Introduction

Task-independent exploration of new environments.

- Classical navigation:
 - Ignore exploration and rely on a human operator for map building.
 - Ignore semantics, sensitive to noise, crucially rely on specialized sensors.
- Learning-based navigation:
 - Only study training environments.
 - Entirely task-driven.

Approach

- Reconstruct the 3D world as the agent moves.
- Project the point cloud onto the ground plane to get the coverage map M_t .
- Calculate an intrinsic coverage reward:

 $R_{int}^{cov}(t) = C(M_{t+1}) - C(M_t)$ $R_{int}^{coll}(t) = -Bump(t+1)$

 $R(t) = \alpha R_{int}^{cov}(t) + \beta R_{int}^{coll}(t)$

- Initialize policy from imitation learning.
- Finetune policy using RL.

Learning Exploration Policies for Navigation Abhinav Gupta¹² Saurabh Gupta² Tao Chen¹ ¹CMU Robotics Institute ²Facebook AI Research (FAIR)



