On Embodied Visual Navigation in Real Environments Through Habitat

Rosano M.[1,3], Furnari A.[1], Gulino L.[3], Farinella G. M.[1,2]

[1]FPV@IPLAB – Department of Mathematics and Computer Science, University of Catania, Italy
[2]Cognitive Robotics and Social Sensing Laboratory, ICAR-CNR, Palermo, Italy
[3]OrangeDev s.r.l, Firenze, Italy
marco.rosano@unict.it, furnari@dmi.unict.it, luigi.gulino@orangedev.it, gfarinella@dmi.unict.it

Abstract

Deep visual navigation models can learn effective policies when trained on large amounts of visual observations through reinforcement learning. To overcome the limited realism of simulators, we propose a tool based on the Habitat simulator which exploits real world images of the environment, together with sensor and actuator noise models, to produce more realistic navigation episodes. We show that our tool can effectively help to train and evaluate navigation policies on real-world observations without running navigation episodes in the real world.

1. Virtual World vs. Real World

Navigation policies trained in simulation do not transfer to real world:

- Differences in the appearance;
- No sensor and actuator noise;
- Simplified physical interaction.

3. Dataset Collection and 3D alignment

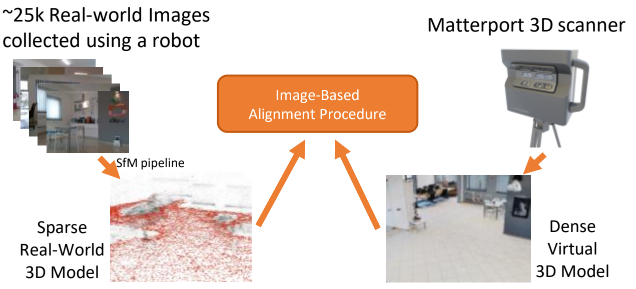

5. Impact of Sensor and Actuator Noise

Sensors noise	Actuators noise	Trained with noise	SPL
No	No	No	0.9127
Small	No	No	0.8173
		Yes	0.8658
Medium	No	No	0.5075
		Yes	0.7114
Large	No	No	0.1552
		Yes	0.3643

- Train in presence of noisy sensors and actuators produced a more robust navigation model;
- Imperfect odometry has a significative impact in performances

Dataset, Code and demo

We released our virtual and real 3D models, the code and additional videos. Check out our project page:

https://iplab.dmi.unict.it/EmbodiedVN

2. Real world Adaptation Framework

The proposed tool allows to generate paired virtual and real-world navigation episodes by sampling visual observations from the aligned 3D model. Sensor and actuator noise models can be provided to enable the simulator to generate more realistic navigation episodes.

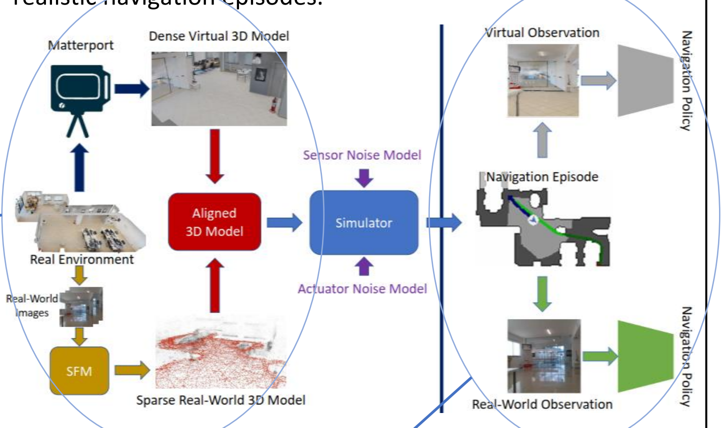

4. Training Setup and Performance

- PointGoal navigation task, updated at each timestep;
- Discrete action space *(0.25m, 10°, STOP)*;
- Trained using the PPO Reinforcement Learning algorithm;
- Pretrained model on Gibson and Matterport 3D datasets;
- Trained on different combination of virtual and real-world images;
- Tested on real-world observations;
- CycleGAN as unsup. domain adaptation baseline;

Training Stages	SPL	Success rate
Virtual	0.0160	0.022
CycleGAN	0.2464	0.3310
Virtual+CycleGAN	0.2648	0.3410
Real	0.7112	0.8590
Virtual+real	0.8001	0.9700
CycleGAN+real	0.7665	0.8880
Virtual+CycleGAN+real	0.7553	0.9360

- Results suggest that knowledge learned in a simulated environment can be transferred to real-world observations using appropriate tools.