[Habitat Challenge]

Mid-Level Visual Representations (RGB Track Winner)

Alexander Sax* Jeffrey O. Zhang* Bradley Emi Amir R. Zamir Silvio Savarese Leonidas Guibas Jitendra Malik

Berkeley UNIVERSITY OF CALIFORNIA · Stanford University

CVPR LONG BEACH CALIFORNIA June 16-20, 2019

https://perceptual.actor

Overview

- **What should lie between pixels and control**? An agent-understandable abstraction of the world. We use mid-level visual representations to learn **better policies with less data.**

Learning *Tabula Rasa* | Learning with Visual Priors

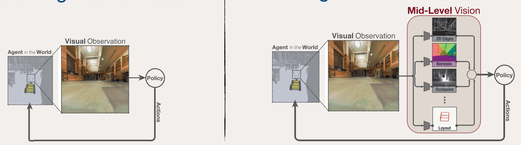

- **Problem Statement:** *For a given task (POMDP: (S, A, r, T, Ω, O)) find a policy π: $\{o_i\}_1^t \mapsto a_{t+1}$ that maximizes the expected test-set reward.*
- When are different representations useful? We examine
  I. sample efficiency
  II. generalization
  III. universality of representations

Max-Coverage Min-Distance Feature Set

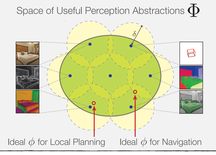

Solve BIP sequence:

$$\text{minimize: } \mathbb{1}^T x,$$
$$\text{subject to: } Ax \succeq \delta \text{ and } x \in \{0,1\}^m$$

Results: Max-Coverage Min-Distance Feature Set

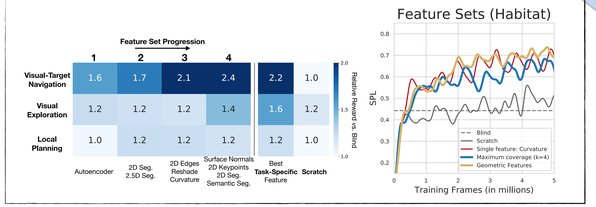

Methodology

Perceptual Features

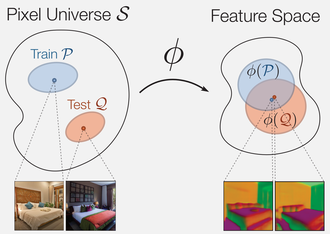

Feature Dictionary

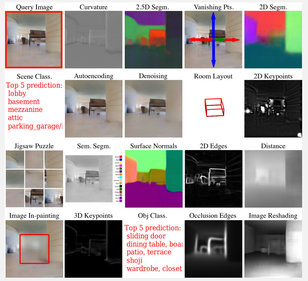

Reinforcement Learning Setup

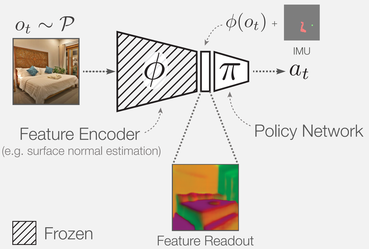

- NO recurrent memory
- NO HP tuning
- NO architecture search
- NO random restarts
- NO ensembling
- NO initialization from behavior cloning
- NO reward shaping
- NO domain adaptation

Hypothesis III: No Universal Feature

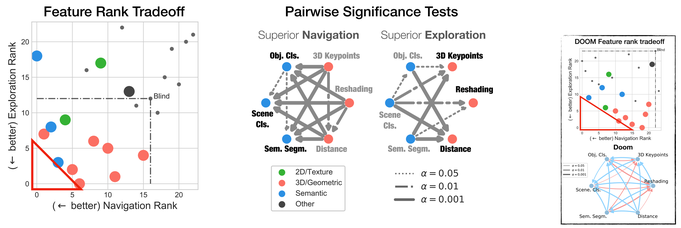

Results

Hypothesis I: Better Sample Efficiency

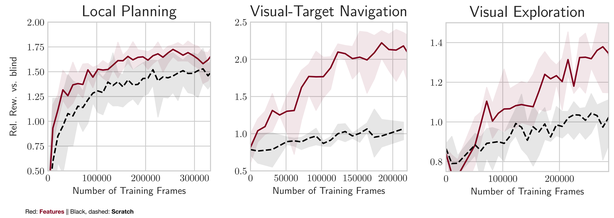

Hypothesis II: Stronger Generalization

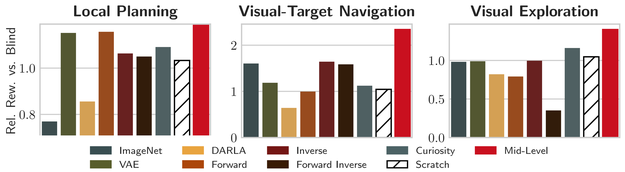

Baselines

- Random
- Blind
- Scratch
- Pixels-as-state

Not All Features are Equal

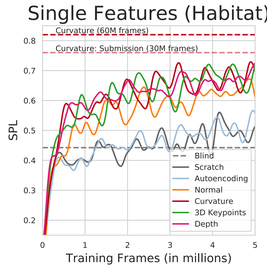

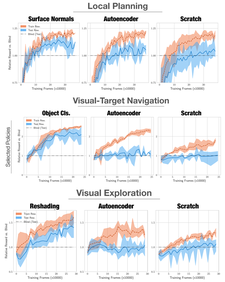

Environments

Gibson Simulator

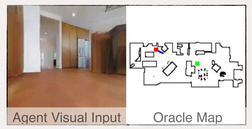

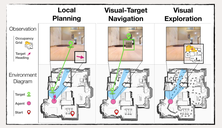

Habitat

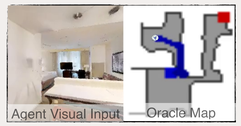

Doom

Conclusions and Limitations

- Introducing the appropriate visual bias improves generalization and reduces sample complexity
- Results reproduce across multiple environments and tasks
- Choice of vision feature is essential to performance, and the correct choice depends on downstream task
- Our study focuses on locomotive tasks, but other tasks require manipulation and have dynamics environments

Website, demos, results:
http://perceptual.actor
I: https://perceptual.actor/policy_explorer/
II: https://perceptual.actor/generalization_curves/
https://arxiv.org/pdf/1812.11971

Code/Toolkit/Models:
https://github.com/alexsax/midlevel-reps

* Authors contributed equally