

Towards Sample Efficient Agents through Algorithmic Alignment (Student Abstract)

Mingxuan Li*, Michael L. Littman,

Department of Computer Science, Brown University, Providence, RI 02912-1910
{mingxuan.li, michael.littman}@brown.edu

Abstract

In this work, we propose and explore Deep Graph Value Network (DeepGV) as a promising method to work around sample complexity in deep reinforcement-learning agents using a message-passing mechanism. The main idea is that the agent should be guided by structured non-neural-network algorithms like dynamic programming. According to recent advances in algorithmic alignment, neural networks with structured computation procedures can be trained efficiently. We demonstrate the potential of graph neural network in supporting sample efficient learning by showing that Deep Graph Value Network can outperform unstructured baselines by a large margin in solving Markov Decision Process (MDP). We believe this would open up a new avenue for structured agents design. See <https://github.com/drmeerkat/Deep-Graph-Value-Network> for the code.

Introduction

Deep reinforcement-learning algorithms have produced breakthroughs in recent years. However, agents with powerful non-linear function approximators also require large amounts of experience to learn. In this work, the question we focus on is how to get the best of both neural networks and non-neural network algorithms (e.g. Dynamic Programming) to learn effectively.

One intuitive solution is to execute those algorithms with neural networks. Different datasets favor different network structures and computational process. Recent work has shown that better algorithmic alignment improves sample complexity and generalization (Xu, Li et al. 2020). Given the fact that GNNs align well with dynamic programming (DP), we presume value iteration, as a probabilistic version of Bellman-ford algorithm, should also be solvable with GNNs. Thus, we propose Deep Graph Value Network (DeepGV), a message-passing framework, to robustly solve the given MDP. We empirically verify its effectiveness as a general MDP solver and its potential towards building general structured agents.

*Contact: +1 4012256526, Blog: <https://drmeerkat.github.io>
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Algorithmic Alignment

We adopt the theoretical framework of (Xu, Li et al. 2020). In this section, we will briefly review the basic definitions. For more details, please refer to Xu, Li et al. (2020).

Definition 2.1. (PAC learning and sample complexity). Fix an error parameter $\epsilon > 0$ and failure probability $\delta \in (0, 1)$. Suppose $\{x_i, y_i\}_{i=1}^M$ are i.i.d. samples from distribution \mathcal{D} , and the data satisfies $y_i = g(x_i)$ for some underlying function g . Let $f = \mathcal{A}(\{x_i, y_i\}_{i=1}^M)$ be the function generated by a learning algorithm \mathcal{A} . Then, g is (M, ϵ, δ) -learnable with \mathcal{A} if

$$\mathbb{P}_{x \in \mathcal{D}}[|f(x) - g(x)| \leq \epsilon] \geq 1 - \delta. \quad (1)$$

The *sample complexity* $\mathcal{C}_{\mathcal{A}}(g, \epsilon, \delta)$ is

$$\min\{M | \mathbb{P}_{x \in \mathcal{D}}[|f(x) - g(x)| \leq \epsilon] \geq 1 - \delta\}. \quad (2)$$

Definition 2.2. (Algorithmic alignment). Let g be a reasoning function and \mathcal{F} a neural network with n modules f_i . The module functions g_1, \dots, g_n generate g for \mathcal{F} if, by replacing f_i with g_i , the network simulates g . Then, $\mathcal{F}(M, \epsilon, \delta)$ -algorithmically aligns with g if (1) g_1, \dots, g_n generate g and (2) there are learning algorithms \mathcal{A}_i for the f_i such that $n \cdot \max_i \mathcal{C}_{\mathcal{A}_i}(g_i, \epsilon, \delta) \leq M$.

From the definition, it is clear that only if each module of the neural network aligns well with the underlying function modules can we reduce M to its minimum.

Deep Graph Value Network

Due to the space limitation, for general context of MDP and graph neural network, please refer to Sutton and Barto (2018); Battaglia, Hamrick et al. (2018). In each iteration of message distribution and aggregation, every node first updates its embeddings according to information from its neighbours and edges between them. After some iteration of execution, we aggregate all the information and extract answers from it with another network. As shown in Fig. 1, comparing general GNN updates with value iteration, they share quite similar underlying computational steps, which can be simulated by neural networks easily.

The framework of Deep Graph Value Network is as follows. The network takes a complete graph $G(V, E)$ as input with different attributes assigned to nodes and edges. Each node represents a state in the given MDP $\langle S, A, R, \gamma, T \rangle$.




Figure 1: An example of how generally GNNs align with value iteration. Both GNNs and VI share the same loop structure, which does not need to be learned. Thus, it would be relatively simpler for GNNs than generic Multi-layer Perceptrons (MLPs) to learn to execute value iteration.

Edges between nodes represent actions with rewards binding to them.

In iteration $k+1$, the node embedding $h_i^{(k+1)}$ is calculated as follows,

$$f_\psi \left(\text{CAT}_{a \in \mathcal{A}} \sum_{j \in \mathcal{N}(i)} A_{i,j}^{(k+1)} \left(r(i, a, j) + \gamma h_j^{(k)} \right) \right), \quad (3)$$

where $\text{CAT}(\cdot)$ is a concatenate operator, r is the embedded reward attributes attached to edges, γ is the discount factor and $A_{i,j}$ is the transition probability learned by a self-attention mechanism. The part inside the square brackets calculates the Q-value for state i and the local aggregate function f_γ learns to execute max. The attention mechanism is learned as follows,

$$A_{i,j}^{(k+1)} = \frac{\exp(e_{i,j}^{(k+1)})}{\sum_{l \in \mathcal{N}(i)} \exp(e_{i,l}^{(k+1)})}, \quad (4)$$

$$e_{i,j}^{(k+1)} = f_\theta(h_i^{(k)}, h_j^{(k)}). \quad (5)$$


After K iterations, we aggregate information from all the nodes and pass them through an MLP to get value-ranking predictions:

$$\hat{y} = f_w \left(\text{CAT}_{i \in \mathcal{V}} h_i^{(K)} \right). \quad (6)$$


Experiments and Discussion

In this section, we apply our framework to solve MDPs and compare its performance with that of vanilla MLPs. Intuitively, unlike DeepGV, MLP doesn't align well with dynamic programming. Thus, we expect to see that DeepGV can easily be trained and outperforms the MLP by a large margin.

We evaluate the networks on two different types of MDPs. The first one is a slippery $N \times N$ grid world without obstacles and fire states. The second type has obstacles and fire states. We generate different instances of these two kinds of MDPs and evaluate two algorithms on each. To solve a given MDP, the input is randomly initialized state values, state coordinates and the queried state coordinates. The expected output is the value rank of the queried state.



(a) Value ranking accuracy varying grid world size.



(b) Training set size vs. test accuracy.

Figure 2: Test accuracy and sample complexity on two types of grid world. In each sub-figure, the LHS result comes from plain grid world while the RHS result comes from grid world with traps. (a) MLP fails on 8×8 world while DeepGV performs well. (b) As the training set size increases, DeepGV achieves good performance with fewer training examples.

As shown in fig. 2, when we fix the training set size and increase the task difficulty (the size of grid world), DeepGV can still be trained to a satisfying level while MLP performance drops severely in the 8×8 grid world. When we varied the training set size and fixed the size of the grid world at 8×8 , there is a clear gap between MLP and DeepGV performance. These results empirically support Theorem 3.6 from (Xu, Li et al. 2020) that better algorithmic alignment induces better sample complexity. Overall, our experiment verifies the intuition that DeepGV can be trained easily and outperforms the MLP by a large margin.

In conclusion, we presented Deep Graph Value Network (DeepGV), a graph network structure that can efficiently learn to solve MDPs by executing a similar computational procedure to that of value iteration. Our preliminary experimental results show the performance gap between structured networks and unstructured networks in solving relational tasks suggesting the promising potential of message-passing mechanisms in building structured RL agents. We believe this opens an exciting new avenue for future work on designing high-efficiency agents incorporating algorithmic prior.

References

- Battaglia, P. W.; Hamrick, J. B.; et al. 2018. Relational inductive biases, deep learning, and graph networks. *CoRR* abs/1806.01261.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Xu, K.; Li, J.; et al. 2020. What Can Neural Networks Reason About? In *ICLR*.