An Exploration of Embodied Visual Exploration

Project page: http://vision.cs.utexas.edu/projects/exploring-exploration

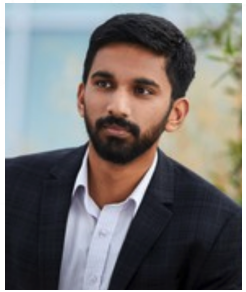

Santhosh Kumar Ramakrishnan[1,3]

[1] The University of Texas at Austin

Dinesh Jayaraman[2,3]

[2] University of Pennsylvania

Kristen Grauman[1,3]

[3] Facebook AI Research

Embodied visual exploration

Egocentric view

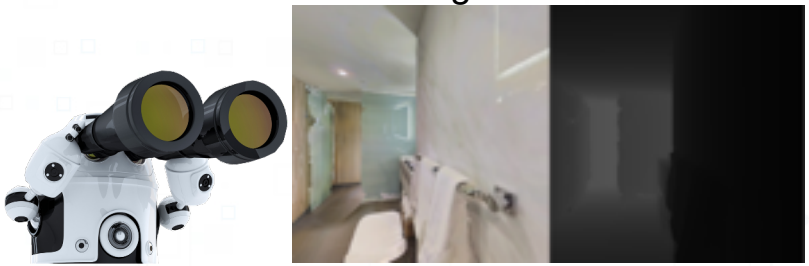

3D environment

Image credits: Matterport3D dataset, Chang et al., 3DV 2017

Embodied exploration

Egocentric view

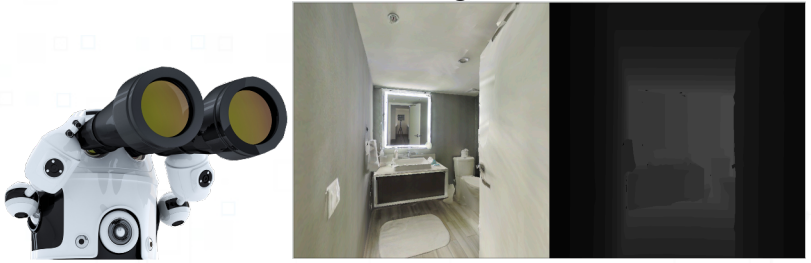

3D environment

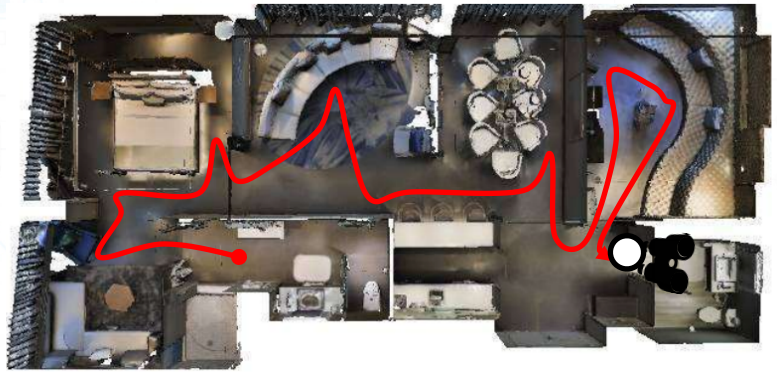

Automatically gather information

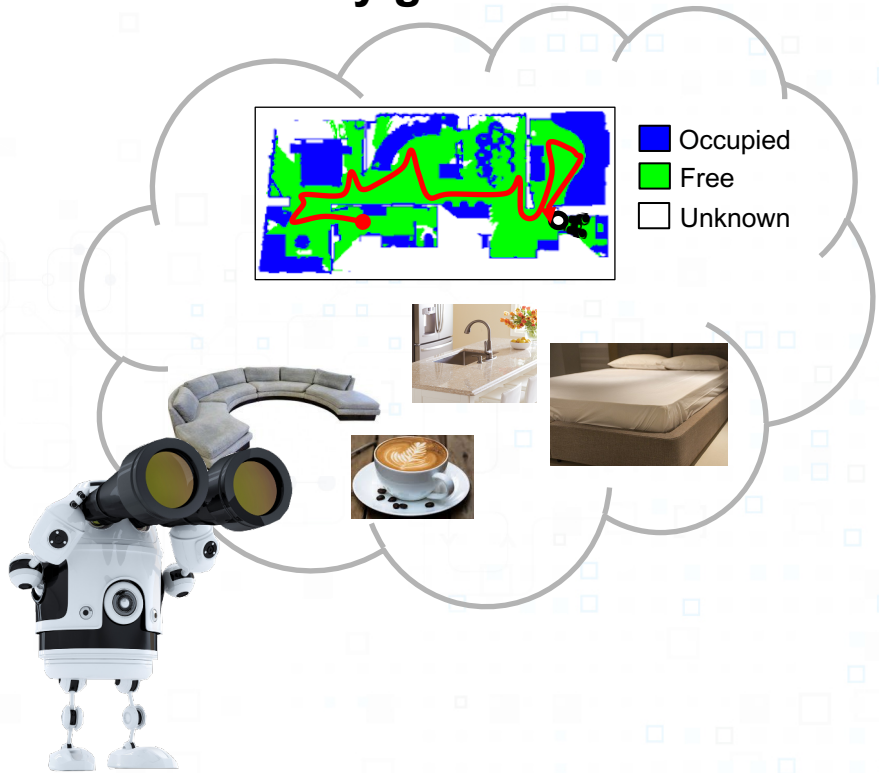

Image credits: Matterport3D dataset, Chang et al., 3DV 2017

Embodied exploration

Automatically gather information

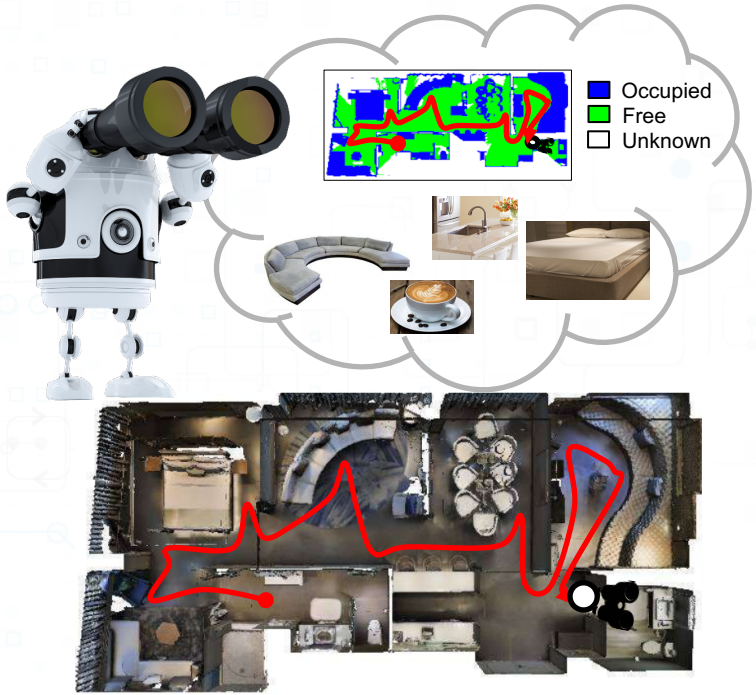

Solve downstream tasks

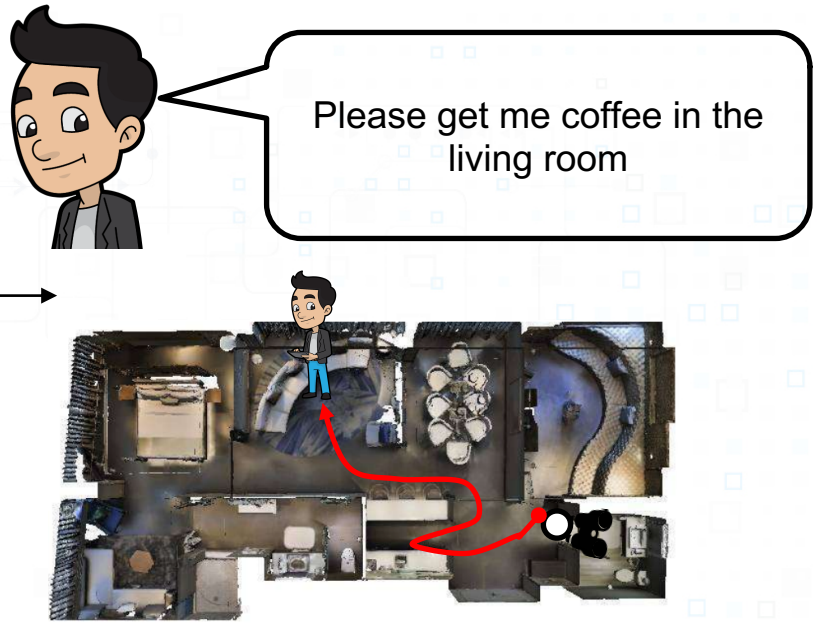

Image credits: Matterport3D dataset, Chang et al., 3DV 2017

Key questions

- Is the spatial exploration problem solved?
- How well do recent approaches compare with each other?
- What are the strengths and weaknesses of different approaches?

Lack of standardized experimental conditions

Architectures, algorithms

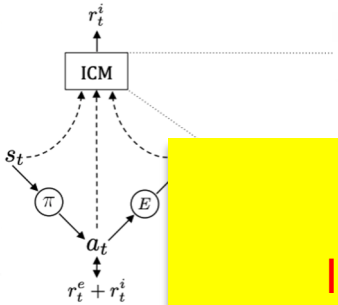

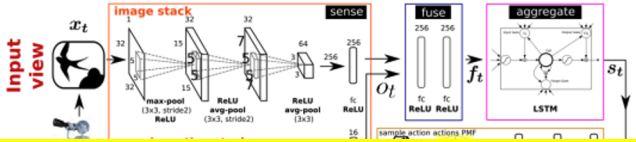

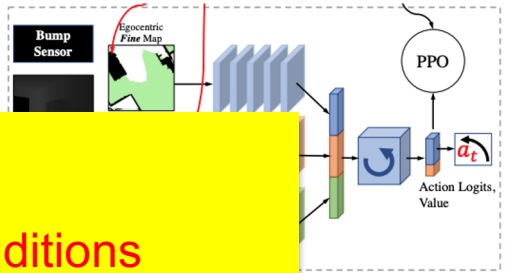

Pathak et al., *LR 2019*

Inconsistencies in experimental conditions
=> Unclear what method works when and why.

etrics

GIBSON

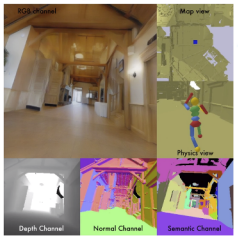

Matterport3D

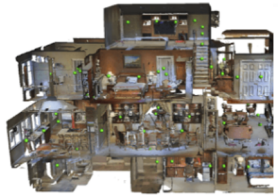

VizDoom

- *area covered*
- *objects covered*
- *reconstruction*
- *navigation*
- *overcoming sparse rewards*

Our contributions

Exploration taxonomy

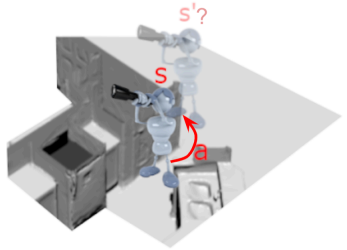

Curiosity

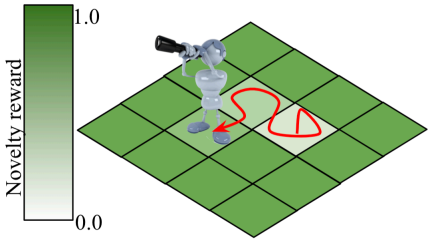

Novelty

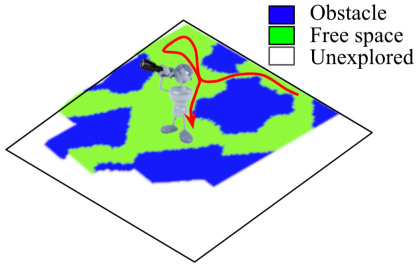

Coverage

Reconstruction

Our contributions

Exploration taxonomy

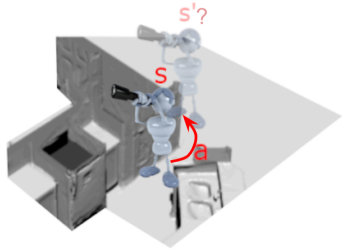

Curiosity

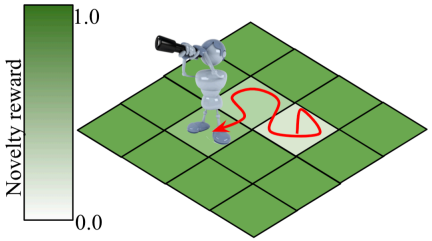

Novelty

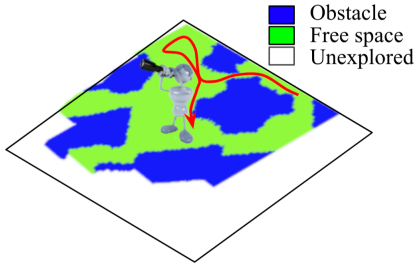

Coverage

Reconstruction

Exploration study framework

- High-quality 3D environments
- State-of-the-art policy architecture
- Diverse evaluation metrics

Our contributions

Exploration taxonomy

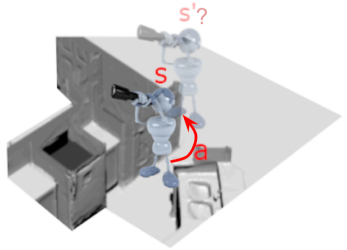

Curiosity

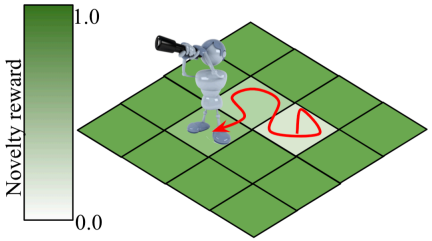

Novelty

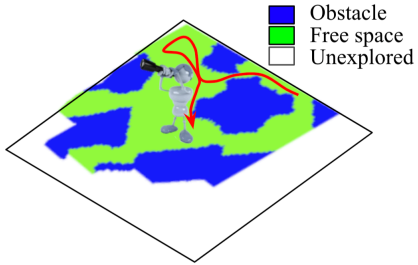

Coverage

Reconstruction

Exploration study framework

- High-quality 3D environments
- State-of-the-art policy architecture
- Diverse evaluation metrics

Exploration taxonomy - novelty

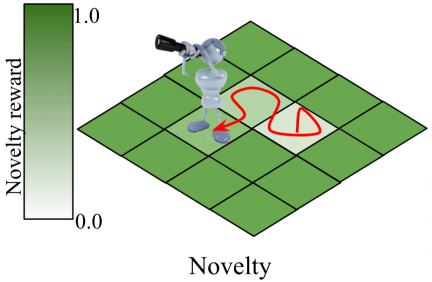

- Rewards visiting states that are infrequently visited

$$r(s) \propto \frac{1}{\sqrt{n(s)}} \qquad n(s) = \text{number of times state } s \text{ has been visited}$$

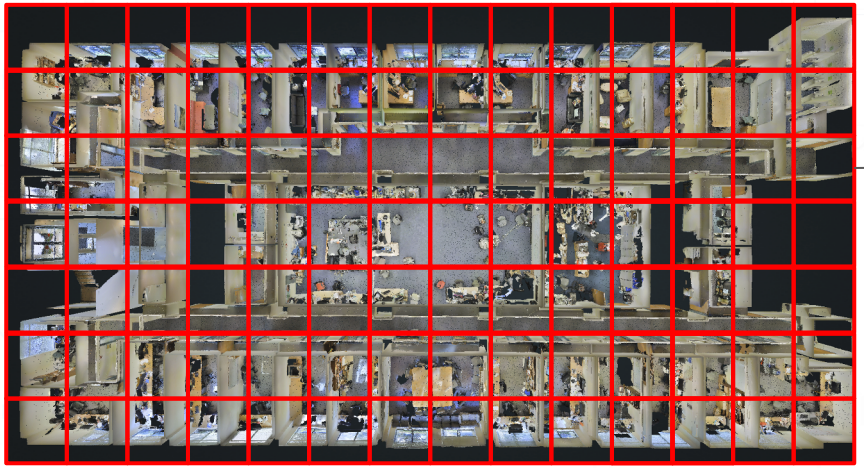

Discretize the world into grids that represent states

Unifying count-based exploration and intrinsic motivation, Bellemare et al, NeurIPS 2016
 # exploration: A study of count-based exploration for deep RL, Tang et al., NeurIPS 2017
Episodic curiosity through reachability, Savinov et al., ICLR 2019

Exploration taxonomy - curiosity

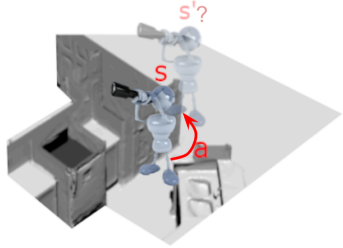

Curiosity

- Learns a forward dynamics prediction model
- Agent is rewarded for visiting states where this model is poor at prediction, i.e., do things you are curious about

$$r(s) \propto ||s' - f(a, s)||^2$$

Exploration in model-based reinforcement learning by empirically estimating learning process, Lopes et al., NeurIPS 2012
Curiosity-driven exploration by self-supervised prediction, Pathak et al., ICML 2017
A large-scale study of curiosity-driven learning, ICLR 2019

Exploration taxonomy - coverage

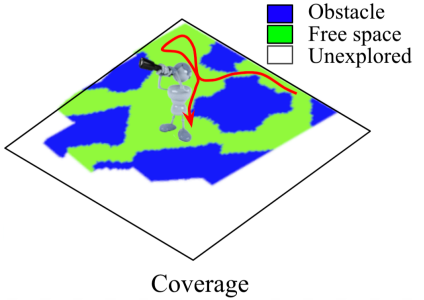

- Maximize the amount of area seen during exploration (A)

$$r(s) = \Delta A$$

Exploration in model-based reinforcement learning by empirically estimating learning process, Lopes et al., NeurIPS 2012
Curiosity-driven exploration by self-supervised prediction, Pathak et al., ICML 2017
A large-scale study of curiosity-driven learning, ICLR 2019

Exploration taxonomy - reconstruction

Reconstruction

- Actively gather information to reconstruct the entire environment

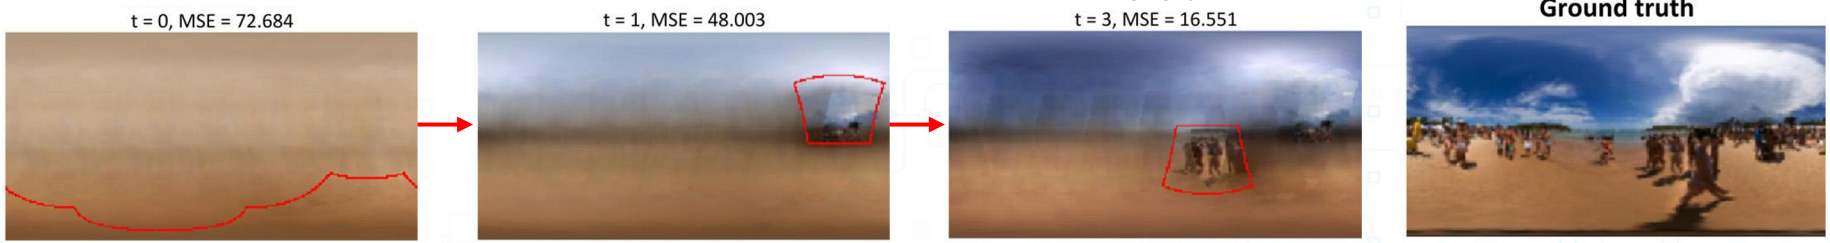

$$r(s) \propto -\Delta \text{MSE}$$

Learning to Look Around: Intelligently Exploring Unseen Environments for Unknown Tasks, Jayaraman & Grauman, CVPR 2018
Sidekick Policy Learning for Active Visual Exploration, Ramakrishnan & Grauman, ECCV 2018
Emergence of Exploratory Behaviors through Active Observation Completion, Ramakrishnan et al., Science Robotics 2019

Our contributions

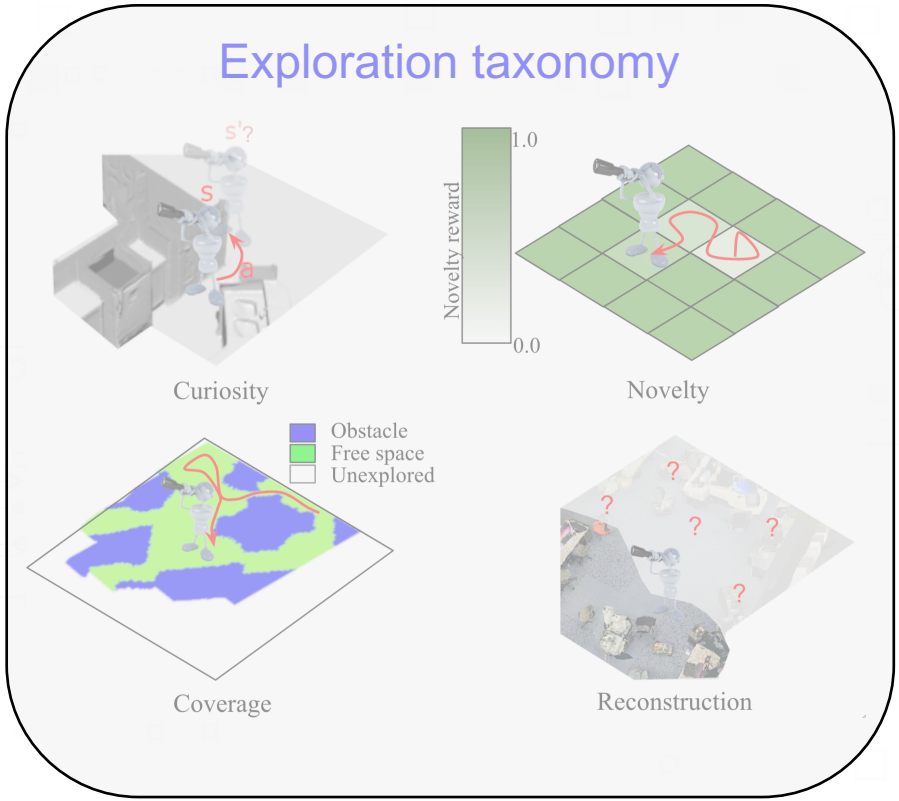

Exploration study framework

- High-quality 3D environments
- State-of-the-art policy architecture
- Diverse evaluation metrics

Exploration study framework

Photorealistic 3D environments

Active Vision Dataset [1]

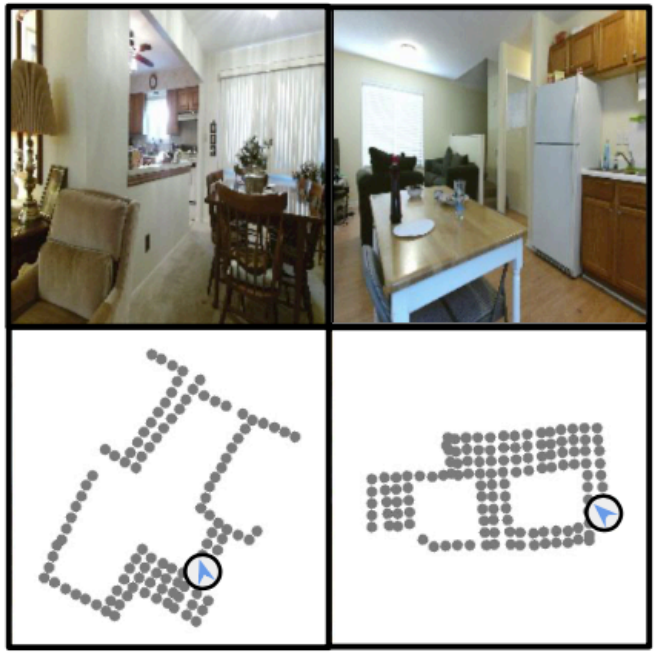

[1] A dataset for developing and benchmarking active vision, Ammirato et al., ICRA 2016
[2] Matterport3D: Learning from RGB-D data in indoor environments, Chang et al., 3DV 2017

Exploration study framework

Exploration policy architecture

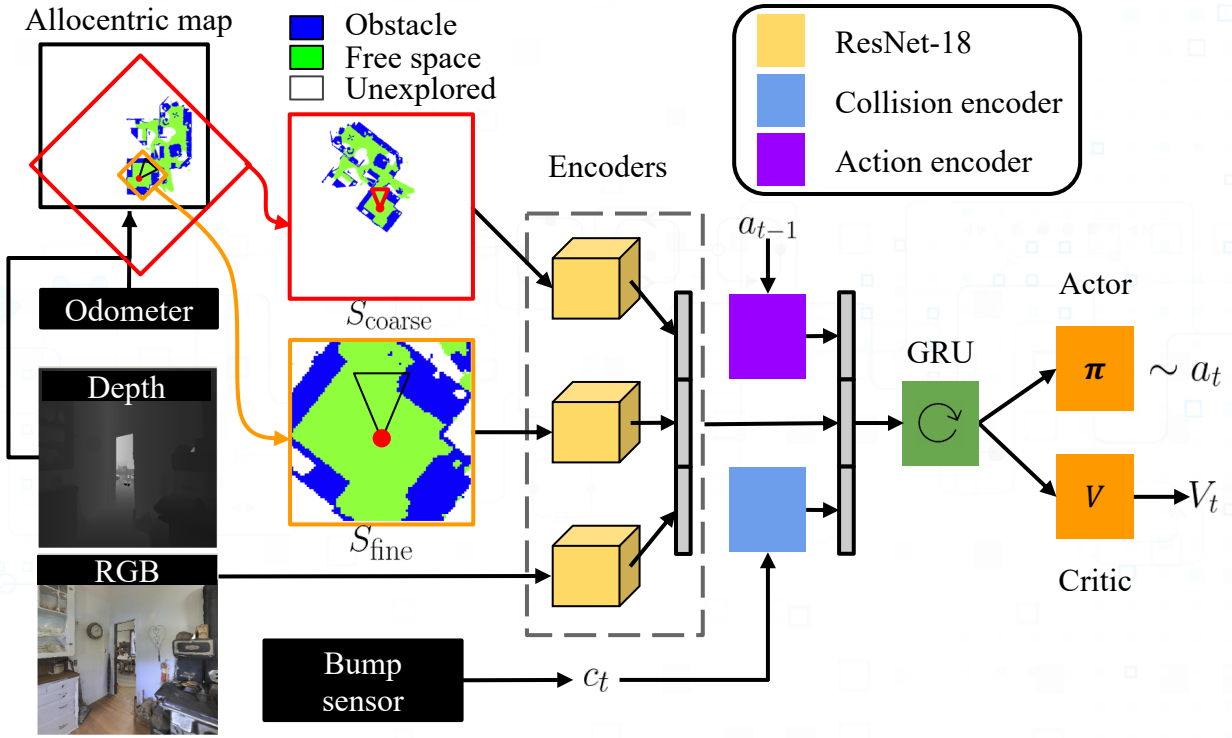

Exploration study framework

Evaluation metrics suite

Visitation metrics

- Area covered
- Objects covered

Downstream task transfer

- Pose estimation
- Concept reconstruction
- Navigation

Experiments

Perfect odometry assumption

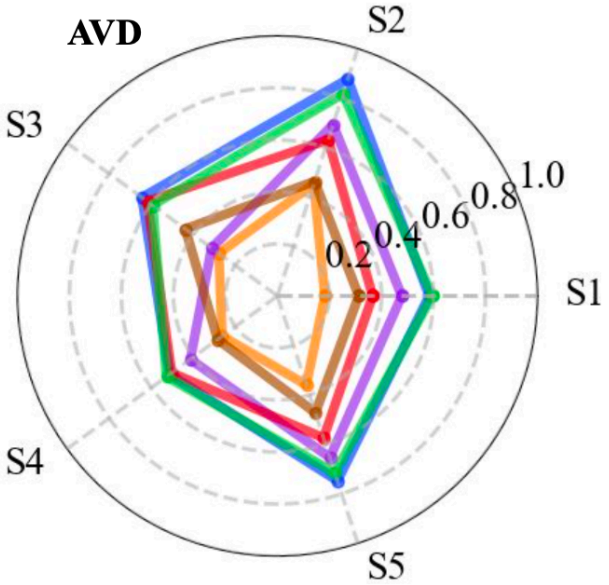

S1 - Navigation S2 - Mapping S3 - Object discovery
S4 - Localization S5 - Reconstruction

Experiments

Impact of noisy odometry on exploration

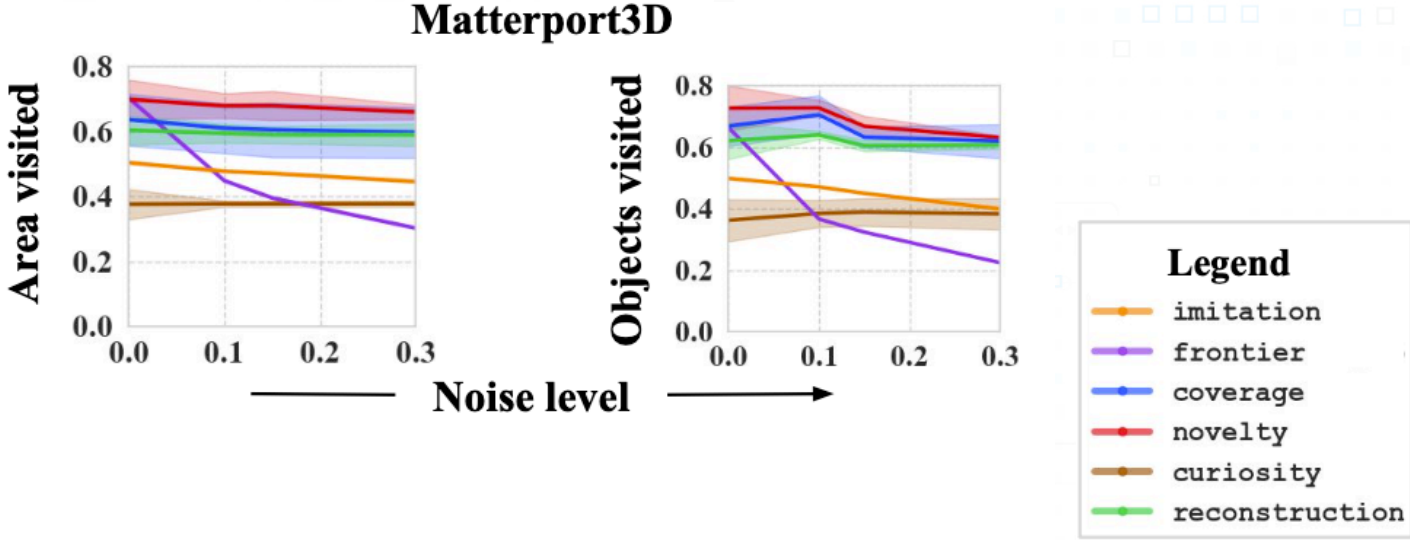

Conclusions

- Taxonomy of exploration algorithms
- Exploration study benchmark
  - Photorealistic 3D environments
  - High-performance exploration architecture
  - Diverse evaluation metrics
- Study reveals strengths and weaknesses of existing approaches.

An Exploration of Embodied Visual Exploration

Project page: http://vision.cs.utexas.edu/projects/exploring-exploration
Code: https://github.com/facebookresearch/exploring_exploration

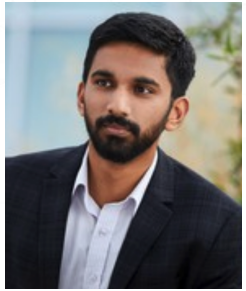

Santhosh Kumar Ramakrishnan[1,3]

[1] The University of Texas at Austin

Dinesh Jayaraman[2,3]

[2] University of Pennsylvania

Kristen Grauman[1,3]

[3] Facebook AI Research