Imitating Shortest Paths for Visual Navigation with Symmetric Siamese Models

Long-Huei Chen, Pratyaksh Sharma, Mohana Prasad

Problem Statement

- **Goal:** Target driven visual navigation with agents trained on simulated virtual environments.
- **Dataset:** Realistic 3D simulator of indoor scenes from CVPR THOR challenge.

Approach

- We propose 3 extensions to Target-driven Siamese model (Yuke et al. 2017) to improve model performance and training efficiency.

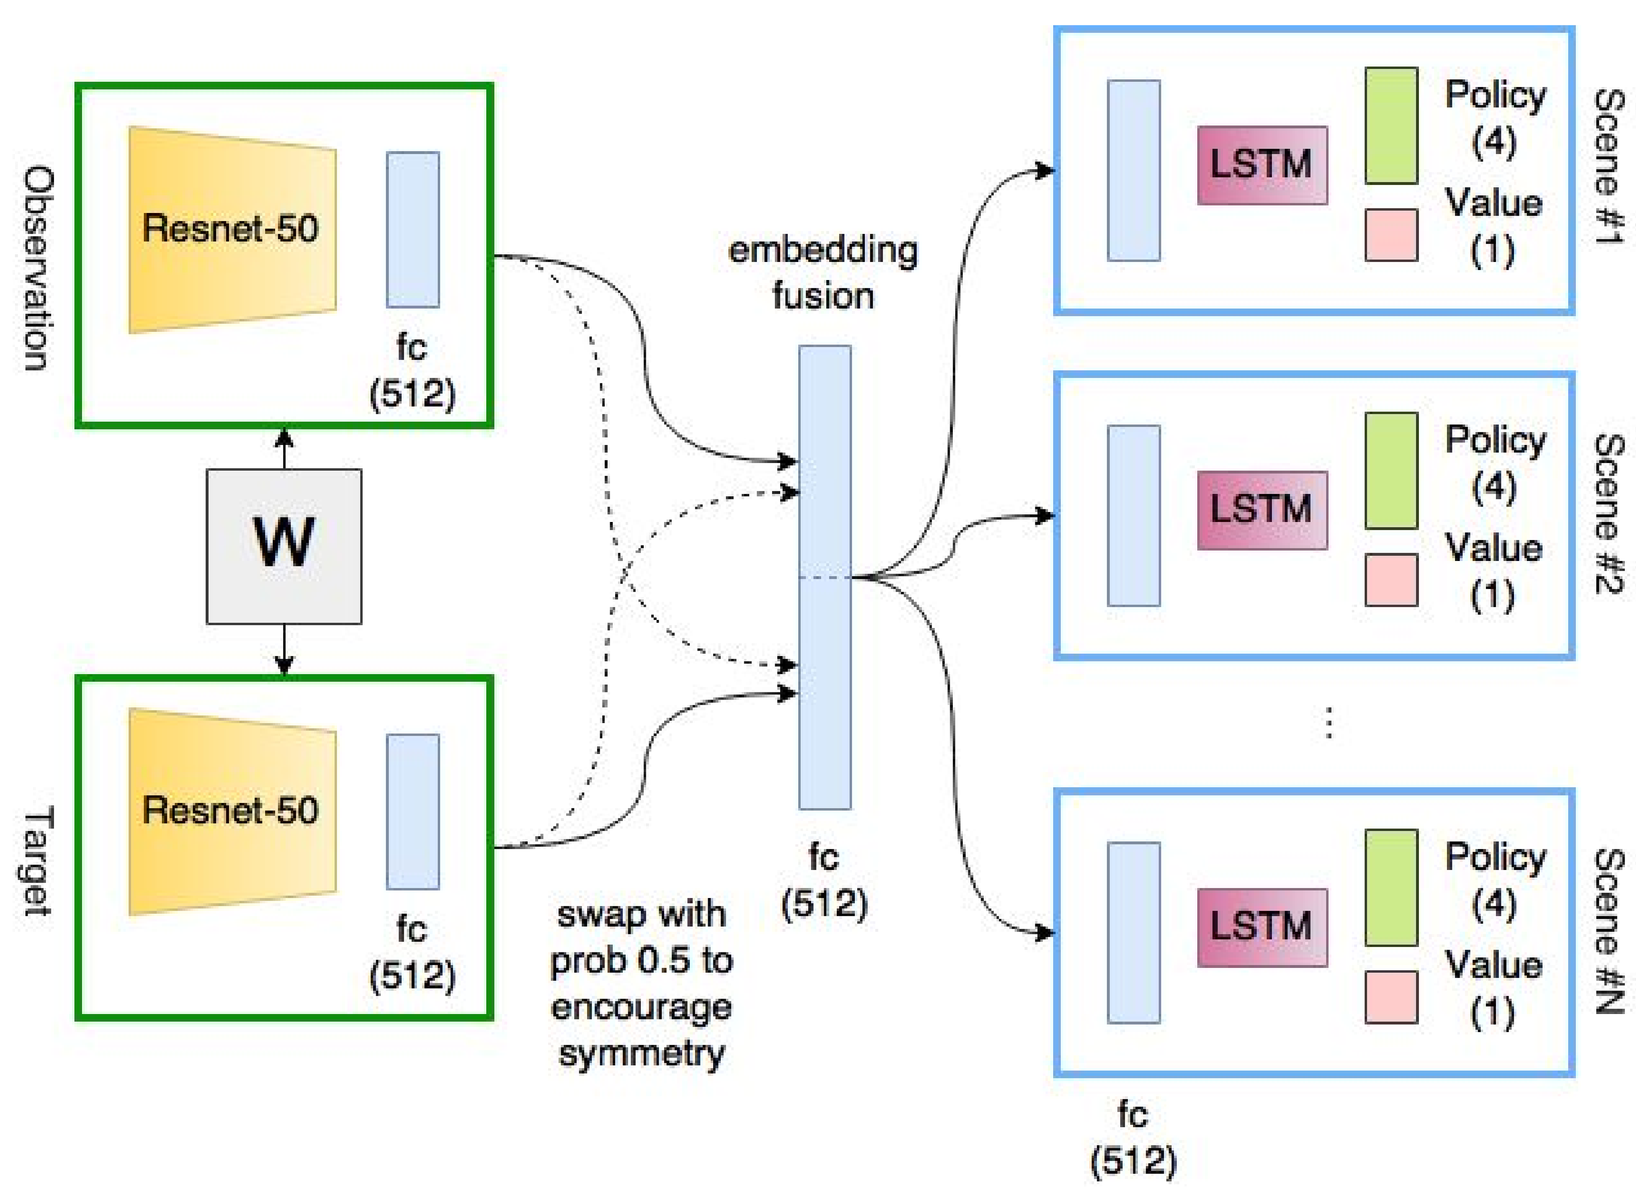

(1) LSTM Extension: Trajectory-aware Memory

- Implicitly learned memory through LSTM cells does not lead to significant improvements over explicit fixed memory (of last 4 observations).

Average path length (100 trials) for 5 targets from 4 scenes

Model \ Scene	Bathroom	Kitchen	Living room	Bedroom
Baseline (Yuke et al.)	7.46	21.54	15.52	14.46
Scene-specific LSTM	92.94	401.57	437.96	548.63

(2) DAagger: Imitation Learning

- Expert supervision of shortest paths enables model to learn shorter paths with lesser training (50x efficiency).

Average path length (100 trials) for 100 targets from 20 scenes

Type	Method	Avg. Trajectory Length
Heuristic	Random Walk	2744.3
	Shortest Path	17.6
Purpose-built RL	1-step Q	2539.2
	A3C	723.5
Target-driven RL	Yuke et al.	210.7
Supervised Learning	**Yuke et al. + DAgger**	**52.7** (oracle:13.1)

(3) Model Symmetry: Generalization

- Encouraging current state/target symmetry with the DAgger model significantly outperforms baseline Target-driven Siamese model on navigating to untrained targets.

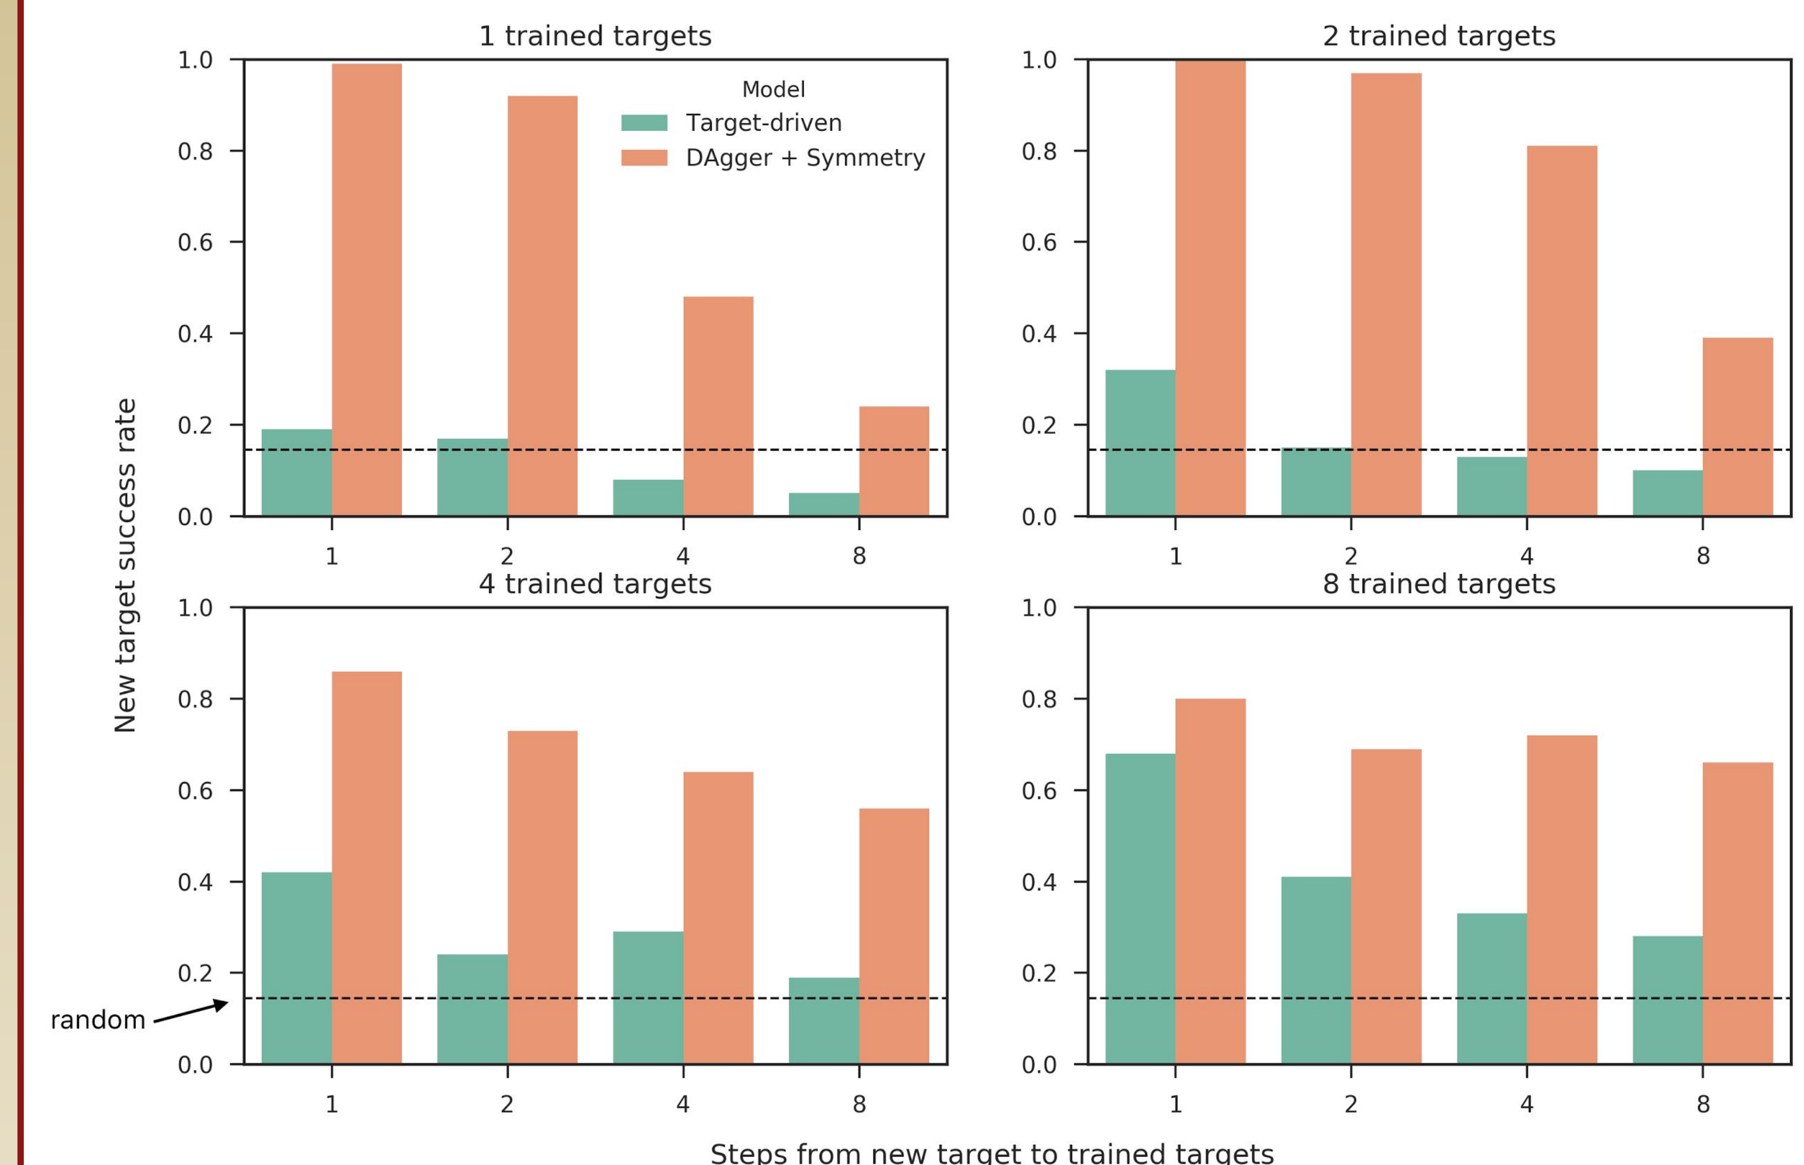

Conclusion & Future Work

(1) **While the model does not benefit from a LSTM extension at the scene specific layer**, more sophisticated memory architectures (like retaining external memories) can be applied.

(2) **Expert supervision of shortest paths speeds up learning and signficantly improves path finding ability.** However, how the model work with imperfect real-world noisy expert supervision (incorrect path estimates, euclidean shortest paths etc.)

(3) **Encouraging symmetry improves target generalization within scenes.** We'd like to explore how the same strategy can benefit learning completely new scenes.

Acknowledgements

We are grateful to Yuke Zhu, Emma Brunskill and Rishi Bedi for their guidance and support throughout the development of this work