Communicating temporal tasks to AI agents is hard!

❄ pre-trained ❄ RAD Embeddings make it easier.

Motivation

1. **Multi-purpose** robots are more capable than ever.
2. Hard to **unambigiously** specify **temporal** tasks for agents.

Proposal

1. Specify tasks as **Deterministic Finite Automata** (DFA) - unambiguous and **as easy to read as a flow chart**.
2. Encode DFA in latent space to condition RL policies on.

Three Problems

1. Encoding large DFAs is computationally expensive.
2. Hard to train encoder and policy at the same time.
3. All training distributions over DFAs are biased.

Solved by encoding Boolean combinations of DFAs

Solved by ...

Pretraining on DFAs that encorage learning to plan paths through DFA

Reach green while avoiding blue and red

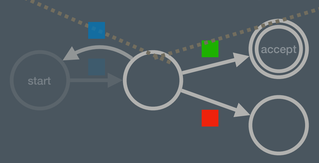

Traversing any path a DFA is a series of reach avoid problems - which itself is representable as DFAs.

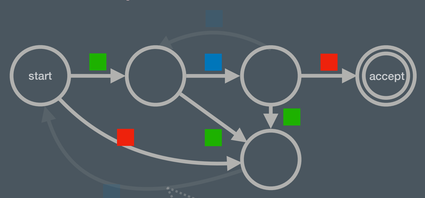

We pretrained a Graph Neural Network on these mutations of these DFA (called **RAD**).

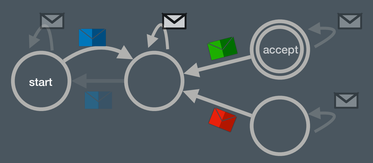

Messages passed in reverse direction and accumulated via attention mechanism (GATv2)

Compositional Automata Embeddings for Goal-Conditioned Reinforcement Learning

Beyazit Yalcinkaya*[1], Niklas Lauffer*[1], Marcell Vazquez-Chanlatte*[2], Sanjit A. Seshia[1]

*Equal contribution
[1]University of California, Berkeley
[2]Nissan Advanced Technology Center - Siilicon Valley

Beats hierarchical planning.

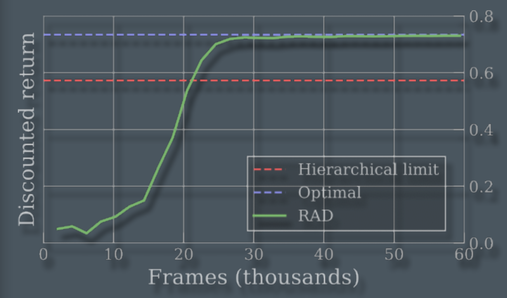

Freezing the embeddings gives the best performance.

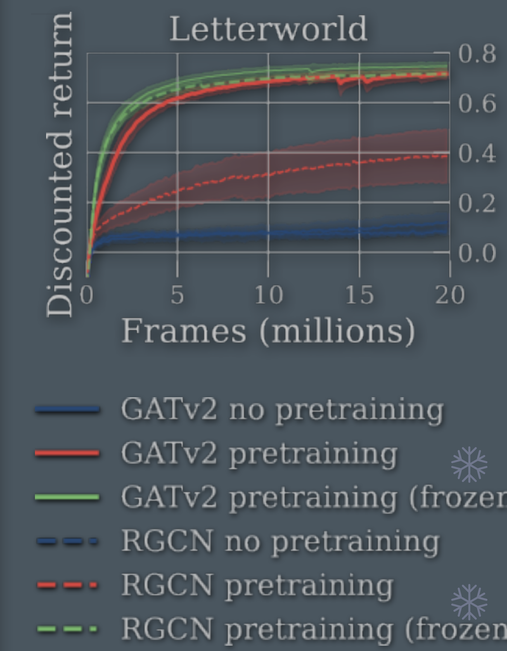

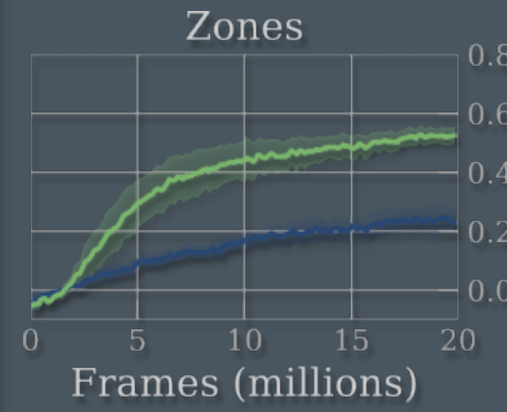

Generalizes to other tasks.

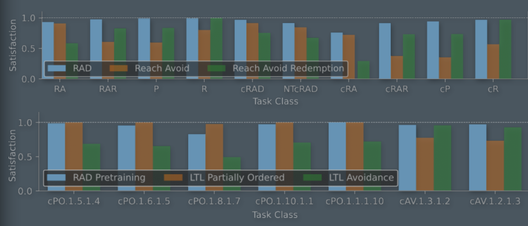

[1] Vaezipoor, Pashootan, et al. "Ltl2action: Generalizing ltl instructions for multi-task rl." International Conference on Machine Learning. PMLR, 2021.