Teaching robots to explore unseen environments

Kamalesh Palanisamy

Saikat Ghosh

Why Exploration?

We cannot always provide a robot the map to an environment.

We need to come up with an efficient method that allows the agent (in this case a physical robot) to build a map of the environment.

Once built, map can be used for various downstream tasks like object navigation, manipulation etc.

Current trend

Reinforcement learning can help the robot achieve this goal because it will enable us to learn the intrinsic characteristic of exploration.

We can train these intrinsic characteristics by designing reward functions that enable the robot to learn them.

Some of these rewards could be curiosity, map coverage(we can use the map for training :)), reconstruction of views etc.

Related Work

An Exploration of Embodied Visual Exploration

The policy architecture consists of three parts:
1. **Spatial memory**
2. **Temporal memory**
3. **Actor-Critic model**

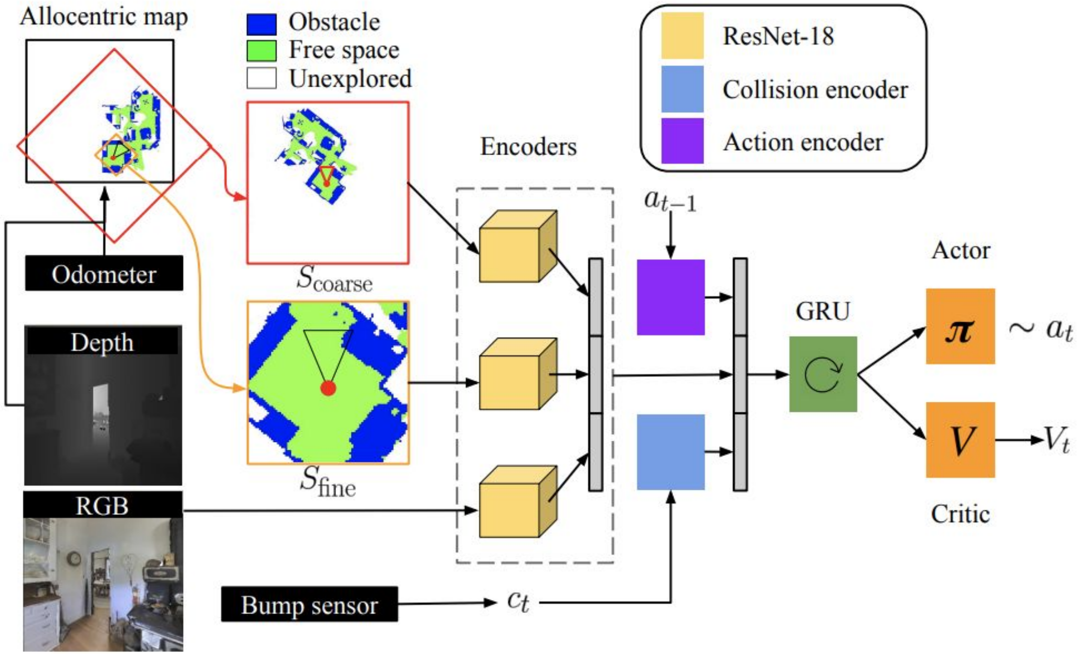

Related Work

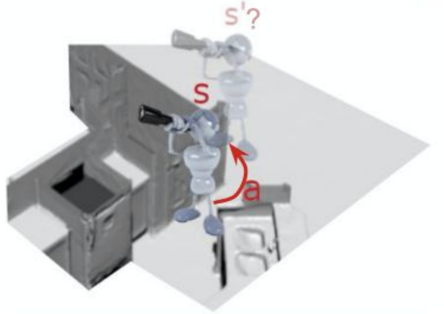

Curiosity

$$r_t \propto ||\hat{s}_{t+1} - s_{t+1}||_2^2.$$

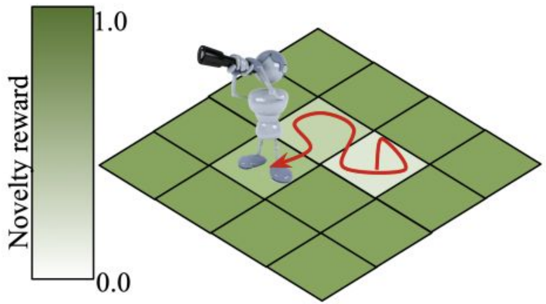

Novelty

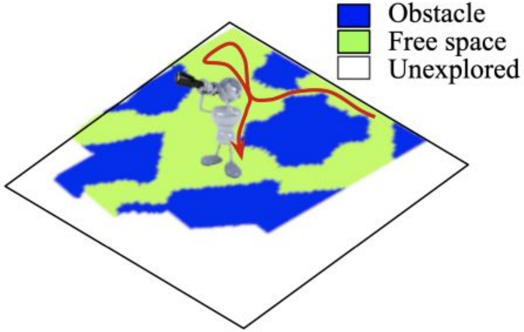

Coverage

Reconstruction

Related Work

Doesn't perform really well? Looks like frontier beats it on most tasks? Who is frontier?

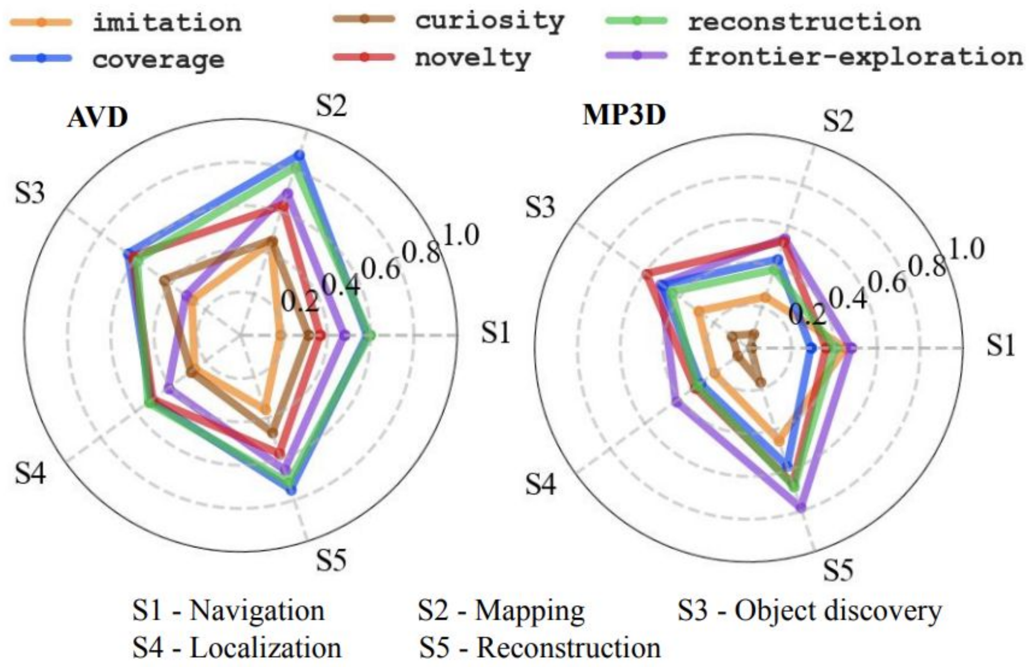

Related Work - Frontier

Frontier is a more traditional algorithm that does not use any of the fancy neural networks.

Purely, based on building a global occupancy map and identifying “frontiers” between free space and unknown space.

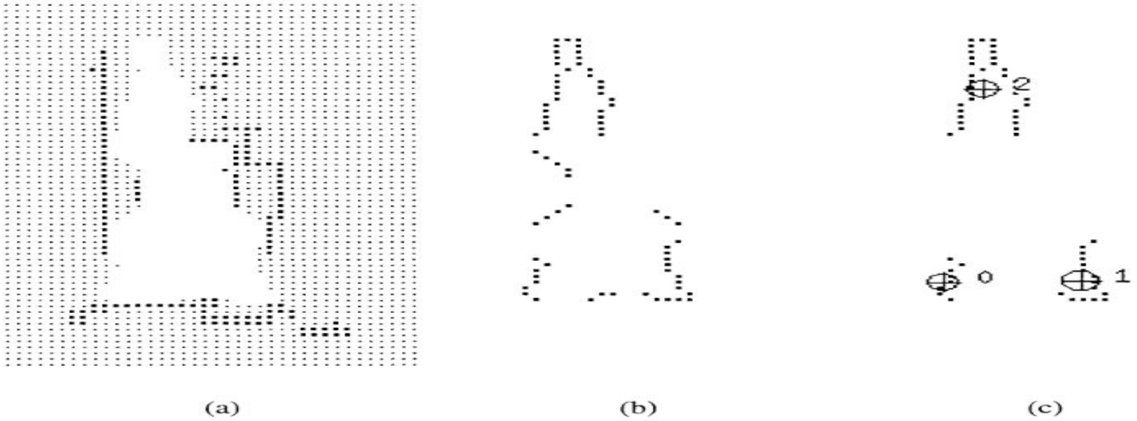

Figure 1: Frontier detection: (a) evidence grid, (b) frontier edge segments, (c) frontier regions

Frontier

Works well in fact performs much better than RL methods. But has a few drawbacks:

1. Since frontier is based purely on LIDAR, so it fails heavily in glass buildings like ECS-West :(.
2. In huge spaces like an airport, it takes a large amount of computation.
3. Additionally, it requires a map at all times (Is this bad?!)

Datasets and Simulator

Datasets

MatterPort3D	Gibson V1	Gibson V2/iGibson

Simulator

Habitat Simulator	iGibson Env/Simulator

Our implementation uses **MatterPort3D + iGibsonEnv**.

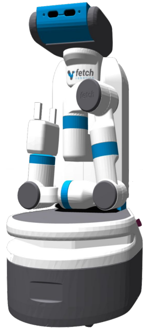

Our method

Design a reward so that the robot can learn to explore rooms. Why rooms though?

Rooms represent different “frontiers” and making the robot understand the difference between the kitchen and the bedroom can help it to explore another room. Design a reward that gives positive values when it visits a previously unseen room and negative if the room was already visited.

How to get this reward?

We used contrastive learning to train a neural network to associate different images from the same room as positive pairs and different rooms as negative pairs.

Flowchart

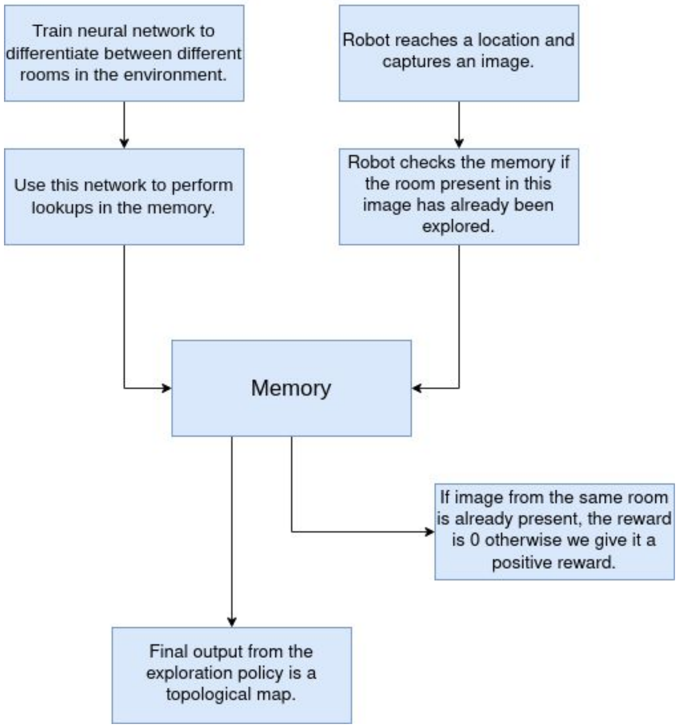

Failed miserably!

Old is gold

It is clear that the RL policy needs to have some pretraining before it can trained with the “room” reward.

Imitation learning is a common pretraining strategy used by other papers. But who should we imitate?

Frontier!

We have implemented frontier in the iGibson simulator. We have modified some parts of the core frontier algorithm. (Not as easy as it sounds)

We use depth images along with LIDAR to improve obstacle detection.

Imitation learning is currently on the works!

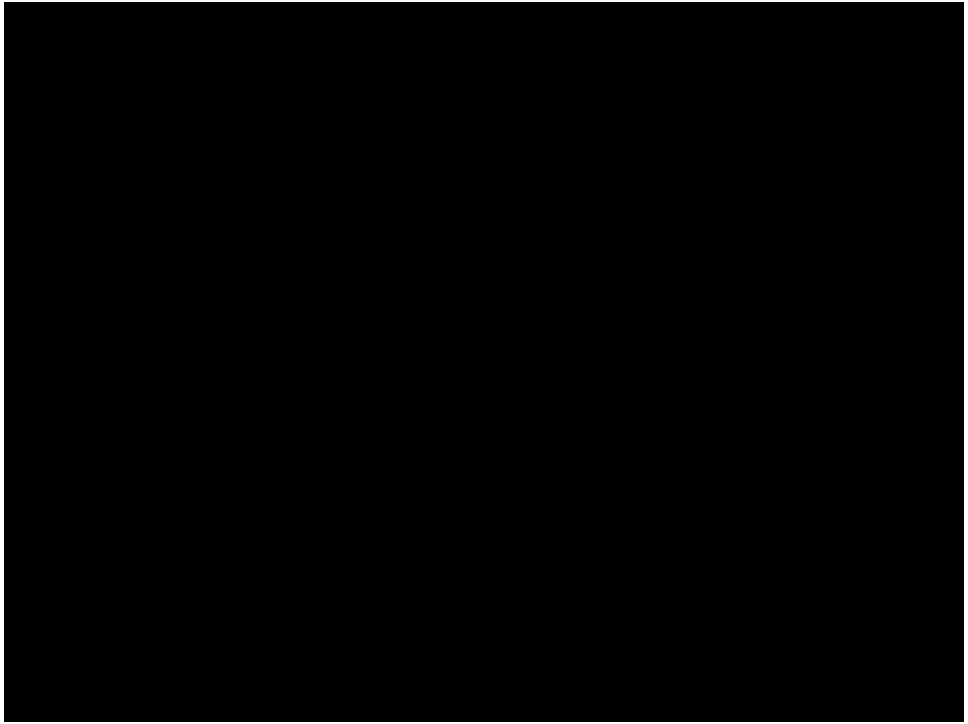

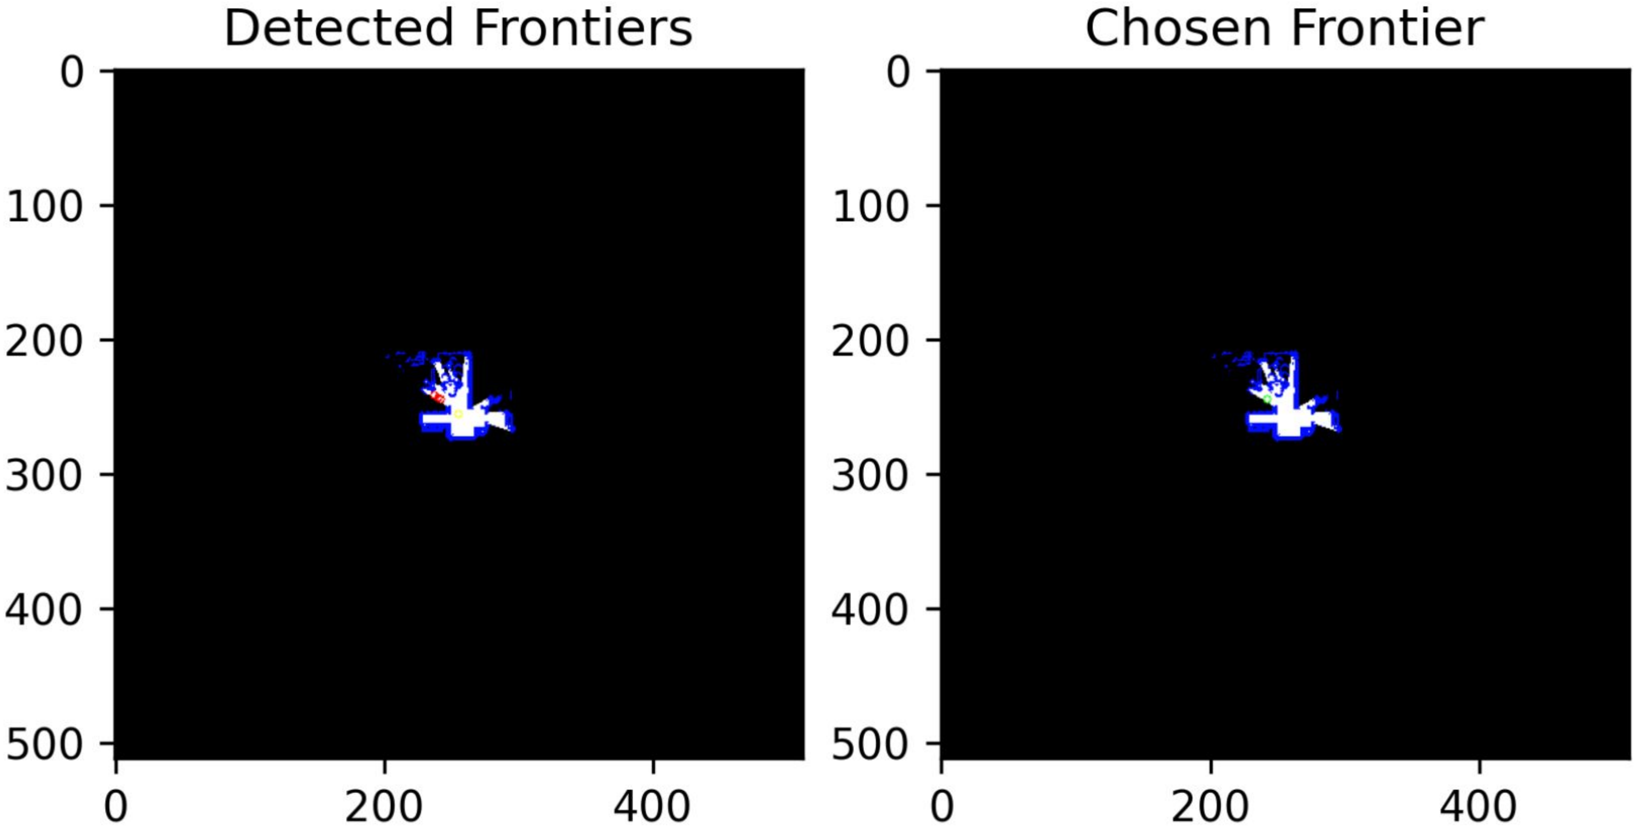
Detected Frontiers
Chosen Frontier
0
100
200
300
400
500
0
200
400

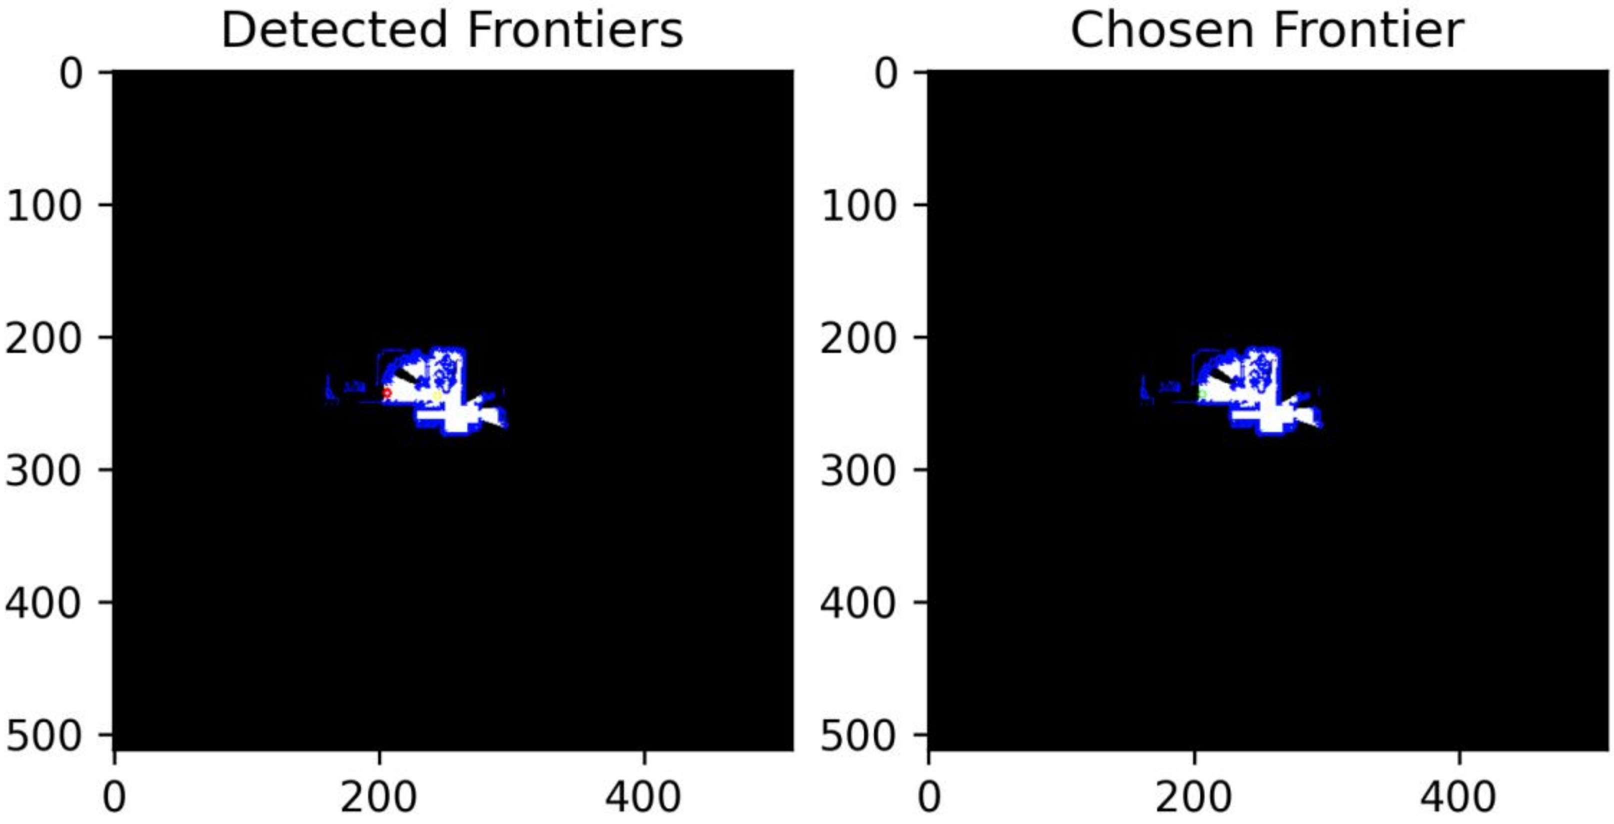

Detected Frontiers
Chosen Frontier
0
100
200
300
400
500
0
200
400

Questions?