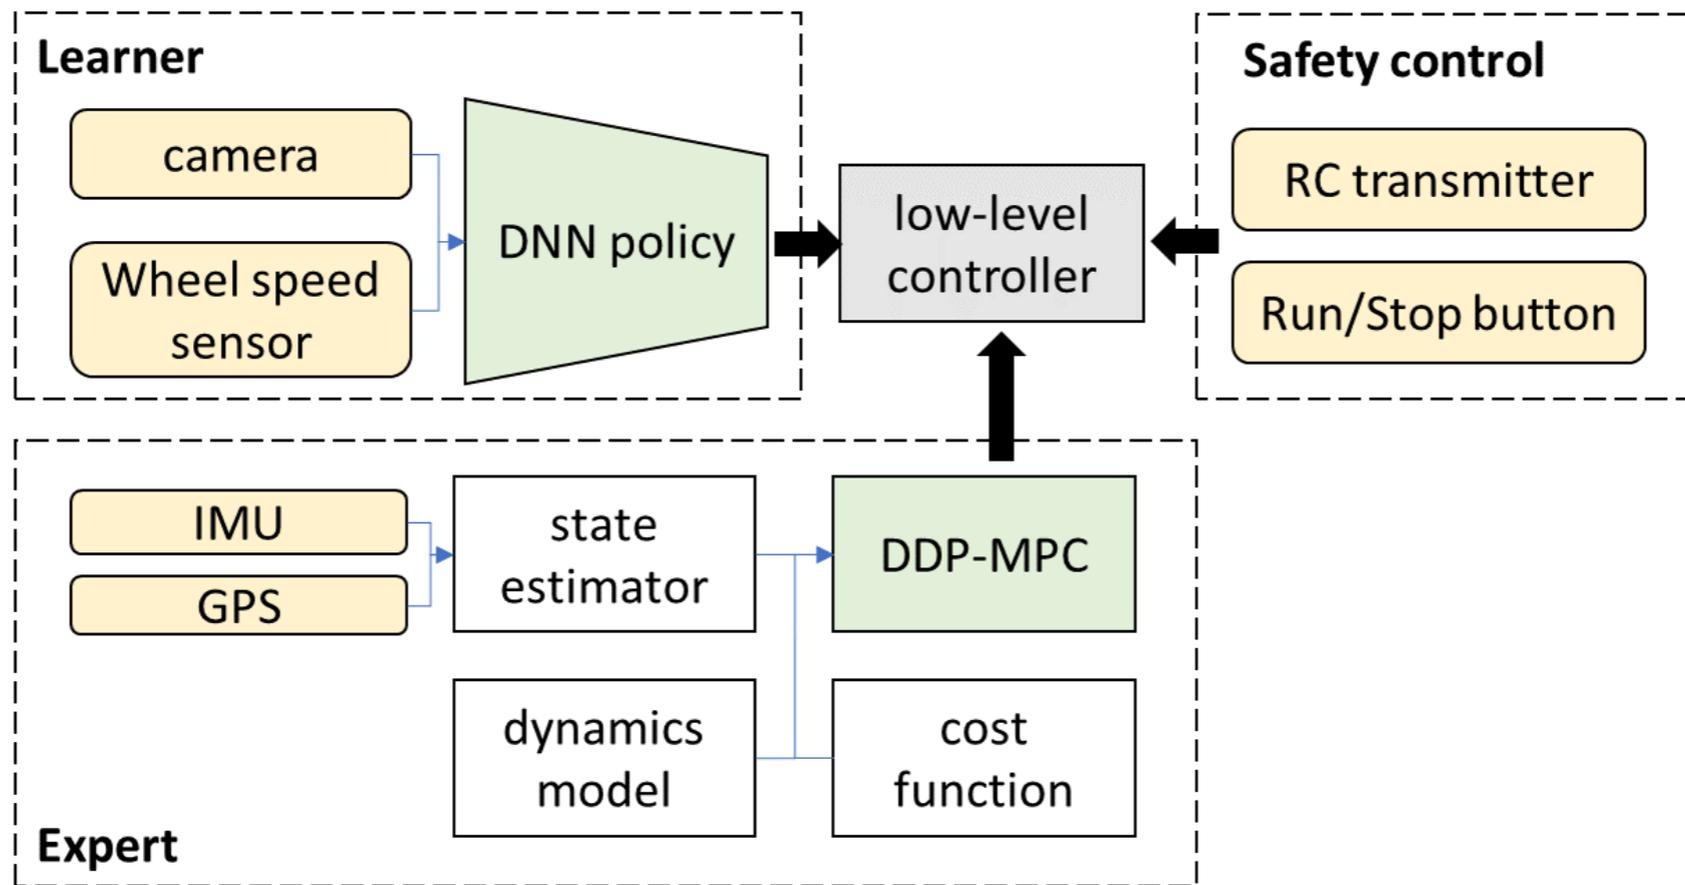
<sup>‡</sup>JD.com American Technologies Corporation, Mountain View, California 94043

{ypan37, cacheng, kamilsaigol, keuntaek.lee, xyan43}@gatech.edu  
evangelos.theodorou@gatech.edu, boots@cc.gatech.edu

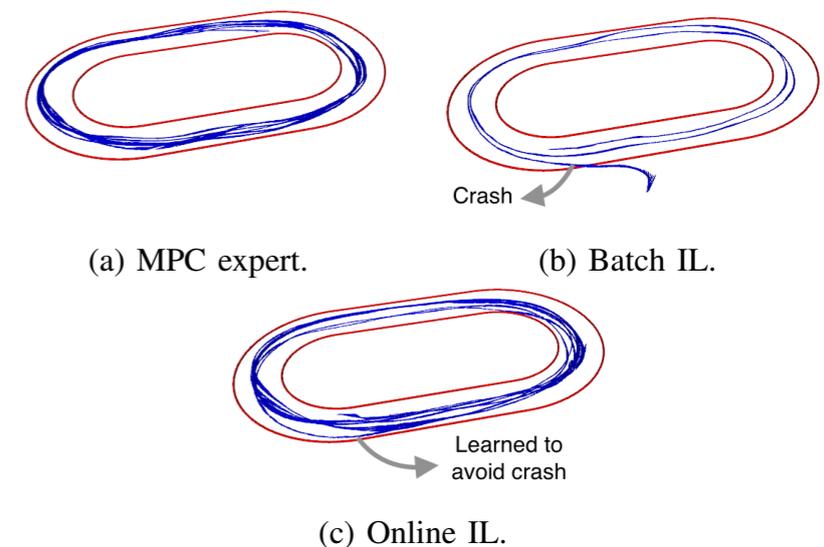




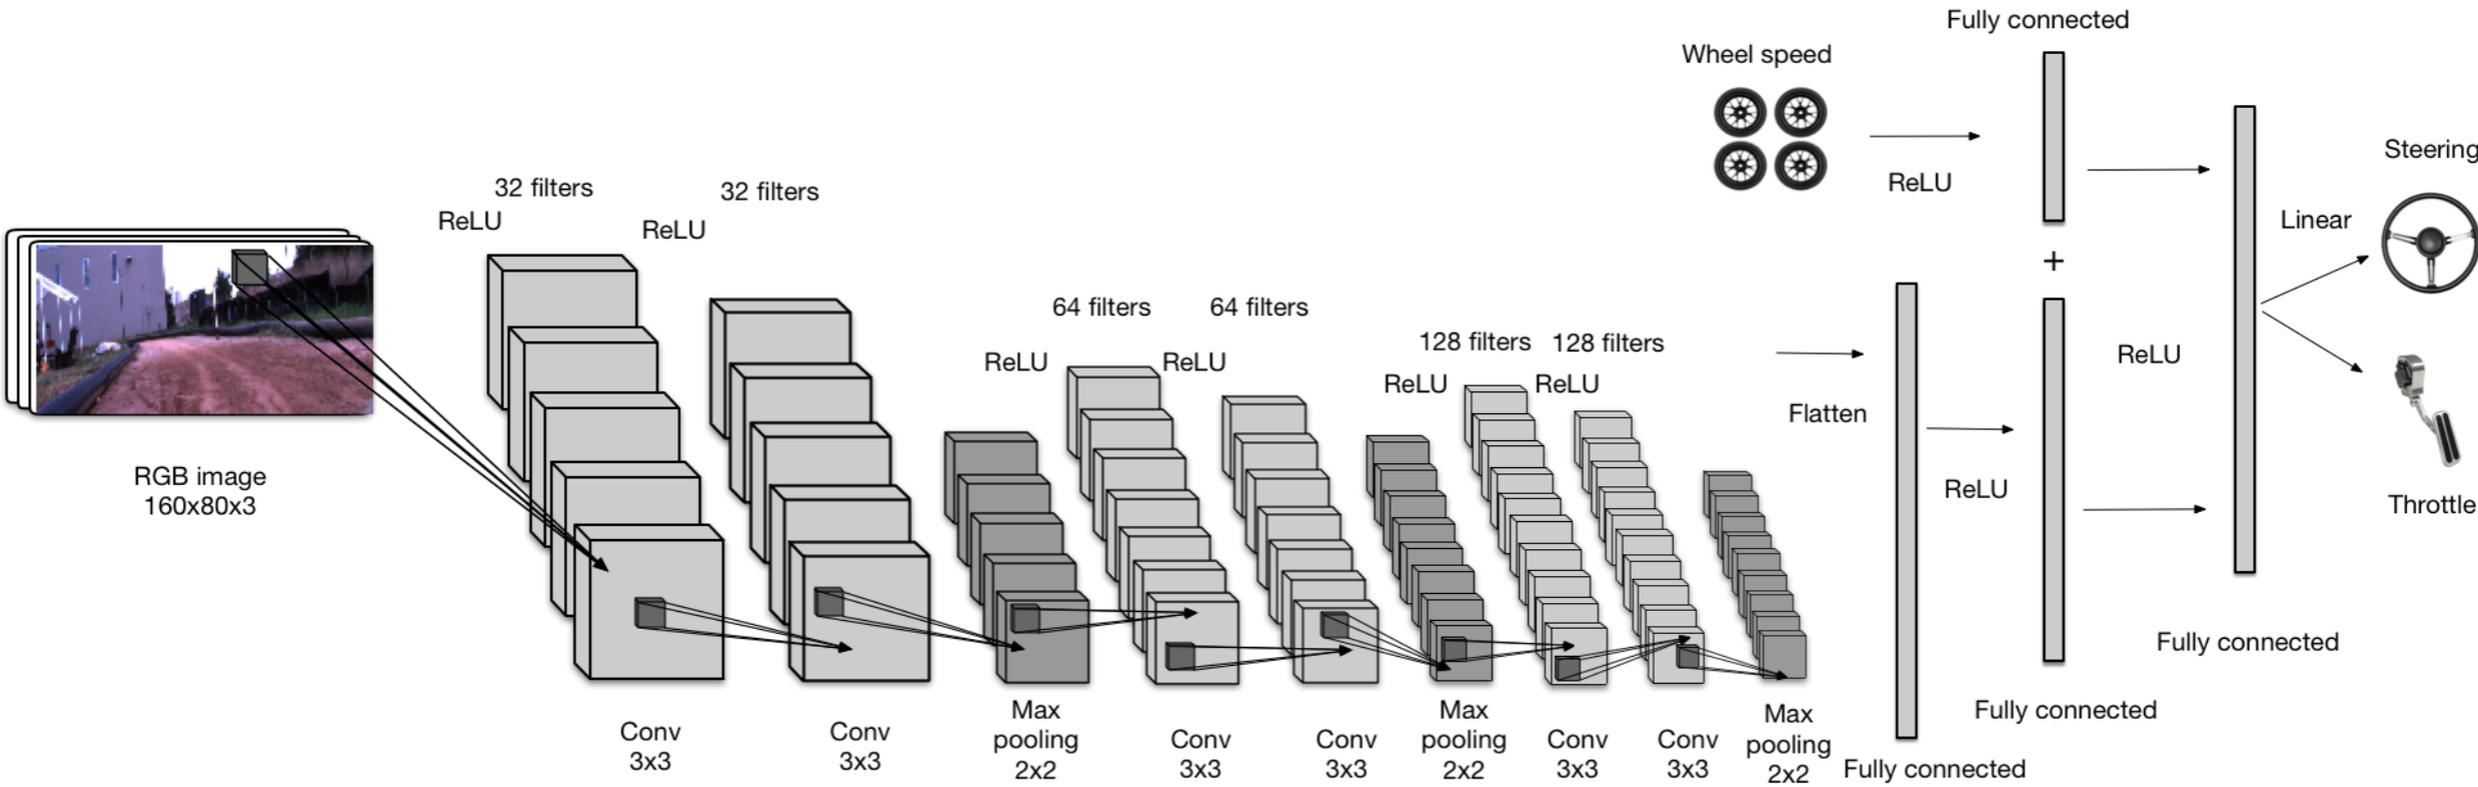
# Overview



- Algorithmic expert uses MPC with expensive state-estimators
- Learner predicts controls output by expert
- Use DAgger, instead of vanilla behavior cloning

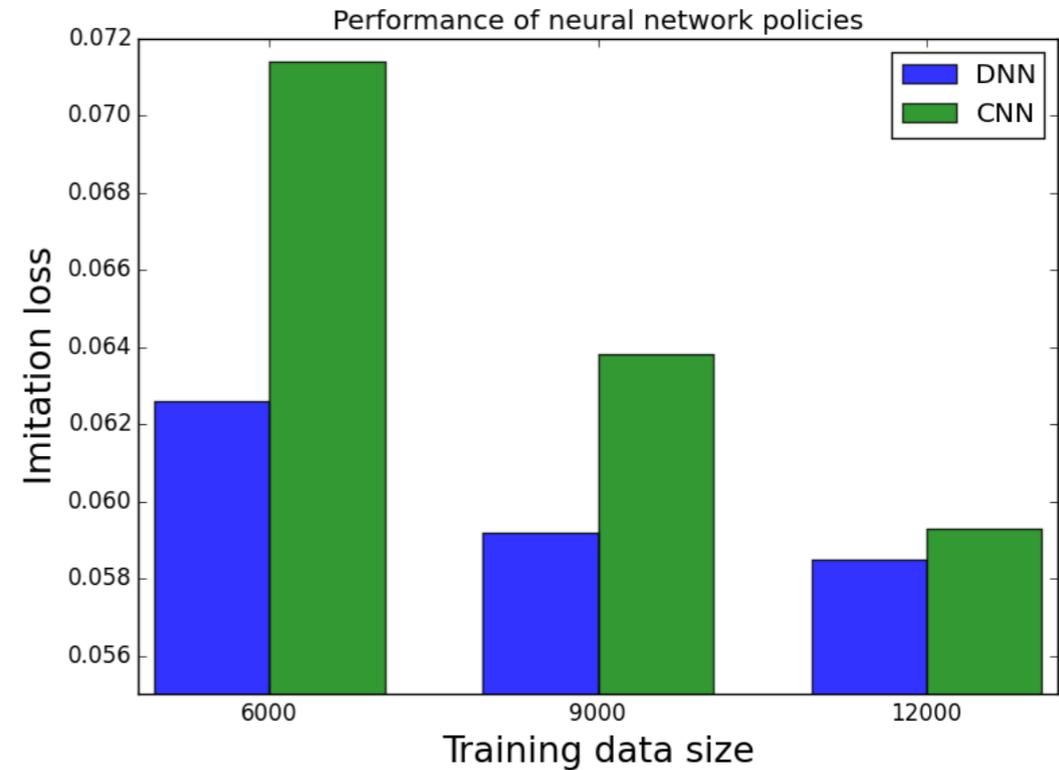
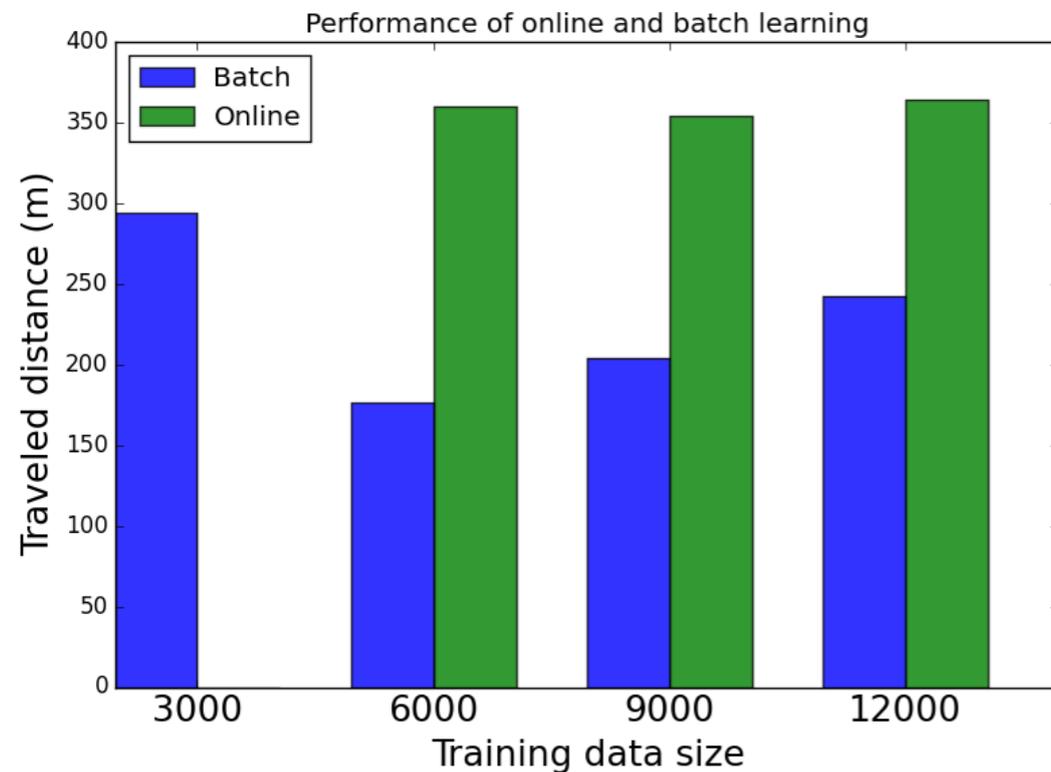


# Policy



# Results

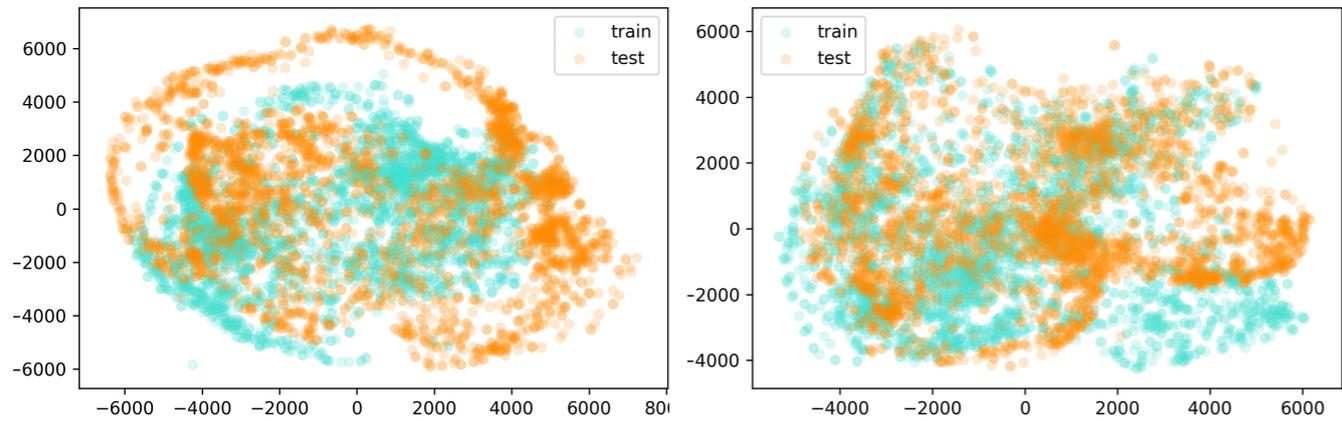
Policy	Avg. speed	Top speed	Training data	Completion ratio	Total loss	Steering/Throttle loss
Expert	6.05 m/s	8.14 m/s	N/A	100 %	0	0
Batch	4.97 m/s	5.51 m/s	3000	100 %	0.108	0.092/0.124
Batch	6.02 m/s	8.18 m/s	6000	51 %	0.108	0.162/0.055
Batch	5.79 m/s	7.78 m/s	9000	53 %	0.123	0.193/0.071
Batch	5.95 m/s	8.01 m/s	12000	69 %	0.105	0.125/0.083
Online (1 iter)	6.02 m/s	7.88 m/s	6000	100 %	0.090	0.112/0.067
Online (2 iter)	5.89 m/s	8.02 m/s	9000	100 %	0.075	0.095/0.055
Online (3 iter)	6.07 m/s	8.06 m/s	12000	100 %	0.064	0.073/0.055



Online IL does better than Batch IL

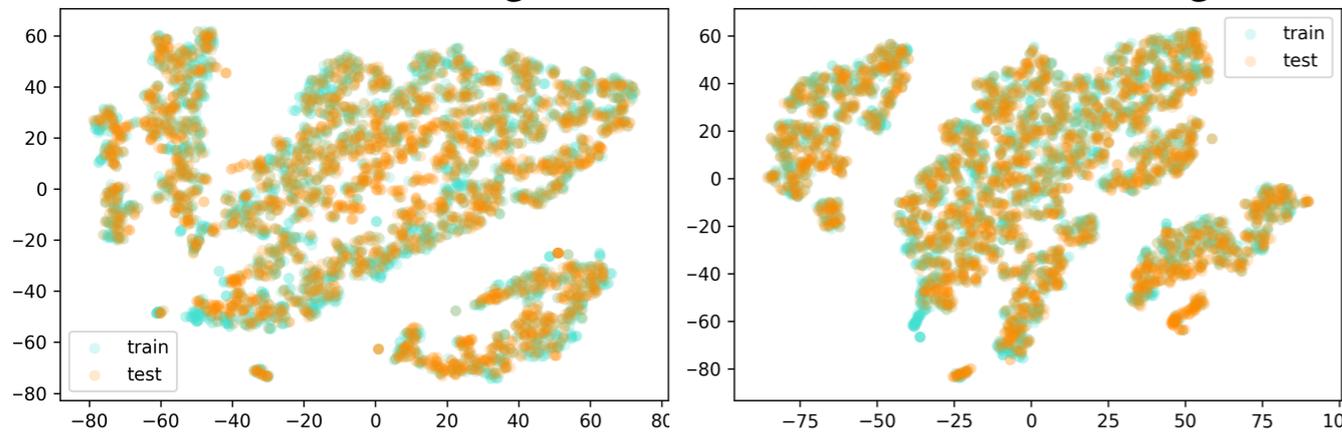
Speed as input to network helps

# Visualizations



(a) Batch raw image

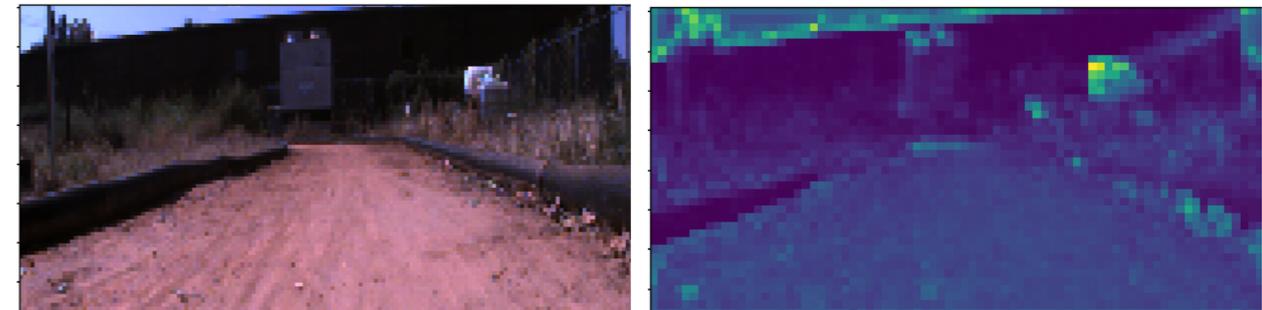
(b) Online raw image



(c) Batch wheel speed

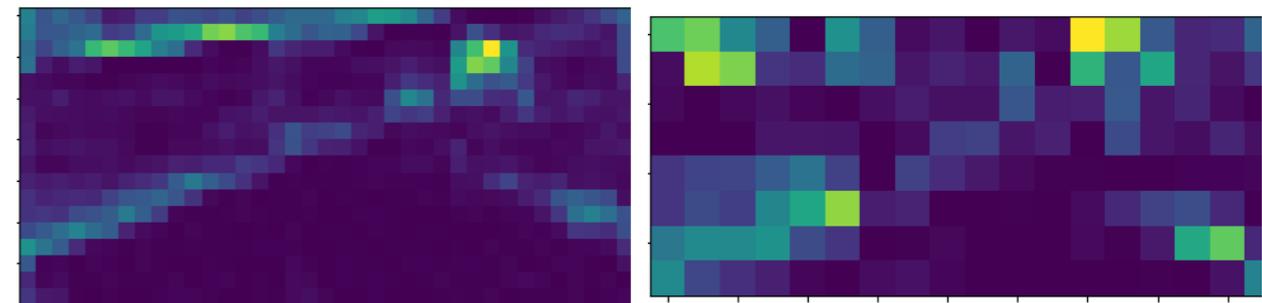
(d) Online wheel speed

Train and test distributions match better for online IL vs batch IL.



(a) raw image

(b) max-pooling1



(c) max-pooling2

(d) max-pooling3

Network seems to attend to places it should attend to.

Thank you