

# A State Aggregation Approach for Solving Knapsack Problem with Deep Reinforcement Learning

**Citation for published version (APA):**

Refaei Afshar, R., Zhang, Y., Firat, M., & Kaymak, U. (2020). A State Aggregation Approach for Solving Knapsack Problem with Deep Reinforcement Learning. In *Proceedings of The 12th Asian Conference on Machine Learning (ACML2020)* (pp. 81-96). (Proceedings of Machine Learning Research; Vol. 129). <https://arxiv.org/abs/2004.12117>

**Document status and date:**

Published: 01/01/2020

**Document Version:**

Accepted manuscript including changes made at the peer-review stage

**Please check the document version of this publication:**

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

**General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

[www.tue.nl/taverne](http://www.tue.nl/taverne)

**Take down policy**

If you believe that this document breaches copyright please contact us at:

[openaccess@tue.nl](mailto:openaccess@tue.nl)

providing details and we will investigate your claim.

# A State Aggregation Approach for Solving Knapsack Problem with Deep Reinforcement Learning

Reza Refaei Afshar<sup>1</sup>[0000-0003-1558-8380], Yingqian Zhang<sup>1</sup>[0000-0002-5073-0787],  
Murat Firat<sup>1</sup>[0000-0001-8483-3533], and Uzay Kaymak<sup>1</sup>[0000-0002-4500-9098]

Eindhoven University of Technology, Eindhoven, Netherlands

**Abstract.** This paper proposes a Deep Reinforcement Learning (DRL) approach for solving knapsack problem. The proposed method consists of a state aggregation step based on tabular reinforcement learning to extract features and construct states. The state aggregation policy is applied to each problem instance of the knapsack problem, which is used with Advantage Actor Critic (A2C) algorithm to train a policy through which the items are sequentially selected at each time step. The method is a constructive solution approach and the process of selecting items is repeated until the final solution is obtained. The experiments show that our approach provides close to optimal solutions for all tested instances, outperforms the greedy algorithm, and is able to handle larger instances and more flexible than an existing DRL approach. In addition, the results demonstrate that the proposed model with the state aggregation strategy not only gives better solutions but also learns in less timesteps, than the one without state aggregation.

**Keywords:** Knapsack Problem · Deep Reinforcement Learning · State Aggregation.

## 1 Introduction

Heuristic algorithms for solving Combinatorial Optimization Problems (COPs) achieve acceptable solutions in a polynomial time. These algorithms relies on handcrafted heuristics that conduct the process of finding the solutions. Although these heuristics work well in many COPs, they mostly rely on the nature of problems and they need to be modified for different class of problems. In this paper, we aim to learn and improve the handcrafted heuristics to improve the quality of the solutions. We study *knapsack problem (KP)*, which is one of the well-known benchmark problems in COPs. KP is defined as a set of items, each with a value and a weight. The objective is to select a subset of items with maximum total value to fill a knapsack such that the cumulative items weight does not exceed its capacity. This problem has many applications such as cargo loading, cutting stock and capital budgeting [28].

Recently, there is a great progress in the Artificial Intelligence (AI) community in developing machine learning (ML) methods to solve COPs [3], where a

popular ML based method is *Deep Reinforcement Learning (DRL)*. DRL is the integration of *Reinforcement Learning (RL)* and *Deep Neural Networks (DNN)* [1,18,13]. Several DRL based approaches have been proposed to solve the Traveling Salesman Problem (TSP), e.g. [2,14,17], where a discrete representation of TSP is used as states and the solution is a sequence of the inputs. These approaches work well for the TSP problem, however, in the Knapsack problem, the values and weights of items are continuous, which entails an extremely large state space when the number of items increases. Hence, the existing DRL based approaches for solving the problems with discrete nature such as TSP might not work well for KP, as shown in [2]. The authors of [2] solve a Knapsack problem using the policy gradient algorithm with pointer networks. Although they show optimal solutions can be obtained for instances up to 200 items, the following limitations are identified: (1) intractability to large instances: the state space grows rapidly with increasing number of items, and (2) generality to other sizes of instances: the trained model is applicable for solving the problems that have the exactly same knapsack capacity and same number of items. In this paper, we introduce a DRL approach with state aggregation that boosts the capability of the typical greedy algorithm and improves the heuristic to overcome these two limitations.

In our approach, propose a state aggregation method to discretize the feature values of items. We construct a feature table by assigning a row for each problem instance and a column for each item’s information. A tabular reinforcement learning is used to learn the best operation strategy for each item. This resulting discretization of features not only provides a discrete representation of the problem instances, but also reduces the state space by reducing the number of unique values. However, the state space is still large despite state aggregation. Therefore we exploit DRL as a powerful function approximation approach. We use *Advantage Actor Critic (A2C)* algorithm to learn the policy of selecting items. A2C makes use of two DNNs for learning policy and value functions [19]. The policy DNN has  $N$  outputs which is equal to the number of items. By following a greedy or softmax algorithm on the output of the policy DNN, a sequence of items are selected until the knapsack is full.

The experimental results show that the proposed approach finds optimal solutions with two decimal places for the problem instances of same size used in [2]. Moreover, we show the method obtains close to optimal solutions for three different types of instances with at most 50, 300 and 500 items. We also demonstrate that the proposed DRL with state aggregation performs better than the DRL without aggregation in terms of both learning rate and the solution quality. We summarize our contributions as follows.

- We develop a state aggregation strategy to derive state embedding that reduces the state space size. We show this general strategy effectively speed up learning on solving KP.
- Our DRL-based approach to solve KP improves the heuristic greedy algorithm for 0-1 KP and shows better performance than the existing DRL

approaches. The developed method can be trained once for  $N$  items and it can be used for any KP instances with size up to  $N$ .

## 2 Related Work

It has been proven by reduction that most of COPs are NP-Hard problems. Their optimal solutions can not be found in polynomial time and exact algorithms take exponential time to find optimal solutions [15,4]. Knapsack Problem (KP) has gained a remarkable attention in the literature. Despite the fact that the fractional KP is optimally solvable by the heuristic greedy algorithm, the 0-1 knapsack problem is NP-Hard [5], and a large variety of KPs remain hard to solve [21]. Moreover, it has been shown by empirical evidence that solving instances near the phase transition are challenging for humans [30]. The phase transition emerges around critical values of items and capacity so that the probability of having a solution for an instance change from zero to one. Many algorithms, ranging from dynamic programming algorithms (e.g. [7]) to meta-heuristics (e.g. [9]) have been proposed to solve KP.

Cleverly searching and branch and bound methods can prune the search tree and reduce the time complexity of COPs [29]. However, these methods are still prohibitive for large instances. Polynomial time approximation schemes and integer linear programming (ILP) based approaches are the other helpful methods [25,8]. Although the approximation algorithms might be performed in reasonable time, they rely on handcraft heuristics and the methods need to be revised when the problem settings change. Furthermore, they suffer from weak optimality for some problems. In order to cope with this limitations, Machine Learning (ML) based and data driven methods are developed.

In recent years, it has been shown that DRL can be used for learning good heuristics for solving COPs. In [26], the Pointer Network architecture is introduced where the output layer of the deep neural network is a function of the input. In [2], the pointer network is used with RL to solve the Traveling Salesman Problem (TSP). They use policy gradient and a variant of Asynchronous Advantage Actor-Critic (A3C) algorithm of [19] to train a DNN, and show close to optimal solutions are found for up to 100 cities. In [16] another neural network framework is introduced for graph-based COPs, where *structure2vec* [6] is used to derive an embedding for the vertices of the graph. The *structure2vec* computes a  $p$ -dimensional feature embedding for each node and a parametric  $Q$  function is trained using Q learning algorithm. In [17], the pointer network is incorporated with attention layers. With the REINFORCE algorithm, they obtained close to optimal solutions for the TSP instances of up to 100 nodes.

Most of ML-based research on solving COPs focuses on TSP. COPs like TSP and Vehicle Routing Problem that have gained high attentions in past few years, require a sequence of the input as the solution and sequence-to-sequence neural architectures might be proper approaches for solving them [23]. However, the solutions of COPs like KP and Weighted Vertex Cover are a subset of the input. This issue makes the original sequence-to-sequence approaches inapplicable for

solving KP. Recently, a pointer network deep learning approach is presented for solving 0-1 KP [10]. This method is based on supervise learning and optimal solutions which is not available in most of the cases. In this paper we propose a DRL framework for subset selection problems.

### 3 Problem Definition and Modeling

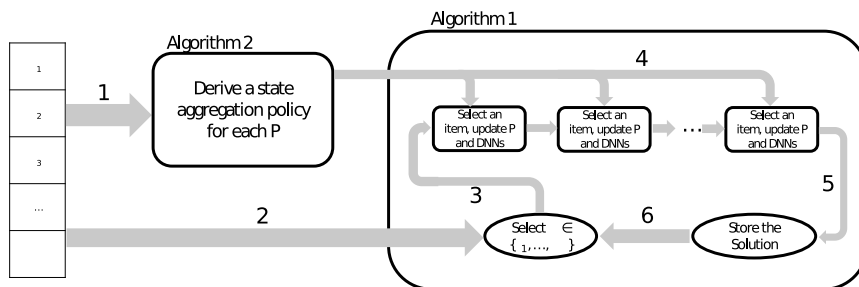
We consider the following *0-1 Knapsack Problem* instance  $P$ : We are given a set  $\mathcal{I}_P$  containing  $n_P$  items and a knapsack of capacity  $W_P$ . Each item  $i$  has value  $v_i$  and weight  $w_i$ . The goal is to fill the knapsack with a selected subset of items such that the total weight of the selected items does not exceed  $W_P$  and the total value is maximized. Since  $P$  is a 0-1 KP, selecting a fraction of an item is not possible.

Our method for solving this variant of KP is based on deep reinforcement learning. We assume that the number of items is variable and a constructive solution can solve the problem only by considering the capacity constraint. Therefore, the process of selecting a subset of items  $\mathcal{I}'_P \subseteq \mathcal{I}_P$  is modeled as a sequential decision process. The policy DNN is trained with A2C introduced in [19] on a set of problem instances with at most  $N$  items. The information of each problem instance consists of  $|\mathcal{I}_P| = n_P \leq N$  items with value  $v_i$  and weight  $w_i$  for each  $i \in \mathcal{I}_P$  and together with  $W_P$ , they are the inputs of DNN. The DNN has  $N$  outputs that each being associated with a value of selecting a specific item  $i \in \mathcal{I}_P$ . The policy is to select an item in each step. After selecting item  $i$ , it is removed from the original problem instance  $P$  and a new problem instance  $P'$  with a reduced item set  $\mathcal{I}_{P'} = \mathcal{I}_P \setminus \{i\}$  and capacity  $W_{P'} = W_P - w_i$  is generated. For the cases where  $i$  cannot be added to the knapsack because of the capacity constraint, the new instance  $P'$  is generated by simply removing  $i$  from the item set, without altering  $W_P$ . In this way, when the policies are trained with problem instances of at most  $N$  items, the policies can be used to find solutions for new instances as long as their item sizes are no greater than  $N$ .

Such KP problems can be found in different applications. For example, an online ad publisher faces with a set of advertisements. Assuming a fixed upper bound for the number of ads, the problem is to select a subset of them to show to the users. In this example, the values are relevance scores and the weights are the size of ads banners. The goal is to fill a slot of a certain size with the ads.

### 4 DRL-based KP Solver

Figure 1 shows the overview of our proposed method. It consists of two components. The first component includes a formulation of KP to MDP, which is solved using a DRL approach (Algorithm 1). The second component is a state aggregation method (Algorithm 2), which learns a aggregation policy to discretize states that are serves as inputs to DRL. We first discuss how to formulate the KP problem as MDP.



**Fig. 1.** The overview of the KP solver method. 1) A set of problem instances are used for deriving an aggregation policy for item information. 2) The same set of problems are used in the second step which is DRL. 3) A problem instance is selected for training. 4) Items are selected sequentially until finding a solution. At each step the updated  $P$  is aggregated to find the state. The parameters of value and policy DNNs are updated using A2C. 5) The best solution is stored. 6) Another problem instance is selected for training. The process continues for a certain number of timesteps.

#### 4.1 Deep Reinforcement Learning method

In order to solve the 0-1 KP, DRL is used to derive a policy through that the items are sequentially added to the solution. We define the states, actions and rewards of DRL modeling of KPs for an instance  $P'$  which is a representation of  $P$  after selecting some items, as follows.

*States  $s(P)$ :* A complete set of information of instance  $P'$  containing  $n_{P'}$ ,  $v_i$  and  $w_i$  for  $n_{P'}$  items, capacity  $W_{P'}$ , the total value of the items ( $Sv = \sum_{i \in \mathcal{I}_{P'}} v_i$ ), and the total weight of the items ( $Sw = \sum_{i \in \mathcal{I}_{P'}} w_i$ ) makes a feature vector of  $2n_{P'} + 4$  features. Since  $n_{P'} \leq N$  for all  $P'$ , the feature vector of the instances that have  $n_{P'} < N$  items consists of  $2N + 4$  features such that the first  $2n_{P'} + 4$  ones carry the information of the problem instance and the remaining ones are zero. Section 4.2 will reduce this feature vector by a state aggregation strategy.

*Actions:* There are  $N$  actions  $A_1, A_2, \dots, A_N$ , each corresponding to select one item. At each decision moment, a state is fed to the policy DNN and an action is selected according to the output of the DNN.

*Reward Function:* The reward function is defined based on three criteria. First, if item  $i$  can be added to the knapsack without exceeding the capacity limit,  $v_i$  is used as a positive reward. Second, if  $w_i$  is greater than  $W_{P'}$ , i.e.  $i$  cannot be added to the knapsack,  $-w_i$  is set as a negative reward. Third, for each instance  $P'$  where  $n_{P'} < N$ , the first  $n_{P'}$  outputs of DNN correspond to the items of  $P'$  and the next  $N - n_{P'}$  outputs are undefined actions because the corresponding items do not exist. Therefore, a large penalty i.e.  $-W_{P'}$  is used for the reward of choosing undefined actions. We separate the reward of undefined action and

heavy items because an action with  $i > n$  is always undefined, however items with  $w_i > W_P$  could be added to the knapsack if they were selected in earlier steps. Therefore, their penalty is lower. Equation (1) shows the reward of state  $s(P')$  and action  $A_i$ .

$$r(s(P'), A_i) = \begin{cases} -W_{P'} & \text{if } i > n_{P'} \\ v_i & \text{if } w_i \leq W_{P'} \\ -w_i & \text{if } w_i > W_{P'} \end{cases} \quad (1)$$

Employing these definitions of states, actions and rewards, the A2C algorithm is used for training policy and value DNNs [19], where two DNNs are used for policy ( $\pi$ ) and value ( $V$ ) functions. The advantage value is obtained by subtracting state values ( $V$ ) from state action values ( $Q$ ) which is defined by  $r + \gamma V(s_{t+1})$ . This value is used in gradient function to update the parameters of the DNNs using Equations (2) and (3) [19,12].

$$\theta^{t+1} \leftarrow \theta^t + \nabla_{\theta^t} \log \pi(A_i | s(P), \theta^t) [r_t + \gamma V(s(P'), \theta_v^t) - V(s(P), \theta_v^t)] \quad (2)$$

$$\theta_v^{t+1} \leftarrow \theta_v^t + \frac{\partial (r_t + \gamma V(s(P'), \theta_v^t) - V(s(P), \theta_v^t))^2}{\partial \theta_v^t} \quad (3)$$

where,  $\theta^t$  and  $\theta_v^t$  are the parameters of policy and value DNNs in decision moment  $t$  respectively. The corresponding state of a problem instance  $P$  is fed to the policy DNN and the items can be selected by following a policy according to the output of the policy DNN. Upon selecting an item,  $P'$  is obtained from  $P$  and it is again fed to the policy DNN to select the next item. This process is continued until filling the knapsack or exceeding the weight constraint. Algorithm 1 shows the DRL-based knapsack solver method.

## 4.2 State Aggregation

As the number of items increases, the state space grows up exponentially and this affects the performance of function approximation with DNN. In order to shrink the state space and boost the method to have the capability of solving large problem instances, a new state embedding is derived by state aggregation. The feature values of states are divided into subsets and the values of each subset are converted to a certain value. Instead of manually testing different number of subsets to find the one that has the best performance, the process of finding appropriate number of subsets is considered as a sequential decision making problem and reinforcement learning is used for solving the problem. Before developing the RL framework, we first pre-process the problem instances.

*Preparing data.* A set of problem instances are used for deriving the state embedding. Each problem instance is identified by a set of feature values which are the items information and capacity. The first step in aggregating the states is to

---

**Algorithm 1** DRL-based Knapsack Solver

---

**Input:**  $M$  Problem Instances each having at most  $N$  items**Output:** Values of solutions of the  $M$  instances

```

1: Initialize a policy DNN with  $2N + 4$  inputs,  $N$  outputs and parameters  $\theta$  as policy
    $\pi(A_i|s, \theta)$ 
2: Initialize a value DNN with parameters  $\theta_v$  as  $V(s, \theta_v)$ 
3:  $t_{max} = 3N \times 10^4$ ,  $t = 0$ 
4: Initialize  $Val$ : a list of length  $M$ , all 0
5: while  $t < t_{max}$  do
6:   Select a problem instance  $P$  with capacity  $W_P$ .
7:    $ow = 0$  {Total weight of selected items}
8:    $ov = 0$  {Total values of selected items}
9:    $P' \leftarrow P$ ,  $n_{P'} \leftarrow n_P$ ,  $W_{P'} \leftarrow W_P$ 
10:  while  $ow < W_{P'}$  and  $n_{P'} > 0$  do
11:    Find  $s(P')$  using state aggregation strategy (Eqn. (8))
12:    Perform action  $i$  according to policy  $\pi(A_i|s(P'), \theta^t)$  and observe  $r(s(P'), A_i)$ 
13:    if  $i \leq n_{P'}$  and  $w_i + ow \leq W_{P'}$  then
14:       $ow \leftarrow ow + w_i$ 
15:       $ov \leftarrow ov + v_i$ 
16:       $W_{P'} \leftarrow W_{P'} - w_i$ 
17:    end if
18:     $P' \leftarrow P' \setminus \{i\}$ ,  $n_{P'} \leftarrow n_{P'} - 1$ 
19:    Update  $\theta$  and  $\theta_v$  using Eqns. (2) and (3)
20:     $t \leftarrow t + 1$ 
21:  end while
22:  if  $ov > Val[P']$  then
23:     $Val[P'] \leftarrow ov$ 
24:  end if
25: end while
26: return  $Val$ 

```

---

generate random solutions for each problem instance. As mentioned before, an episode is a sequence of states and actions that each action selects an item and the solution is the set of selected items. These instances can be shown in a table in which each row corresponds to a problem instance and the columns are items information.

One issue in selecting the feature vector of original items information as states is that different KP instances are not comparable because the values and weights of items might be very different. As an example, assume that values and weights of an instance are integer numbers between 1 to 10, while these values and weights lies between 100 and 110 for another instance. Generalization based on these different values is difficult, although their ratio are similar. In order to solve this issue, for each item of instance  $P$ , all  $v_i$  are normalized through dividing by the product of  $w_i$  and  $W_P$  as shown in Eqn. (4). Furthermore, the ratio between  $w_i$  and  $W$  is also calculated based on Eqn. (5). The  $v_i$  and  $w_i$  for each item are replaced with these two ratios in the feature vector of  $P$ . This modification



makes the items of different problems comparable. The learned policy network in this way would boost the capability of the well-known heuristic greedy algorithm which is optimal for fractional KP.

$$vr_i(v_i, w_i, W_P) = \frac{v_i}{w_i \times W_P} \quad (4)$$

$$wr_i(v_i, w_i, W_P) = \frac{w_i}{W_P} \quad (5)$$

where,  $vr_i$  and  $wr_i$  are the normalized value and normalized weight respectively. For a problem instance  $P$ ,  $vr_i$ ,  $wr_i$ ,  $W_P$ ,  $Sv$  and  $Sw$  construct a feature vector  $F(P)$ .

$$F(P) = (F_1, \dots, F_{2N+4}) = (n_P, W_P, Sv, Sw, vr_1, wr_1, \dots, vr_{n_P}, wr_{n_P}) \quad (6)$$

where,  $F(P)$  is the feature vector of  $P$ ,  $Sv$  and  $Sw$  are the sum of remained values and weights respectively.

After obtaining a table of problem instances with comparable items, we sort for each row (i.e. each problem instance) the columns (i.e.  $vr_i$  and  $wr_i$ ) in descending order with respect to  $vr_i$ . In other words, the first two columns of each row, i.e.  $vr_1$  and  $wr_1$  correspond to the item with highest normalized value. The second two columns which are  $vr_2$  and  $wr_2$ , correspond to the of item with the second highest normalized value and so on. For a problem  $P$  with  $n_P < N$ , the items information are located from the columns  $vr_1$  and  $wr_1$  to  $vr_{n_P}$  and  $wr_{n_P}$  respectively. The values of  $vr_{n_P+1}$  to  $vr_N$  and also the values of  $wr_{n_P+1}$  to  $wr_N$  are zero. This ordering helps to aggregate all the highest  $vr_i$  of all problem instances with a single aggregation strategy because the problem instances are comparable and the highest  $vr_i$  is in a certain column. This explanation holds for second, third and other highest  $vr_i$ . Each column is called a feature and the next step is to derive an aggregation strategy for the values of each feature.

*State aggregation through Q-Learning.* The idea of the aggregation is to reduce the number of unique values for all features. We do such reduction by splitting the values of one feature into several groups and then mapping each group's value to a particular integer. The proper number of splits for each feature is learned by reinforcement learning. For each feature  $F_k$  that  $k \in \{1, \dots, 2N + 4\}$ , let action  $d_{F_k}$  be the number of splits on the values of the feature  $F_k$ , and  $F_{k,P}$  be the value of feature  $F_k$  for problem instance  $P$ . Among all features, we perform state aggregation on  $vr_i$  and  $wr_i$  of item  $i$ .

For aggregating the values of  $vr_i$  of all  $M$  problem instances, action  $d_{vr_i} \in \{1, 2, \dots, x\}$  is the number of splits where its optimal value i.e.  $d_{vr_i}^*$  is obtained by Algorithm 2. Using  $d_{vr_i}^*$  splits, the values of  $vr_i$  are divided into  $d_{vr_i}^* + 1$  subsets and all the subsets except the last one have  $\lceil \frac{M}{d_{vr_i}^* + 1} \rceil$  values. The last subset has  $M - (\lceil \frac{M}{d_{vr_i}^* + 1} \rceil d_{vr_i}^*)$  values. Then, all values of each subset is converted to an integer starting from 0. This process transforms the values of feature  $vr_i$  to a set

of integers  $\{0, 1, \dots, d_{vr_i}^*\}$ . As an example, assume there are  $M = 7$  problem instances that the values of  $vr_1$  are  $(1, 2, 6, 3, 1, 2, 5)$  and  $d_{vr_1}^*$  is 2. These values need to be divided into  $d_{vr_1}^* + 1 = 3$  subsets. First they are sorted in order to acquire the sorted values  $(1, 1, 2, 2, 3, 5, 6)$ . Then, 3 subsets ( $\{1, 1, 2\}, \{2, 3, 5\}, \{6\}$ ) are obtained that each has  $\lceil 7/3 \rceil = 3$  values except the last one which has one value. Finally, the values of  $vr_1$  are aggregated and the new values are  $(0, 0, 2, 1, 1, 0, 1)$ . The aggregation reduces 5 unique values of  $vr_1$  to 3 unique values.

For all  $wr_i$ ,  $d_{wr_i}^*$  is 2 and the split points are 0.5 and 1. The motivation of this hard setting is separating illegal, light and heavy weights. Illegal weights are the weights with  $wr_i > 1$  that cannot be added to the knapsack. Similarly,  $wr_i \leq 0.5$  and  $0.5 < wr_i \leq 1$  determine light and heavy items respectively. The aggregation process is performed by the function  $map(F_k, P, d_{F_k}^*)$  that gets  $F_k, P$  and returns an integer which corresponds to a subset based on  $d_{F_k}^*$  splits.

We used heuristics to define the reward function  $R(F_k, d_{F_k})$  which is shown in Eqn. (7).

$$R(F_k, d_{F_k}) = \frac{\prod_{j=1}^{d_{F_k}+1} l_{F_k,j}}{(d_{F_k} + 1) \times c_{F_k, d_{F_k}}} \quad (7)$$

where,  $l_{F_k,j}$  is the size of  $j^{th}$  subset, and  $c_{F_k, d_{F_k}}$  is the number of all common values between all subsets. Three main motivations of designing rewards are as follows.

- We aim to define the reward function such that it reduces the size of state space. The number of unique states for each feature is  $d_{F_k} + 1$  after applying  $d_{F_k}$  splits and this value inversely relates to the reward of each action.
- For feature  $F_k$  and  $d_{F_k}$  splits, let  $j \in \{0, 1, 2, \dots, d_{F_k}\}$  be a subset based on  $d_{F_k}$  splits and  $l_{F_k,j}$  be the difference between maximum and minimum values of  $j^{th}$  subset. As larger  $l_{F_k,j}$  entail in aggregating more values, their rewards are higher than those for smaller  $l_{F_k,j}$ . However, unequal subsets contain unequal number of values. For example, if the feature values are uniformly dispersed between 0 and 10, creating two subsets with lengths 5 and 5 are better than two subsets with length 1 and 9. Therefore, the product of the  $l_{F_k,j}$  for all  $j$  is in the numerator of the reward function.
- Distinct states help an agent to derive a deterministic policy because states have dissimilar features. Likewise, two subsets with less overlapped values represent different sets of states and the policy can better distinguish them. For example, for the subsets ( $\{1, 1, 2\}, \{2, 3, 5\}, \{6\}$ ), 2 is common between two subsets and it can be assigned to both subsets. Assigning this value to different subsets entails a different policy that may have different performance. In order to reduce the number of common values between two groups, we define  $c_{F_k, d_{F_k}}$  as the total number of common values in different subsets.

A  $Q$  table is constructed for the states and actions and it is filled by the  $Q$ -learning algorithm [24] as shown in Algorithm 2. Each  $vr_i$  is a state and the

next state is the  $vr_{i'}$  which  $i'$  is an arbitrary state. Finally, an optimal decision is found by using the  $Q$  table for each feature. The algorithm is used for aggregating  $vr_i$  and we denote  $d_{vr_i}^*$  as the optimal aggregation action for each  $vr_i$ . The state embedding derived by this strategy is a feature vector consisting of aggregated features and this state embedding is used in line 11 of algorithm 1. Equation (8) shows  $s(P)$ , the state embedding of  $P$ .

$$s(P) = \{map(F_{k,P}, d_{F_k}^*) : \forall F_k \in F(P)\} \quad (8)$$

---

**Algorithm 2** Q-Learning for State Aggregation
 

---

**Input:** Feature table of problem instances  $P_1, \dots, P_M$

**Output:** The number of optimal split points for all  $vr_i$

- 1: Initialize a  $Q$  table with  $N$  rows and  $x$  columns. States are features and actions are the number of split points
  - 2: Select item  $i$  randomly
  - 3: **repeat**
  - 4:   Select  $i'$  randomly as the next item
  - 5:   Select  $d_{vr_i} \in \{1, \dots, x\}$  according to  $\epsilon$ -greedy policy
  - 6:   Find  $R(vr_i, d_{vr_i})$  using Eqn. 7
  - 7:   Update  $Q(vr_i, d_{vr_i}) \leftarrow Q(vr_i, d_{vr_i}) + \alpha[R_{vr_i, d_{vr_i}} + \gamma \max_{d'} Q(vr_{i'}, d') - Q(vr_i, d_{vr_i})]$
  - 8:    $i = i'$
  - 9: **until** Convergence
  - 10: **return**  $d_{vr_i}^* = argmax_d Q(vr_i, d) \forall i$
- 

## 5 Experiments

The DRL based knapsack solver is applied on three different types of problem instances. We tested different algorithms for training the policy DNN such as Deep Q Network (DQN) [20], Advantage Actor Critic (A2C) [19], Proximal Policy Optimization [22] and Sample Efficient Actor-Critic with Experience Replay [27], and the A2C algorithm is selected because it provided better solutions. We used *stable-baseline* tools for implementing the A2C algorithm [12]. The policy network consists of two layers of 64 nodes and the method is trained on  $10^4$  episodes which are selected from  $M$  instances. The DRL with aggregation algorithm is compared with (1) greedy algorithm and (2) DRL without aggregation. The problem instances and code used for experiments are available in URL<sup>1</sup>.

---

<sup>1</sup> The link is not shown due to the blind review.

## 5.1 Problem Instances

The three different types of instances are called *Random Instances*, *Fixed  $W_P$  Instances* and *Hard Instances*. A set of  $M$  problem instances makes a dataset that the maximum number of items over all instances in the dataset is  $N$ . Each dataset contains the instances of one of the following types.

*Random instances (RI)*: A dataset of random instances has  $M$  problem instances that each instance  $P$  has  $n_P \in \{1, 2, \dots, N\}$  items. For an item  $i$ ,  $v_i$  and  $w_i$  are randomly generated integers from one to  $R$  that is a fixed upper bound for  $v_i$  and  $w_i$ . The  $W_P$  is a random integer between  $R/10$  and  $3R$ . Three datasets of random instances are generated with  $M = 1000$ . For these three datasets,  $N$  is 50, 300 and 500, and  $R$  is 100, 600 and 1800 respectively.

*Fixed  $W_P$  Instances (FI)*: In [2] a set of KP instances with fixed capacity and fixed item set size are used for evaluation. We generated three datasets of the same instances with  $M = 1000$ . The  $N$  for these three datasets is 50, 300 and 500 respectively. The values and the weights of all items in the three datasets are random real numbers between zero and one. The  $W_P$  is fixed for all the instances and it is 12.5 for  $N = 50$ , 37.5 for  $N = 300$  and 37.5 for  $N = 500$ .

*Hard instances (HI)*: In [21], a group of hard to solve problem instances were introduced that for each item  $i$ ,  $v_i$  is strongly correlated with  $w_i$ . Specifically,  $w_i$  is a random integer in  $[1, R]$ ,  $v_i = w_i + R/10$  and  $W_P = \frac{p}{M+1} \sum_{i=1}^{n_P} w_i$  where  $p$  is the id of  $P$ . Three datasets of  $M = 1000$  hard instances are generated. For the first dataset,  $N$  is 50 and  $R$  is 100. Likewise,  $N$  is 300 and 500, and  $R$  is 600 and 1000 for the second and the third datasets respectively.

## 5.2 Evaluation Metrics

The following metrics are considered to evaluate the performance of using DRL based KP solver for solving the instances introduced in 5.1.

*Average values of solutions ( $\overline{Val}$ )*. For each dataset of  $M$  problem instances,  $\overline{Val}$  is the average of all solution values (total values of the selected items). Likewise,  $\overline{Val}_{opt}$  is the average values of optimal solutions, which are obtained using the optimization solver *Gurobi* [11].

*Learning rate*. In order to calculate the learning rate, the rate of increasing in  $Val$  is calculated per timesteps and the result is shown for each instance type when  $N = 300$ .

*Number of optimally solved instances ( $\#_{opt}$ )*. In order to evaluate the performance of the method on the individual problem instances, the number of instances that the method finds their optimal solution is computed for each dataset.


**Table 1.** Results of different algorithms and datasets of  $M = 1000$  problem instances. The method of [2] is not applicable on RI and HI while it is optimal for small  $N$  as well as the DRL w/ aggregation method. It is possible that two approaches find the optimal solution for a certain instance. Hence, the total number of optimally solved instances is not necessarily 1000.

Dataset	Method	$N$	$\overline{Val}$	$\#_{opt}$	$\#_{highest}$	$\overline{Val}_{opt}$	$\frac{Val}{Val_{opt}}$
RI	Greedy	50	429.10	596	0	434.78	98.694%
	DRL w/o aggregation	50	434.09	893	7	434.78	99.843%
	DRL w/ aggregation	50	<b>434.50</b>	<b>959</b>	<b>41</b>	434.78	99.937%
	Greedy	300	1144.96	418	0	1151.58	99.425%
	DRL w/o aggregation	300	1150.83	830	21	1151.58	99.934%
	DRL w/ aggregation	300	<b>1151.10</b>	<b>878</b>	<b>47</b>	1151.58	99.958%
	Greedy	500	15216.51	345	0	15285.56	99.548%
	DRL w/o aggregation	500	15273.47	701	30	15285.56	99.920%
	DRL w/ aggregation	500	<b>15278.44</b>	<b>786</b>	<b>80</b>	15285.56	99.953%
FI	Greedy	50	20.10	172	0	20.15	99.738%
	DRL w/o aggregation	50	20.14	740	36	20.15	99.931%
	DRL w/ aggregation	50	<b>20.15</b>	<b>773</b>	<b>54</b>	20.15	99.959%
	Greedy	300	86.26	202	0	86.31	99.942%
	DRL w/o aggregation	300	86.27	226	24	86.31	99.961%
	DRL w/ aggregation	300	<b>86.29</b>	<b>330</b>	<b>205</b>	86.31	99.976%
	Greedy	500	111.68	204	0	111.73	99.945%
	DRL w/o aggregation	500	111.63	64	31	111.73	99.871%
	DRL w/ aggregation	500	<b>111.70</b>	<b>261</b>	<b>144</b>	111.73	99.970%
HI	Greedy	50	772.428	134	0	802.72	96.226%
	DRL w/o aggregation	50	799.036	655	113	802.72	99.540%
	DRL w/ aggregation	50	<b>799.438</b>	<b>689</b>	<b>147</b>	802.72	99.591%
	Greedy	300	27778.03	37	0	27965.76	99.328%
	DRL w/o aggregation	300	27947.11	<b>370</b>	161	27965.76	99.933%
	DRL w/ aggregation	300	<b>27952.63</b>	336	<b>233</b>	27965.76	99.953%
	Greedy	500	80779.23	25	0	81103.99	99.781%
	DRL w/o aggregation	500	81022.60	<b>217</b>	168	81103.99	99.899%
	DRL w/ aggregation	500	<b>81064.99</b>	166	<b>304</b>	81103.99	99.951%

*Number of instances with highest solution value ( $\#_{highest}$ ).* This metric compares the value of solutions of DRL with aggregation and DRL without aggregation and counts the number of times that each one is higher. This value is calculated for the last  $M/2$  instances of the datasets. These instances are more difficult to solve among the hard problem instances because  $W_P = \frac{p}{M+1} \sum_{i=1}^{n_P} w_i$  is increasing with respect to  $p$ . Therefore, with larger  $W_P$ , the set of feasible solutions is bigger and hence finding the optimal solution is more difficult.

### 5.3 Results

We ran the algorithm with 1000 problem instances. Table 1 shows the quality of the solutions of different types of KP instances: RI, FI, and HI, that are obtained



**Fig. 2.** Random Instances

by DRL algorithms with (i.e. w/) and without (i.e. w/o) aggregation, and the greedy algorithm (Greedy).

Table 1 contains the ratio of  $\overline{Val}$  and  $\overline{Val}_{opt}$ . These values show that the ratios of the solutions provided by our proposed method (DRL w/ aggregation) and the optimal solutions are most of the times more than 99.9%. This ratio does not change considerably when the number of items increases. Hence we conclude that our DRL based approach is able to find very close to optimal solutions for all instances we tested.

*Comparison with [2].* The pointer network based DRL method [2] is also able to find close to optimal solutions for problem size up to  $N = 200$ . However, the method of [2] can only be applied to solve the instances with exactly same number of items  $N$ , and in addition, with exactly same capacity value  $W_P$ . In comparison, our DRL formulation allows to solve instances of any size up to and including  $N = 500$ , and of any capacity value  $W_P$ .

*Comparison with Greedy and DRL without aggregation.* The results show that the proposed DRL-based methods, with or without aggregation, always perform better than the greedy algorithm, in terms of the average solution quality ( $\overline{Val}$ ), the number of optimally solved instances ( $\#_{opt}$ ), and the number of instances with highest solution value ( $\#_{highest}$ ).

When we evaluate the advantage of having state aggregation, we notice that the state aggregation strategy improves the solutions especially for large instances, which is clearly observable in the solutions of instances RI and FI. Regarding to hard instances HI, the DRL with aggregation method is better than

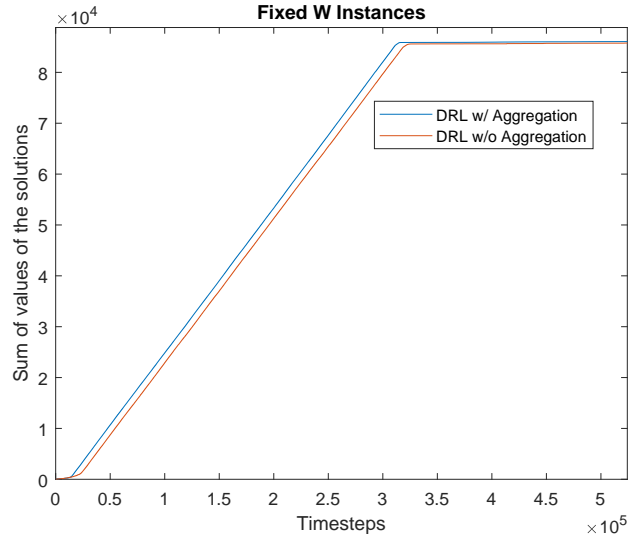


Fig. 3. Fixed  $W_P$  instances




Fig. 4. Hard instances

the one without aggregation strategy in terms of solution quality ( $\overline{Val}$ ). For the large instances of sizes 300 and 500, the DRL without aggregation finds more optimal solutions than the one with aggregation, with 370 vs 336 for  $N = 300$ , and 217 vs 166 for  $N = 500$ . However, when looking at their performances in terms of how many times they have the highest solution values for 500 more difficult instances in HI, DRL with aggregation performs better than without aggregation, with 72 more wins for  $N = 300$ , and 133 more wins for  $N = 500$ . We investigate these instances in HI further. We have mentioned that  $W_P$  is increasing with respect to  $p$  which is the identifier of each problem instance  $P$ , for  $M$  problem instances, this identifier ranges from one to  $M$ . When  $p$  is small,  $W_P$  is also small. Since  $vr_i$  indirectly relates to  $W_P$ , it would be large when  $W_P$  is small. For small  $W_P$ , the number of feasible solutions is low because less number of items can be fit into the knapsack. These problem instances are hence not actually “hard”. In this case, state aggregation is not beneficial as aggregating  $vr_i$  of the items of the instances with small  $p$  leads to sub-optimal solutions. However, problem instances with large  $p$  have larger feasible solution space, and hence aggregation is beneficial as it reduces the state size to enhance generalization.

The other benefit of DRL with aggregation method is that it is able to find the high quality solutions in less time steps. As it can be observed from Figures 2, 3 and 4, the learning rate of DRL with aggregation method is higher than the DRL without aggregation. Hence, in general, it not only provides better solutions, but also the solutions are found in around 10,000 less timesteps.

## 6 Conclusion and future work

In this paper we developed a DRL-based method for boosting the heuristic greedy algorithm and solving KP. In the DRL based KP solver, a policy DNN and a value DNN are trained using A2C algorithm and the policy DNN is used for sequentially selecting items to find a solution. The states in DRL modeling of KP contain the information of the instances that are aggregated to reduce the state space. The state aggregation policy is derived by solving a tabular RL problem. Using this aggregation policy, a state embedding is obtained and this state embedding is used with another RL framework to train the parameters of the policy network.

We applied this method on three types of problem instances named random instances, fixed  $W_P$  instances and hard instances. Three datasets with 50, 300 and 500 are generated for each type. The DRL with aggregation method found close to optimal solutions for the instances. It also found optimal solutions for fixed capacity instance with small number of items as well as the method developed by [2].

The proposed method can be generalized to other COPs. For instance, the TSP consists of some cities and the goal is to find the minimum length tour that visits every city exactly once. The cities might be the items and an aggregation strategy could reduce the state space by aggregating the coordinates. As another



example, in minimum vertex cover problem, the items may be the vertices and aggregation can be performed by grouping the weight of the vertices.

In this paper, we use RL to automate the reduction of the state space, as a pre-processing step of the DRL based approach for KPs. It is also interesting to investigate in the future how better reward functions can be derived through learning. This might be very helpful for problems that many tuning processes are needed for deriving a strong reward function. In general, automating the state, reward and action derivation for RL problems are interesting topics for research in the future.

## References

1. Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A.: A brief survey of deep reinforcement learning. arXiv preprint arXiv:1708.05866 (2017)
2. Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S.: Neural combinatorial optimization with reinforcement learning. In: ICLR (Workshop) (2017), <https://academic.microsoft.com/paper/2560592986>
3. Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. arXiv preprint arXiv:1811.06128 (2018)
4. Cook, S.: The p versus np problem. The millennium prize problems pp. 87–104 (2006)
5. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to algorithms. MIT press (2009)
6. Dai, H., Dai, B., Song, L.: Discriminative embeddings of latent variable models for structured data. In: International conference on machine learning. pp. 2702–2711 (2016)
7. Dasgupta, S., Papadimitriou, C.H., Vazirani, U.V.: Algorithms. McGraw-Hill Higher Education (2008)
8. Du, D.Z., Pardalos, P.M.: Handbook of combinatorial optimization: supplement, vol. 1. Springer Science & Business Media (2013)
9. Feng, Y., Yang, J., Wu, C., Lu, M., Zhao, X.J.: Solving 0–1 knapsack problems by chaotic monarch butterfly optimization algorithm with gaussian mutation. Memetic Computing **10**(2), 135–150 (2018)
10. Gu, S., Hao, T.: A pointer network based deep learning algorithm for 0–1 knapsack problem. In: 2018 Tenth International Conference on Advanced Computational Intelligence (ICACI). pp. 473–477. IEEE (2018)
11. Gurobi Optimization, L.: Gurobi optimizer reference manual (2020), <http://www.gurobi.com>
12. Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y.: Stable baselines. <https://github.com/hill-a/stable-baselines> (2018)
13. Huang, T., Ma, Y., Zhou, Y., Huang, H., Chen, D., Gong, Z., Liu, Y.: A review of combinatorial optimization with graph neural networks. In: 2019 5th International Conference on Big Data and Information Analytics (BigDIA). pp. 72–77. IEEE (2019)
14. Joshi, C.K., Laurent, T., Bresson, X.: An efficient graph convolutional network technique for the travelling salesman problem. arXiv preprint arXiv:1906.01227 (2019)

15. Karp, R.M.: Reducibility among combinatorial problems. In: Complexity of computer computations, pp. 85–103. Springer (1972)
16. Khalil, E., Dai, H., Zhang, Y., Dilkina, B., Song, L.: Learning combinatorial optimization algorithms over graphs. In: Advances in Neural Information Processing Systems. pp. 6348–6358 (2017)
17. Kool, W., van Hoof, H., Welling, M.: Attention, learn to solve routing problems! arXiv preprint arXiv:1803.08475 (2018)
18. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
19. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. In: International conference on machine learning. pp. 1928–1937 (2016)
20. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602 (2013)
21. Pisinger, D.: Where are the hard knapsack problems? *Computers & Operations Research* **32**(9), 2271–2284 (2005)
22. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
23. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Advances in neural information processing systems. pp. 3104–3112 (2014)
24. Sutton, R.S., Barto, A.G., et al.: Introduction to reinforcement learning, vol. 2. MIT press Cambridge (1998)
25. Vazirani, V.V.: Approximation algorithms. Springer Science & Business Media (2013)
26. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In: Advances in Neural Information Processing Systems. pp. 2692–2700 (2015)
27. Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., de Freitas, N.: Sample efficient actor-critic with experience replay. arXiv preprint arXiv:1611.01224 (2016)
28. Wilbaut, C., Hanafi, S., Salhi, S.: A survey of effective heuristics and their application to a variety of knapsack problems. *IMA Journal of Management Mathematics* **19**(3), 227–244 (2008)
29. Woeginger, G.J.: Exact algorithms for np-hard problems: A survey. In: Combinatorial optimization eureka, you shrink!, pp. 185–207. Springer (2003)
30. Yadav, N., Murawski, C., Sardina, S., Bossaerts, P.: Phase transition in the knapsack problem. arXiv preprint arXiv:1806.10244 (2018)