Cognitive Mapping and Planning for Visual Navigation

Saurabh Gupta James Davidson Sergey Levine Rahul Sukthankar Jitendra Malik

Google, UC Berkeley

Code, data & models online!

Problem Statement

Robot navigation in novel environments

Robot equipped with a first person camera

Dropped into a novel environment it has not been in before.

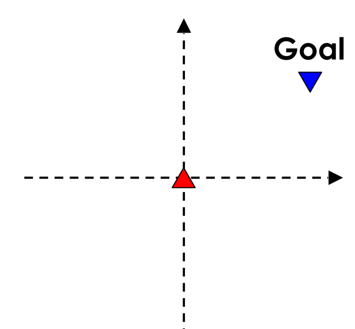

"Go Find me a Chair"

Navigate in the environment

Motivation

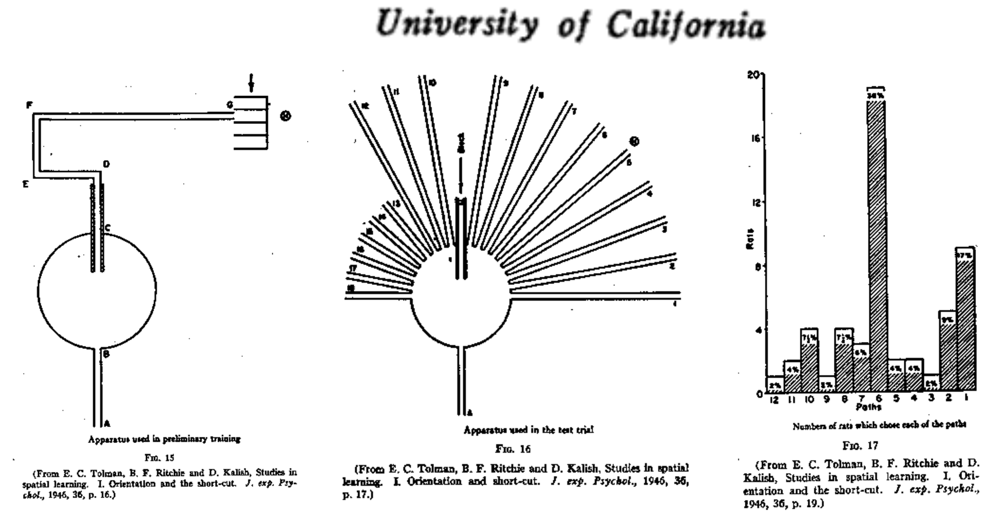

COGNITIVE MAPS IN RATS AND MEN [1]

BY EDWARD C. TOLMAN

University of California

Secondly, we assert that the central office itself is far more like a map control room than it is like an old-fashioned telephone exchange. The stimuli, which are allowed in, are not connected by just simple one-to-one switches to the outgoing responses. Rather, the incoming impulses are usually worked over and elaborated in the central control room into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release.

Classical Approaches

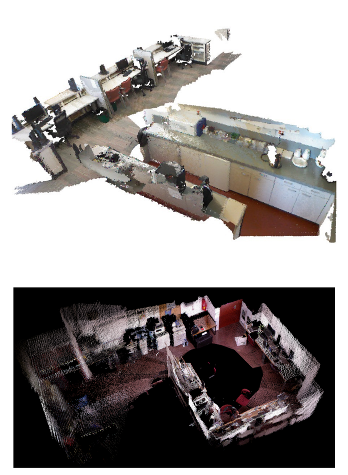

Mapping

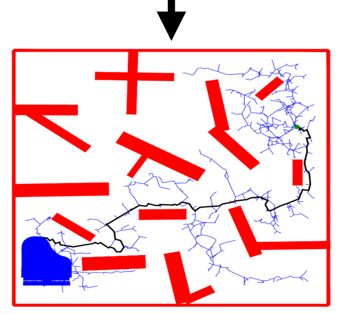

Planning

Modern Approaches

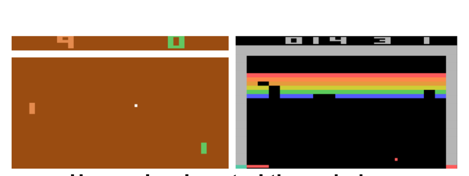

Human-level control through deep reinforcement learning, Mnih et al., Nature 2014

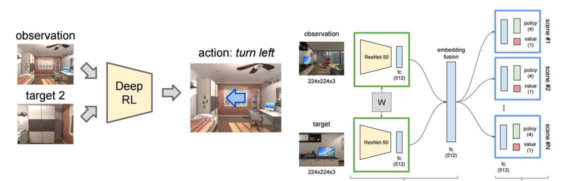

Target-driven Visual Navigation in Indoor Scenes using Deep Reinforcement Learning, Zhu et al., ICRA 2017

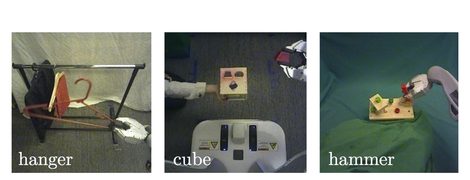

End-to-End Training of Deep Visuomotor Policies, Levine et al., JMLR 2015

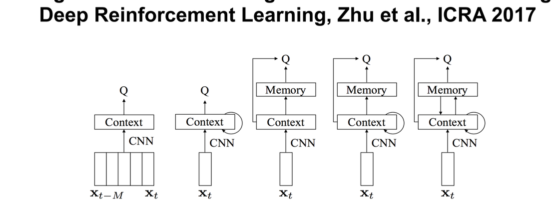

Control of Memory, Active Perception, and Action in Minecraft, Oh et al., ICML 2016

Classical Approaches

- **Over-complete** - Precise reconstruction of everything is not necessary
- **Incomplete** - Only geometry, no semantics. Nothing is known till it is explicitly observed, fail to exploit the structure of the world.
- Separation between mapping and planning.

Modern Approaches

- Ignore structure of the problem

Approach

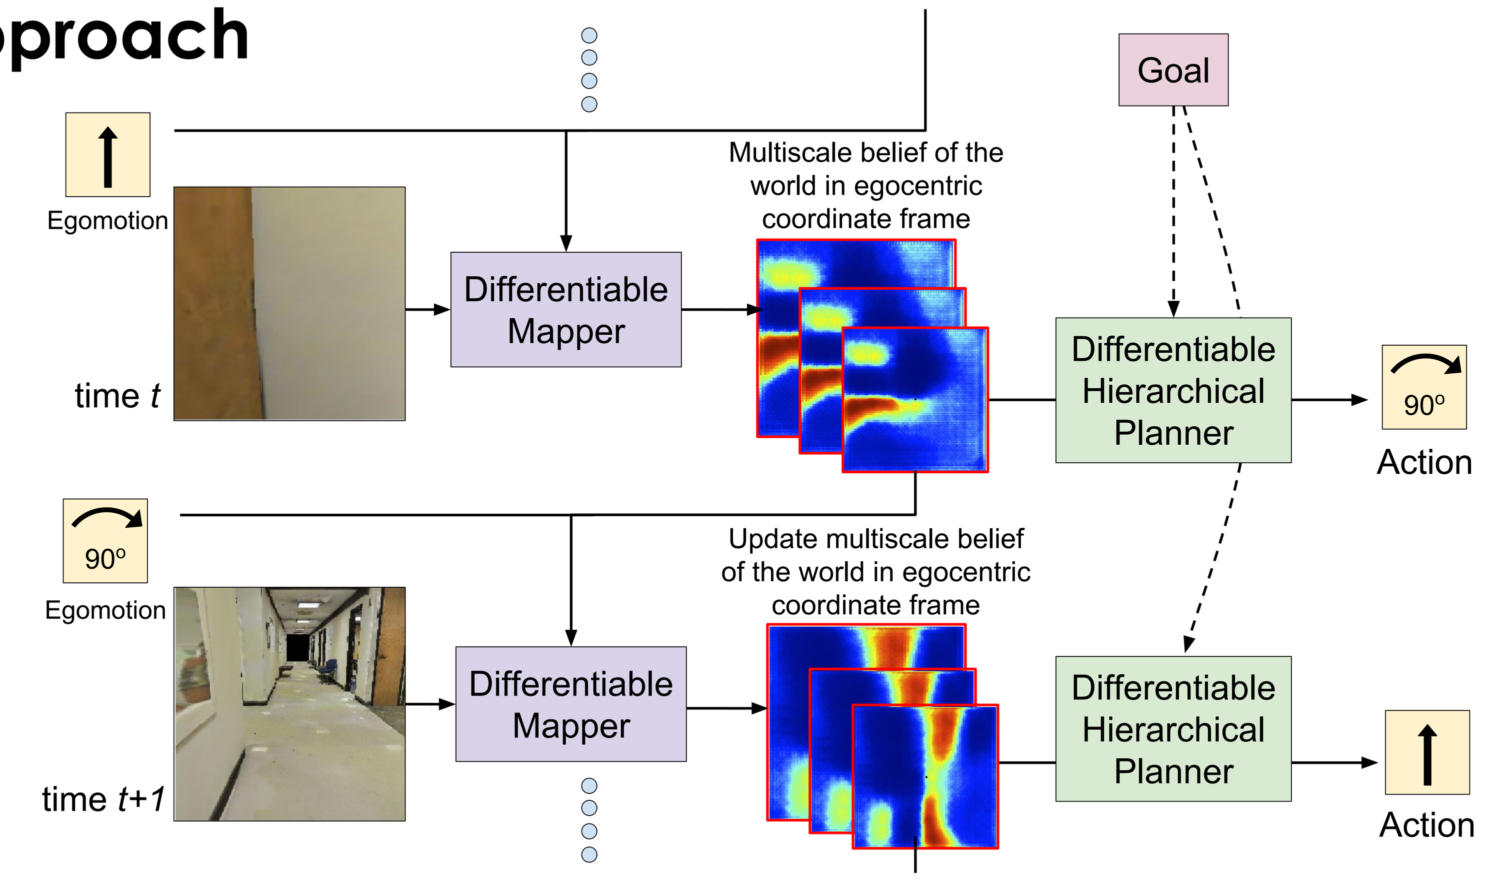

Mapper

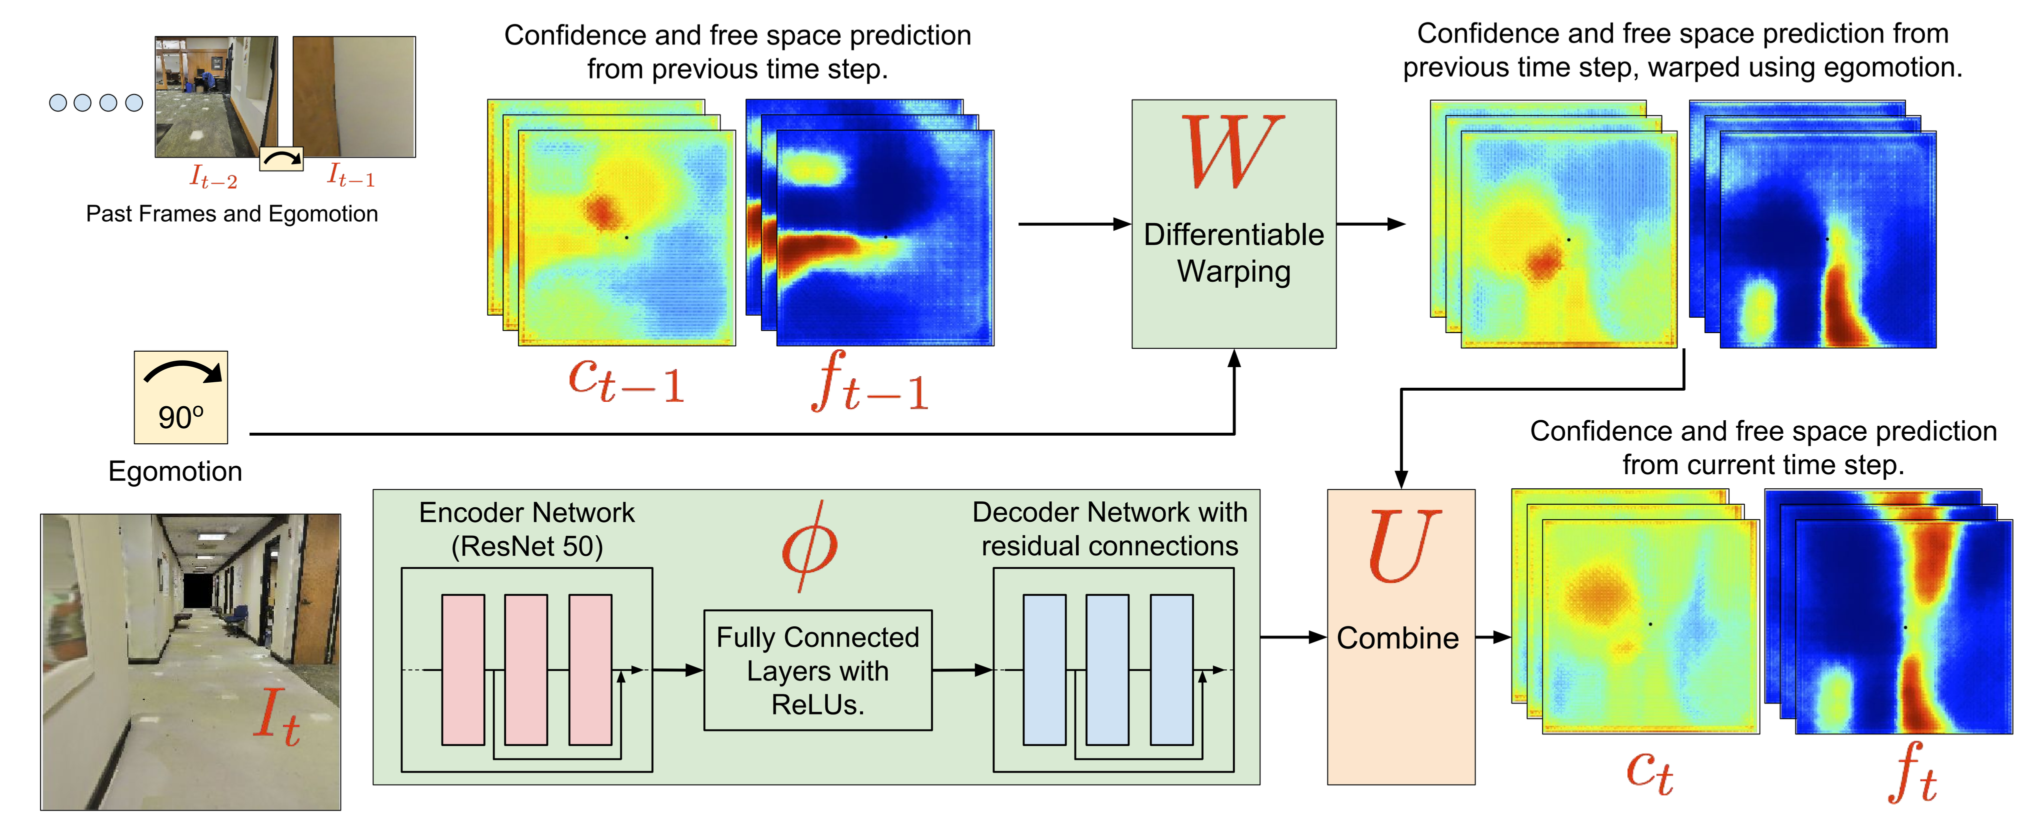

Planner

Express Value Iteration Algorithm as a convolutional neural network making planning trainable, and differentiable [VIN].

$$Q_n(s,a) = R(s,a) + \gamma \sum_{s'} P(s'|s,a) V_n(s')$$

$$V_{n+1}(s) = \max_a Q_n(s,a) \quad \forall s$$

If actions move the agent locally, then can be computed using convolutions

Max Pooling over channels

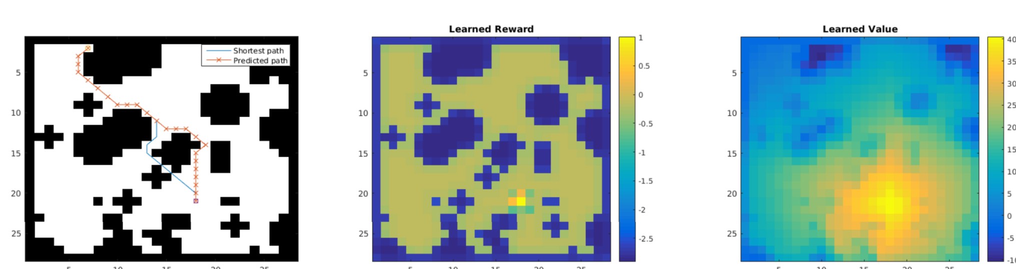

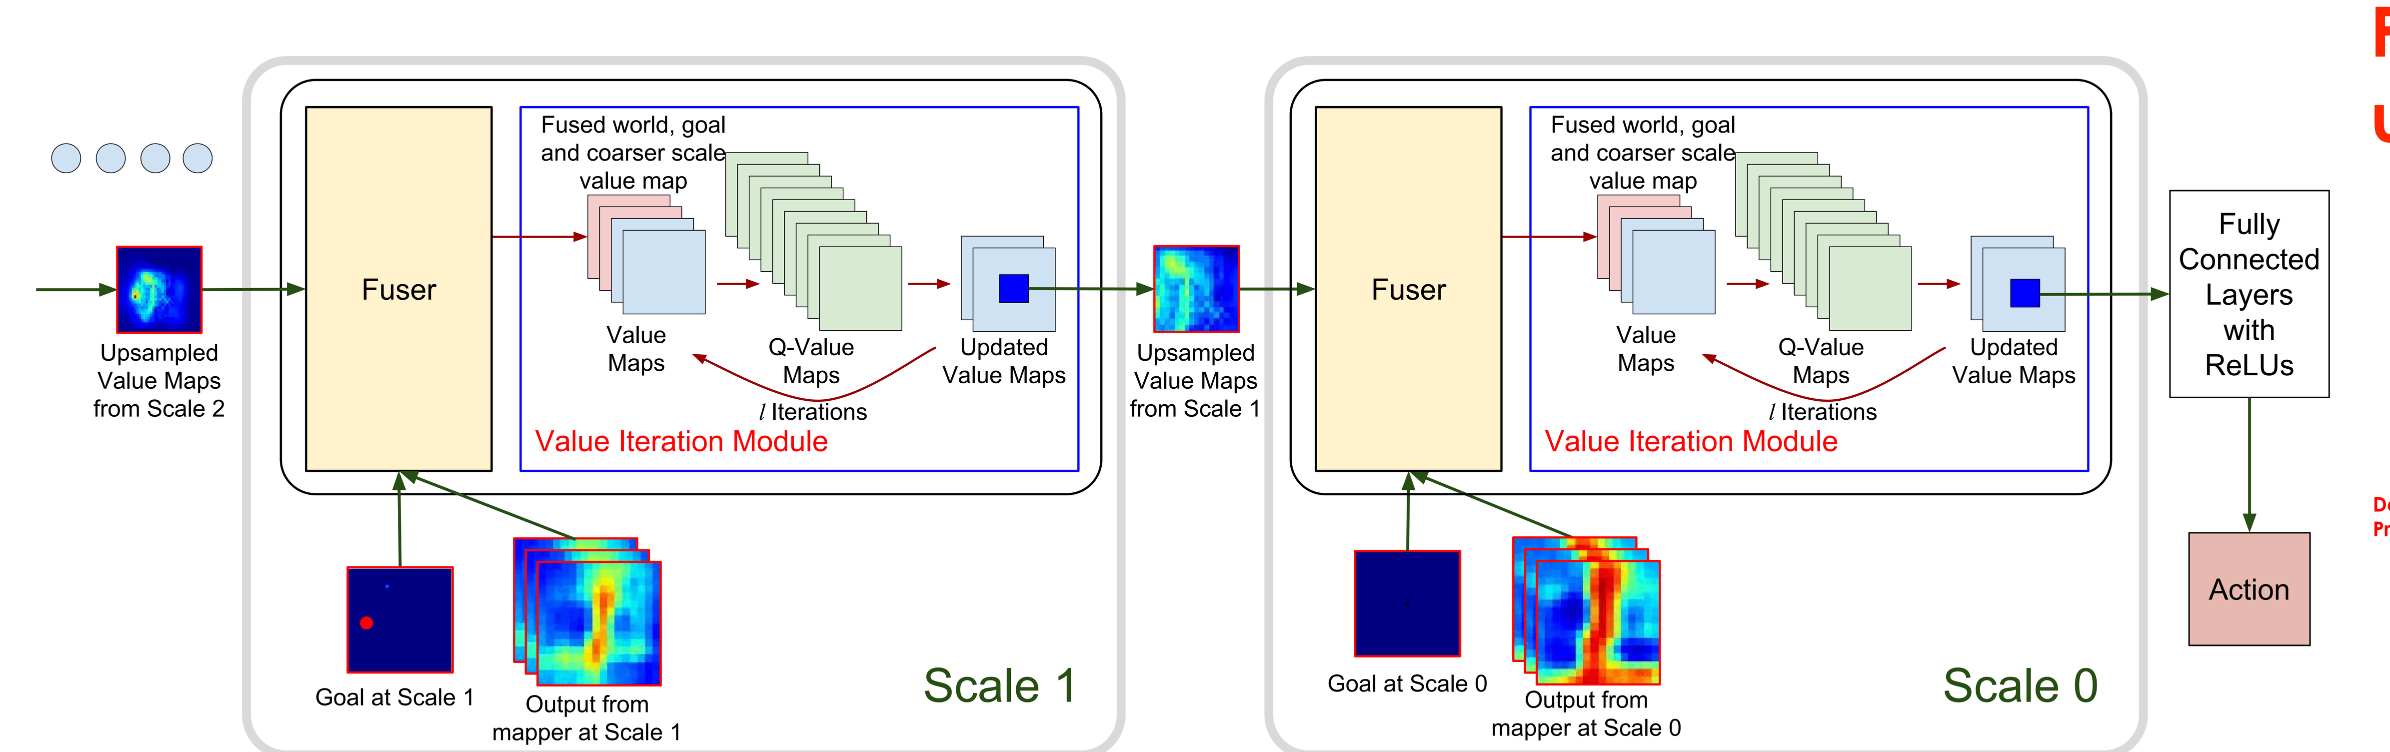

Policy Training using DAGGER

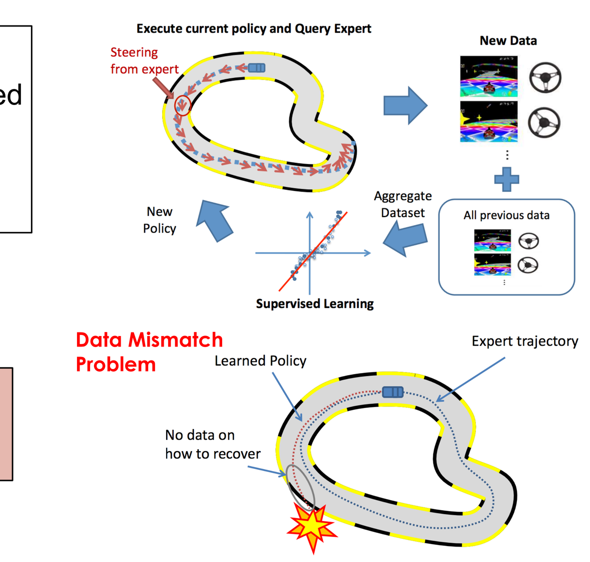

Results

- Trained and tested in static simulated real-world environments.
- Testing environment is different from training environment
- Robot
  - Robot lives in a grid world. Motion is discrete.
  - Robot has 4 macro-actions,
    - Go Forward, Turn left, Turn right, Stay in place.
  - Robot has access to precise ego-motion.
  - Robot has RGB or Depth Cameras
- Geometric Task
  - Goal is sampled to be at most 32 time steps away. Agent is run for 39 time steps
- Semantic Task
  - 'Go to a Chair', agent run for 39 time steps

Value Function Visualization*

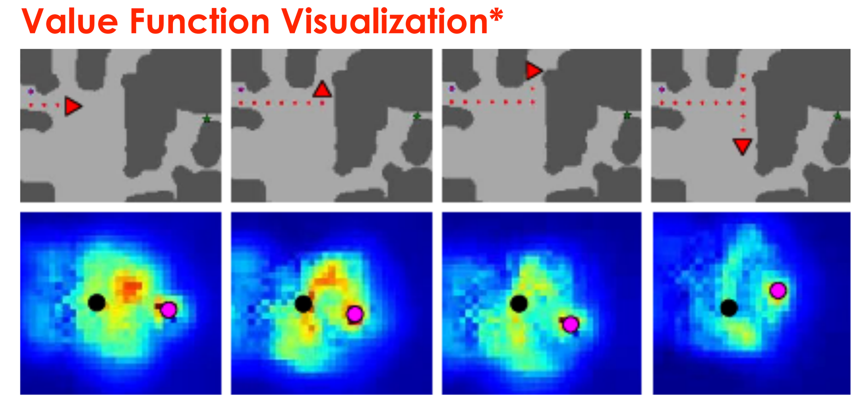

Geometric Task

Methods	RGB Input: Mean Distance	RGB Input: 75th %ile Distance	RGB Input: Success Rate (in %)	Depth Input: Mean Distance	Depth Input: 75th %ile Distance	Depth Input: Success Rate (in %)
Initial	25.3	30	0.7	25.3	30	0.7
No Image	20.8	28	0.7	20.8	28	0.7
React 1	20.9	28	8.2	17.0	26	21.9
React 4	14.4	25	30.4	8.8	18	56.9
LSTM	10.3	21	53	5.9	5	71.8
Our(CMP)	**7.7**	**14**	**62.5**	**4.8**	**1**	**78.3**
Analytic Map				8.0	14	62.9

Semantic Task

Methods	RGB Input: Mean Distance	RGB Input: %ile Distance 50th	RGB Input: %ile Distance 75th	RGB Input: Success Rate (%)	Depth Input: Mean Distance	Depth Input: %ile Distance 50th	Depth Input: %ile Distance 75th	Depth Input: Success Rate (%)
Initial	16.2	17	25	11.3	16.2	17	25	11.3
React 4	14.2	14	22	23.4	14.2	13	23	22.3
LSTM	13.5	13	20	23.5	13.4	14	23	27.2
Our(CMP)	**11.3**	**11**	**18**	**34.2**	**11.0**	**9**	**19**	**40.0**

Successful Navigations

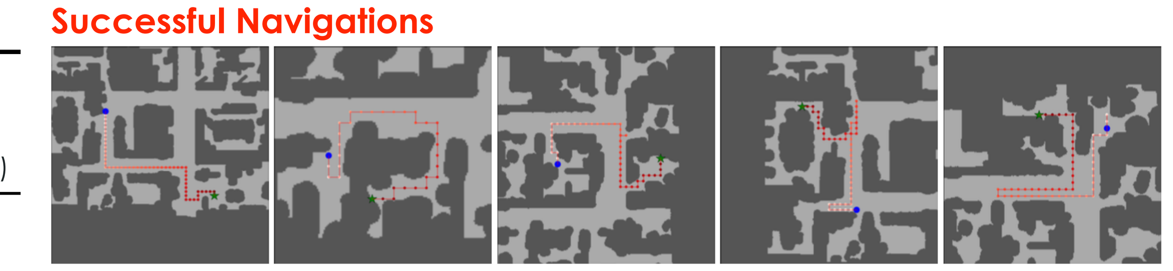

Backtracking

Failed Navigations

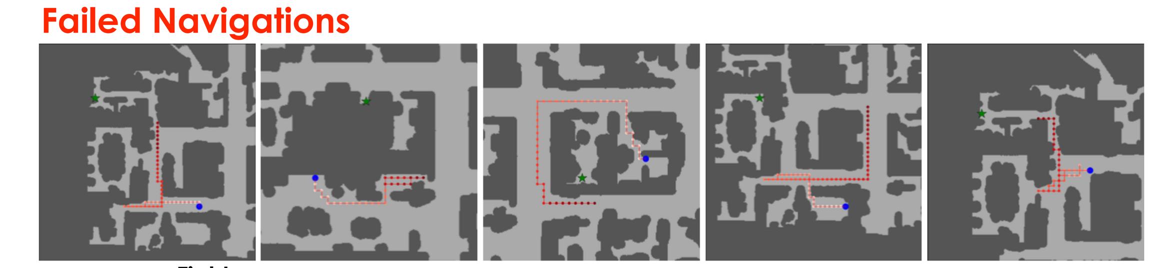

Tight spaces Missed entrances Thrashing

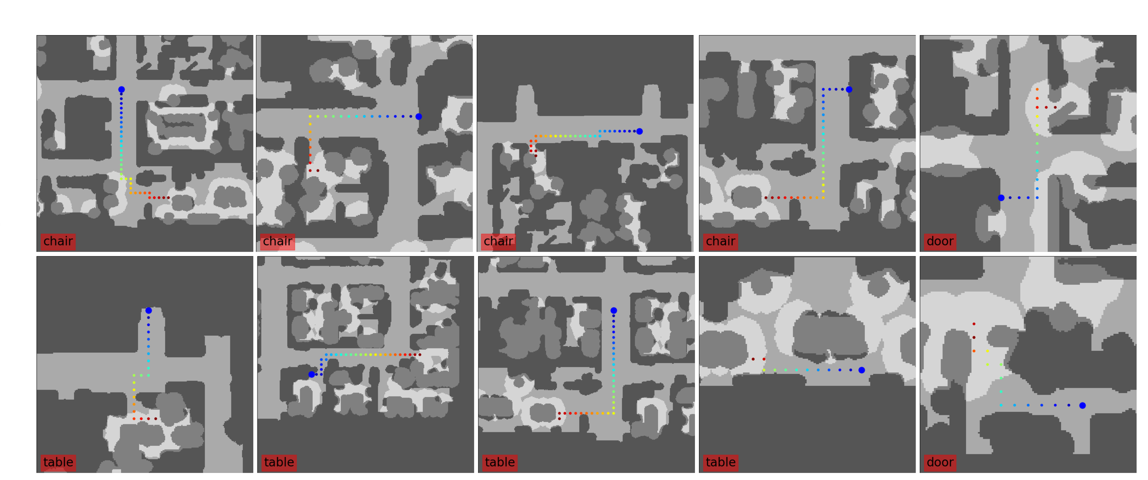

Read Out Mapper Representation

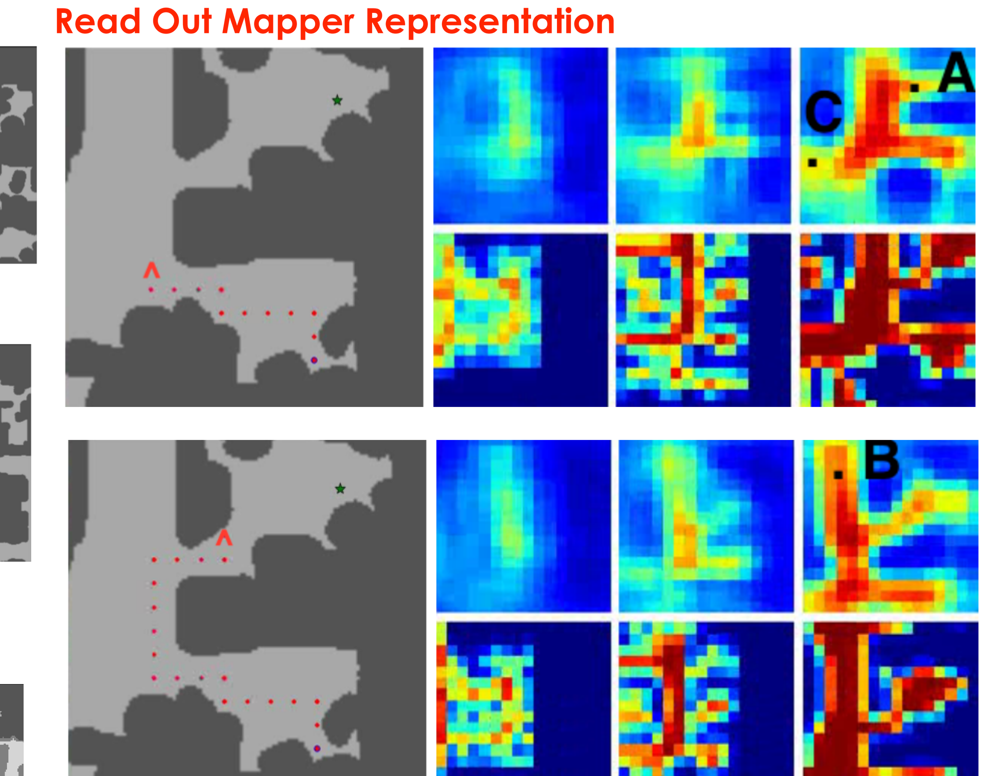

[VIN] Value Iteration Networks. Tamar, Wu, Thomas, Levine, and Abbeel. NIPS 2016.

3D semantic parsing of large-scale indoor spaces. Armeni, Sener, Zamir, Jiang, Brilakis, Fischer, Savarese. CVPR 2016.

A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. Ross, Gordon & Bagnell. AISTATS 2011.