SEQUENCE-AGNOSTIC MULTI-OBJECT NAVIGATION

Nandiraju Gireesh*[1], Ayush Agrawal*[1], Ahana Datta*[1], Snehasis Banerjee[1,2],
Mohan Sridharan[3], Brojeshwar Bhowmick[2], Madhava Krishna[1]

[1]Robotics Research Center, IIIT Hyderabad, India
[2]TCS Research, Tata Consultancy Services, India
[3]Intelligent Robotics Lab, University of Birmingham, UK

ABSTRACT

This paper presents a novel approach for Multi-Object Navigation (MultiON) task using a deep reinforcement learning framework. Our approach is sequence-agnostic, hence it does not rely on a pre-determined sequence of object classes to be explored. We use an actor-critic architecture and a reward specification that rewards the agent's progress towards individual and multiple object classes. Results show that our method outperforms pre-sequenced (PSM) and state-of-the-art ON method when used for MultiON task (M-SemExp) in a 3D simulation environment.

METHODOLOGY

The methodology proposed has the following three modules:

Semantic Mapping: To build a semantic map from the segmented RGB-D and pose observations. It leverages data augmentation techniques (random shift) for generalization.

Encoder Network: Receives the semantic map from the previous module along with the object goal list. It then extracts high-level features to send to an Actor-Critic network to select the long-term goals.

Deterministic local policy: Uses analytical planners to compute low-level navigation actions to reach the long-term goal.

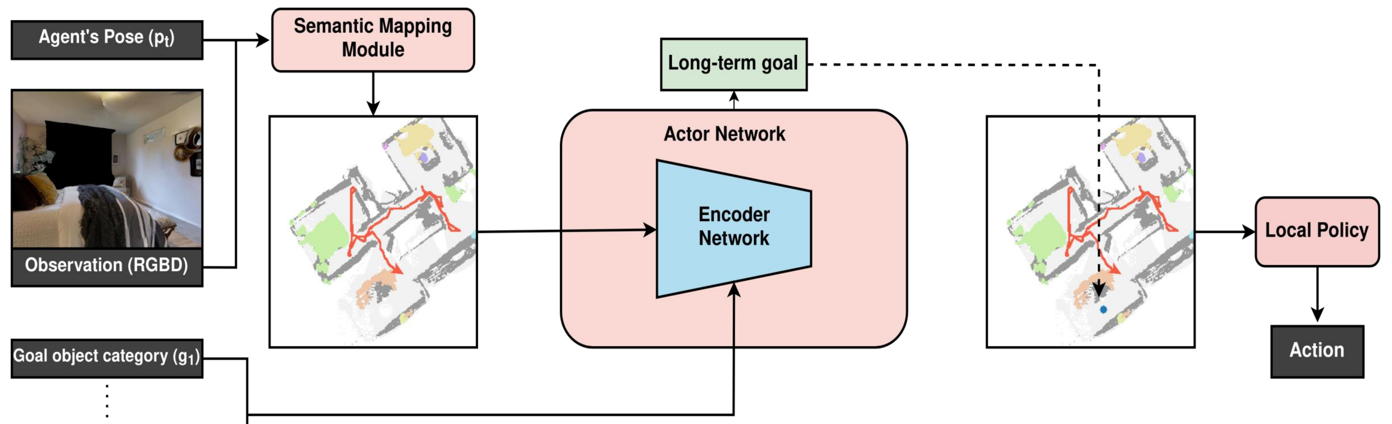

When a long-term goal is provided by the actor-critic network, the local policy module uses the Fast Marching Method(FMM) to guide the robot to this region by utilizing the obstacle channel of the generated semantic map.

Experimental Setup: ObjectNav Challenge's 25 scenes for training by setting up ≈ 1000 episodes (total) of k-ON task, randomly selecting k ∈ [2, 3]; finally testing on 5 scenes.

Hypotheses Supported by this work:

1. Our framework for SAM (sequence-agnostic multi-object navigation) traverses shorter paths and takes fewer timesteps when compared with PSM (pre-sequenced multi-object navigation).
2. Our framework for SAM provides better long-term goals than the state-of-the-art single object navigation baseline extended to multi-object settings.

REWARD STRUCTURE

The reward function formulated captures the following three characteristics:

1. Gives credit for finding a sub-goal, i.e., an instance of each target object class ($R_{sub\text{-}goal}$).
2. Credits the robot for reducing the distance to multiple object classes concurrently ($R_{process}$). $\alpha_{process}$ is the hyperparameter.
3. Penalizes the agent for longer episodes (*CNR* or Constant Negative Reward assigned at every timestep).

$$Reward = R_{sub-goal} + \alpha_{process} * R_{process} + CNR$$

$$R_{sub-goal} = \begin{cases} r_{sub-goal} & \text{if a sub-goal is reached} \\ 0 & \text{otherwise} \end{cases}$$

where $r_{sub\text{-}goal} = 2$ (tuned experimentally)

$$R_{process} = \begin{cases} \frac{n}{N} + d_t & \text{if } dtg \text{ of } n \text{ classes decreases} \\ d_t & \text{otherwise} \end{cases}$$

where d_t is the sum of decrease in geodesic distance to the nearest instance of each of the *n* object classes out of total *N* target classes (*N*-ON task) and *dtg* is the distance of the agent to an instance of each target class.

RESULTS

M-SemExp

Timesteps: 182
Global Path Length: 19.3m

Achieved Sub-Goal	✗	✗	✓	✗	✗	✓
Reduced Distance to All remaining Sub-Goals	✗	✓	✓	✗	✗	✓

SAM Framework

Timesteps : 86
Global Path Length: 11.3m

Achieved Sub-Goal	✓	✓	✗	✓
Reduced Distance to All remaining Sub-Goals	✓	✓	✓	✓

Our model is able to select a better set of long-term goals which ensures an improved Multi-Object Navigation. This enhanced performance is attributed to the reward function that teaches the agent to reduce its distance to multiple object classes simultaneously.

REFERENCES

1. Chaplot et al., "Object goal navigation using goal-oriented semantic exploration," in Neural Information Processing Systems, 2020.
2. Savva et al., "Habitat: A Platform for Embodied AI Research," in International Conference on Computer Vision, 2019.
3. J. A. Sethian, "A fast marching level set method for monotonically advancing fronts." National Academy of Sciences, vol. 93, no. 4, pp. 1591–1595, 1996.

METRICS

We used five performance metrics to evaluate the performance of our solution.

For comparison with PSM:

- Timesteps
- Global Path Length

For comparison with M-SemExp:

- Success (1 *iff* all objects found within maximum allowed timesteps)
- Sub-success (fraction of objects found)
- Global-SPL (ratio of shortest possible path to the actual path taken).

COMPARATIVE ANALYSIS

Vs Pre-sequenced setting (PSM)

Scene Name	Timesteps ↓		Global Path Length (m) ↓	
	PSM	SAM	PSM	SAM
Collierville	242	**122**	29.53	**16.16**
Corozal	336	**179**	46.23	**27.28**
Darden	248	**117**	31.43	**16.08**
Markleeville	272	**140**	35.41	**18.87**
Wiconisco	389	**224**	52.81	**33.94**

Performance of SAM formulation compared with the PSM formulation, with numbers averaged over 200 paired episodes for each previously unseen scene.

Vs M-SemExp

Method	Success (%) ↑		Sub-success (%) ↑		G-SPL (%) ↑	
	2-ON	3-ON	2-ON	3-ON	2-ON	3-ON
Random	3.3	4.7	11.5	14.2	0	0
M-SemExp	60.5	61.7	73.1	76.6	30.5	29.8
Ours	**70.7**	**72.3**	**82.5**	**86.9**	**39.3**	**39.3**

SAM framework compared with the Random and M-SemExp baselines, averaged over the 200 episodes of each of the five testing scenes (unseen scenes).

FUTURE WORK

Future work includes application on a physical robot and including clustering as the number of objects increases. So far we have experimented with Household scenes, however this project would find great relevance even in other types of scenes like: Supermarket or Hospital.

Presenting author: ahana.datta@research.iiit.ac.in