Generalized Planning with Deep Reinforcement Learning

TECHNION
Israel Institute of Technology

Or Rivlin, Erez Karpas and Tamir Hazan
[srivlin@campus., karpase@, tamir.hazan@]technion.ac.il

Generalized Planning

- Generalized planning aims at finding plans that work for many problems
- We researched learning policies that generalize very well, acting as approximate generalized plans

State Representation

- We represent states and goals as a graph, with nodes, edges and global features
- The graph is complete, and structure is encoded through the features
- State features are concatenated to goal features

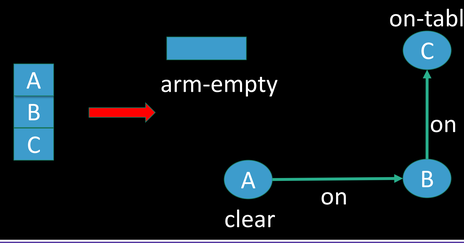

State:
arm-empty
clear b1
on-table b3
on b1 b2
on b2 b3

Goals:
on b1 b3

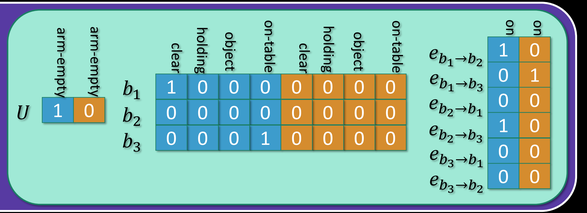

Policy Representation

- Our policy is composed of two modules; a graph embedding module and an action module
- The graph embedding module is a graph neural network (GNN) which performs message passing between the graph nodes

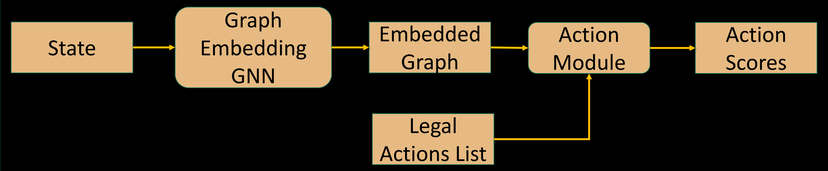

- Embedded elements are clustered by their type and concatenated with an indicator vector
- Finally the effect vectors are scattered back to their actions and aggregated to from action embeddings

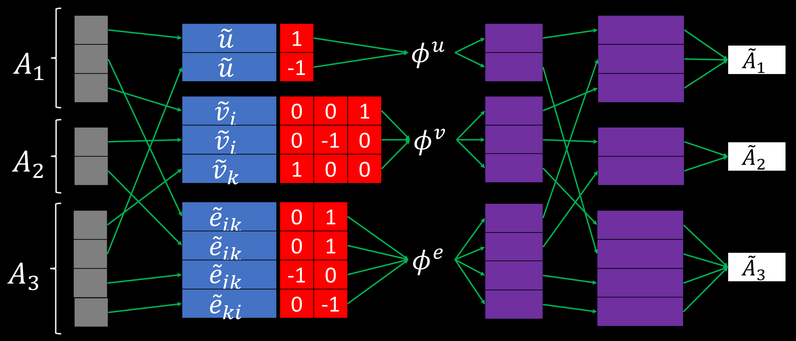

Tree Search Algorithm

- We use GBFS with a full policy rollout at each expanded leaf
- Only the first state-action of the rollout is added to the open list

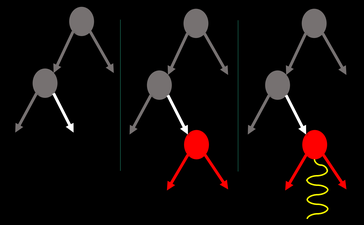

Experimental Results

- Test success rate against number of expanded states: