SCIENCE CHINA
Information Sciences

• **RESEARCH PAPER** •

Learning Practically Feasible Policies for Online 3D Bin Packing

Hang ZHAO[†1], Chenyang ZHU[†1], Xin XU[1], Hui HUANG[2] & Kai XU[*1]

[1]*National University of Defense Technology, Changsha* 410073*, China;*
[2]*Shenzhen University, Shenzhen* 518060*, China*

arXiv:2108.13680v1 [cs.RO] 31 Aug 2021

Abstract We tackle the Online 3D Bin Packing Problem, a challenging yet practically useful variant of the classical Bin Packing Problem. In this problem, the items are delivered to the agent without informing the full sequence information. Agent must directly pack these items into the target bin stably without changing their arrival order, and no further adjustment is permitted. Online 3D-BPP can be naturally formulated as Markov Decision Process (MDP). We adopt deep reinforcement learning, in particular, the on-policy actor-critic framework, to solve this MDP with constrained action space. To learn a practically feasible packing policy, we propose three critical designs. First, we propose an online analysis of packing stability based on a novel stacking tree. It attains a high analysis accuracy while reducing the computational complexity from $O(N^2)$ to $O(N \log N)$, making it especially suited for RL training. Second, we propose a decoupled packing policy learning for different dimensions of placement which enables high-resolution spatial discretization and hence high packing precision. Third, we introduce a reward function that dictates the robot to place items in a far-to-near order and therefore simplifies the collision avoidance in movement planning of the robotic arm. Furthermore, we provide a comprehensive discussion on several key implemental issues. The extensive evaluation demonstrates that our learned policy outperforms the state-of-the-art methods significantly and is practically usable for real-world applications.

Keywords Bin Packing Problem, Online 3D-BPP, Reinforcement Learning

Citation Hang Zhao, Chenyang Zhu, Xin Xu, Hui Huang, Kai Xu. Learning Practically Feasible Policies for Online 3D Bin Packing. Sci China Inf Sci, for review

1 Introduction

The bin packing problem (BPP) is one of the most famous problems in combinatorial optimization. It aims to pack a collection of items with various weights into the minimum number of bins. The total weight of items in each bin is below the bin's capacity c [26]. BPP is a classic NP-hard problem. We are interested in its 3D variant, i.e., 3D-BPP [31] which introduces much more solving complexity. Given item i with 3D "weight" about length, l_i, width w_i, and height h_i, 3D-BPP coordinates planning item's assignment in three dimensions simultaneously. Each 3D dimension has its capacity including $L \geqslant l_i$, $W \geqslant w_i$, and $H \geqslant h_i$. It is assumed that $l_i, w_i, h_i, L, W, H \in Z^+$. 3D-BPP also pursues to pack the set of items $\mathcal{I}$ with the fewest bins.

3D-BPP finds widely practical applications in modern packaging, logistics, and manufacturing. It is especially a core technique in developing palletizing robots for intelligent logistics (Figure. 1). A palletizing robot is designed to pack boxes into rectangular bins of standard dimension. Maximizing the storage use of bins improves production efficiency like inventorying, wrapping, transportation, and warehousing. Due to its computational complexity, 3D-BPP is relatively less explored than 1D-BPP. Especially when the problem scale increases, the exact algorithm (either using integer linear programming or branch-and-bound) cannot give a solution to the problem within a limited time. Solving the medium-scale 3D-BPP still has to resort to heuristic algorithms [9, 24].

* Corresponding author (email: kevin.kai.xu@gmail.com)
† Equal contribution.

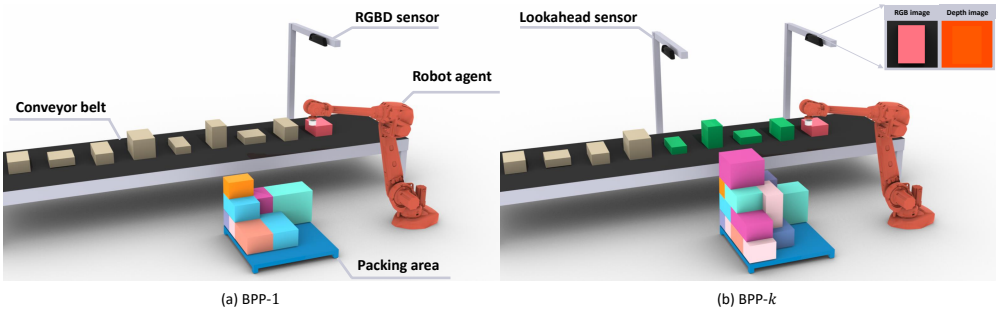

Figure 1 Online 3D-BPP has widely practical applications in logistics, manufacture, warehousing, etc. Left: The agent can only observe the next item to be packed (shaded in red). Right: More items (shaded in green) can be observed with additional sensors.

In many application scenarios, an even more difficult problem setting, *Online 3D-BPP*, is highly demanded. In online 3D-BPP, the information of the full item sequence is not provided to the agent/robot (similar to Tetris). As shown in Figure 1, where a robot packs continuously coming parcels online. The RGB-D sensors placed around the robot can only provide a partial vision of the item sequence due to the limited camera view field. Since the conveyor is forwarding sequentially during the packing, only a small period can be given for the robot to unload and place the parcel. Such constraints make online 3D-BPP not a pure combinatorial optimization problem since the optimal solution cannot be obtained by brute-force enumeration.

We approach online 3D-BPP by formulating it as a sequential decision making problem and solving it with deep reinforcement learning, similar to [50]. Albeit being quite effective, several limitations of [50], such as heuristic analysis of packing stability, limited resolution of spatial discretization, and collision agnostic of packing scheme, hinder the practical applicability of learned policy. To this end, we propose the following substantial enhancements to learn practically more feasible policies for online 3D-BPP:

- *First*, we impose a physically plausible yet highly efficient stability analysis to facilitate the learning of stable packing policies. In particular, we propose an online analysis of packing stability with a novel stacking tree. Stacking tree analysis attains a 99.9% accuracy of stability analysis with an $\mathcal{O}(N \log N)$ complexity (N is item count). It results in several magnitudes of acceleration of stability analysis over traditional static equilibrium analysis methods which are $O(N^2)$, making it especially suited for RL training. Meanwhile, the accurate analysis also leads to a significant boost of space utilization (by about 10%) over [50] based on hand-designed, conservative feasibility evaluation.
- *Second*, to deal with large action space caused by high-resolution bin discretization for accurate packing, we propose a novel decoupled packing policy learning scheme. It learns three packing policies for the length and the width dimensions, and the horizontal orientation of the item to be packed, respectively. The three policies are learned with conditionally probabilistic dependency between each other. This enables our method to work with up to 100×100 discretization resolution which is intractable for the method in [50].
- *Third*, to learn a more practically usable policy, we introduce a reward function which encourages the robot to place items in a far-to-near order. This greatly eases the collision avoidance between the robot and the packed items and simplifies the movement planning of the robotic arm.

Our method is formulated as a constrained Markov decision process (CMDP) [2] and adopts the on-policy actor-critic framework [34, 48]. Different from [50], we compute the feasibility mask for the placement actions based on the stacking tree analysis. The feasibility mask is then provided to the actor networks and then used to modulate the action probabilities output by the actor. We also discuss several practical issues in the real robot implementation of our algorithm. Finally, we conduct extensive evaluations to validate the efficacy of the new designs.

2 Related Work

Bin Packing Problem (BPP) is a long-term concern in combinatorial optimization, and the earliest literature can be traced back to the sixties [23]. The most typical bin packing problem is 1D-BPP which seeks for an assignment of a collection of items with various scalar weights to multiple bins and minimizes the used bin number. Knowing to be strongly NP-hard, most existing literature focuses on

designing good heuristic and approximation algorithms and their worst-case performance analysis [7]. Bin packing problem exists many variants, 2D- and 3D-BPP are natural generalizations of them. For high-dimension BPP, an item has a high-dimension size of the width, height, and/or depth, which differentiates the verification of the packing feasibility. The complexity and the difficulty significantly increase for high-dimension BPP instances. According to the timing statistic reported in [31], exactly solving 3D-BPP of a size matching an actual parcel packing pipeline remains infeasible. There have been several strategies in designing fast approximate algorithms, e.g., guided local search [14], greedy search [11], and tabu search [10, 30]. In contrast, genetic algorithms lead to better solutions as a global, randomized search [28, 40].

Similar strategies have also been adapted to Online BPP works like [17, 20, 46, 50]. Different from the offline setting, the size information of coming items is unknown for the agent to optimize the packing and the packing order can not be adjusted as well. It makes Online BPP a much more challenging problem. Some works [17,46] have employed the hand-coded heuristics based on the human experience. Meanwhile, works like [50] have adopted deep reinforcement learning to learn how to pack things effectively through trials and error optimization.

Stability Estimation is essential in the bin packing problem. However, most previous work focuses on the 3D BPP ignores it in their solution due to the high computational load. Estimating the stability of stack objects for physics engines is widely studied in past decades. [13] introduces novel complementarity formulation, solver, and error correction algorithms for collision detection framework which can enable the stability estimation running in real-time for large scale objects. In contrast, [42] presents a real-time physics engine to improve structured stacking behavior with small-scale objects. [21] present a constraint-based method to stabilize a stack of piles. [19] introduces a hypothesis for physical simulation simplification which is freezing transformations of objects in a random pile does not affect the visual plausibility of a simulation. However, none of them can be efficient enough for stability estimation in a DRL training framework.

Deep Reinforcement Learning (DRL) has demonstrated tremendous success in learning complex behavior skills and solving challenging control tasks with high-dimensional raw sensory state-space [29,34, 35]. The success can be attributed to the utilization of high-capacity deep neural networks for powerful feature representation learning and function approximation. Since the seminal work of deep Q-network (DQN) [35], a large body of literature has emerged. The existing research is largely divided into two lines: value function learning [35, 47] and policy search [3, 38]. Actor-critic methods, designed to combine the two approaches, have grown in popularity. Asynchronous Advantage Actor-Critic (A3C) [34] is a representative actor-critic method. It combines advantage updates with the actor-critic formulation and adopts asynchronously updated policy and value function networks trained in parallel with multiple agents. In A2C [37], on the other hand, the networks of multiple agents are updated synchronously.

We base our actor-critic architecture on ACKTR [48] which applies trust-region optimization to both the actor and the critic. We propose a series of adaptions for solving online 3D BPP. To realize constrained reinforcement learning, we propose a simple approach by projecting the trajectories sampled from the actor to the constrained state-action space. To enable our agent to fit the high-resolution demand for real packing, we decompose the actor into three actor-heads and get them to predict the corresponding actions in sequence. We also feed the last predicted action into the next actor-head as a conditional probabilistic dependency to achieve a more stable training process.

RL for Combinatorial Optimization has booming development in recent years. Bello et al. [4] combine *RL pretraining* and *active search* and demonstrate that RL-based optimization outperforms the supervised learning framework when tackling NP-hard combinatorial problems. Kool et al. [25] encode TSP graphs with Transformer [44] and train this model with RL method. Zhang et al. [49] enable an end-to-end RL agent to master priority dispatching rules for solving Job-shop scheduling problem. Wang et al. [45] apply RL to automatically generating a move plan for indoor scene arrangement. [22] and [27] are both devoted to solving offline 3D-BPP where the main goal is to find an optimal sequence of items.

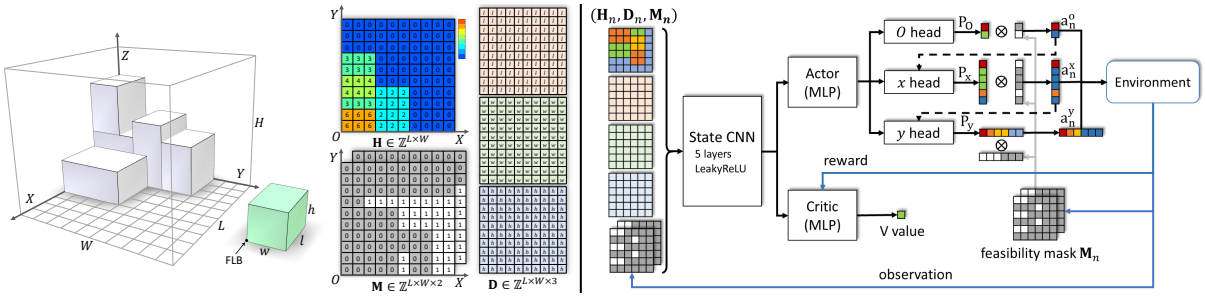

Figure 2 The environment state of the agent. The grey boxes indicate the items already packed, which also represents the bin configuration. The green box is the next item to be packed and it can only be placed at the grid cell where feasibility mask $\mathbf{M}$ is 1. Right: The network architecture with decomposed actor-heads. Note that the training of three actor-heads is coupled with conditional probabilistic dependencies, actor-heads perform their prediction tasks in sequence.

3 Method

We adopt a similar problem configuration for online 3D-BPP as [50]. Each item $n \in \mathcal{I}$ is a cube whose size is $[l_i, w_i, h_i]$. As soon as an item arrives the packing area, the agent needs to place it to the bin immediately while only being aware of the information of next few coming items $\mathcal{I}_o \subset \mathcal{I}$. Our method is implemented based on the architecture of constrained DRL and we will start our problem statement and formulation under the context of DRL.

3.1 Problem Statement and Formulation

The DRL formulation of our method can be expressed as $(\mathcal{S}, \mathcal{A}, P, R)$, which is a Markov decision progress. $\mathcal{S}$ is constructed with a set of environment states $\mathcal{S}$, the action set $\mathcal{A}$, reward function $R : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, and transition probability function $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$. $P(s'|s, a)$ gives the probability of transiting from s to s' for given action a. Our method is model-free since we do not learn $P(s'|s, a)$ explicitly. The policy $\pi : \mathcal{S} \to \mathcal{A}$ is a map from states to probability distributions over actions, with $\pi(a|s)$ denoting the probability of selecting action a under state s. For DRL, we seek for a policy π to maximize the accumulated discounted reward, $J(\pi) = E_{\tau\sim\pi}[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$. Here, $\gamma \in [0, 1]$ is the discount factor, and $\tau = (s_0, a_0, s_1, \ldots)$ is a trajectory sampled based on the policy π.

Environment State A complete 3D-BPP state representation should include the following three parts: the current configuration of the bin, the coming items to be placed, and the feasibility mask. To parameterize the bin configuration, we discretize its bottom area as an $L \times W$ regular grid along the length (X) and the width (Y) directions, respectively. We record at each grid cell the current height of stacked items, leading to a 2D integer *height map* $\mathbf{H}_n \in \mathbb{Z}^{L\times W}$ (see Figure 2). The dimensionality of item n is given as $\mathbf{d}_n = [l_n, w_n, h_n]^\top \in \mathbb{Z}^3$. The feasibility mask $\mathbf{M}_{n,o}$ is a binary matrix of size $L \times W$ indicating the placement feasibility of n with orientation o at each grid cell. Putting together, the current environment state can be written as $s_n = \{\mathbf{H}_n, \mathbf{d}_n, \mathbf{d}_{n+1}, ..., \mathbf{d}_{n+k-1}, \mathbf{M}_{n,o}\}$. We first consider the case where $k = |\mathcal{I}_o| = 1$ (Figure 1 (a)), and name this special instance as **BPP-**1. In other words, **BPP-**1 only considers the immediately coming item n i.e., $\mathcal{I}_o = \{n\}$. We then generalize it to **BPP-**k with $k > 1$ (Figure 1 (b)) afterwards.

Action and State Update We consider only horizontal, axis-align orientations of an item, which means that each item n has two possible orientations $o_n(d_n) = \{[l_n, w_n, h_n]^\top, [w_n, l_n, h_n]^\top\}$. During the packing, the agent places orientated n's front-left-bottom (FLB) corner (Figure 2 (left)) at a certain grid cell or the loading position (LP) in the bin. For instance, if the agent chooses to put n at the LP of (x_n, y_n) with the orientation adjustment o_n, this action is represented as $a_n = (x_n, y_n, o_n)$. The range of $\mathcal{A}$ would increase dramatically with a large resolution of x_n, y_n. In other words, it would be difficult to optimize action while high place accuracy is needed. To solve this problem, we propose an action space decomposition method in Section 3.3. After a_n is executed, $\mathbf{H}_n$ is updated by adding h_n to the maximum height over all the cells covered by n: $\mathbf{H}'_n(x, y, o) = h_{\max}(x, y, o) + h_n$ for $x \in [x_n, x_n + (1 - o_n)l_n + o_n w_n], y \in [y_n, y_n + (1 - o_n)w_n + o_n l_n]$, with $h_{\max}(x, y)$ being the maximum height among those cells.

Feasibility Constraint A practical online BPP solution should also premeditate the stability of a placement besides securing enough valid space for future items. Placing an item in a risky LP will end the packing episode early and even cause economic damage in practice. We propose a stacking tree based stability estimation method in Section 3.2 to form feasibility masks (see Figure 2) as optimization constraints to avoid insecure behaviors. Our problem becomes a constrained Markov decision process (CMDP) [2] in this scenario. Typically, one augments the MDP with an auxiliary cost function $C : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ mapping state-action tuples to costs, and require that the expectation of the accumulated cost should be bounded by c_m: $J_C(\pi) = E_{\tau\sim\pi}[\sum_{t=0}^{\infty} \gamma_C^t C(s_t, a_t)] \leqslant c_m$. Several methods have been proposed to solve CMDP based on e.g., algorithmic heuristics [43], primal-dual methods [6], or constrained policy optimization [1]. While these methods are proven effective, it is unclear how they could fit for 3D-BPP instances, where the constraint is rendered as a discrete mask. In this work, we propose to exploit the mask $\mathbf{M}$ to guide the DRL training to enforce the feasibility constraint without introducing excessive training complexity.

3.2 Stability Estimation

Stack stability estimation is critical in the bin packing problem, especially in 3D. A good stability estimation for LPs of n would not only secure the safety of the placement but also decrease the searching range of the action space. To calculate $\mathbf{M}_{n,o}$, the most straightforward solution is to simulate the force analysis among packed items $\mathcal{I}_{\text{packed}} \subset \mathcal{I}$ while n is placed with orientation o. However, this simulation would become extremely complicated while the $|\mathcal{I}_{\text{packed}}|$ increasing since force analysis with dense structured stacking is an NP-hard problem [13]. To ensure the stack stability estimation for LPs is real-time, we propose a centroid-based approach which improves the efficiency by more than 100 times.

Supported Centroid In industrial bin packing applications, item n to be packed usually has uniform mass distribution. An item n is stable if its centroid is supported in this scenario. Specifically, an LP of n is considered stable if satisfies any of the following conditions: 1) The centroid c_n of n is directly supported by a packed item with this LP. 2) n is supported by a group of items $\mathcal{I}_{\text{support}} \subset \mathcal{I}_{\text{packed}}$ and c_n is inside the convex hull which is constructed by contact points of $\mathcal{I}_{\text{support}}$, as illustrated in Figure 3. The centroid-based discrimination is easy to implement with the boolean operation and can be highly parallelized.

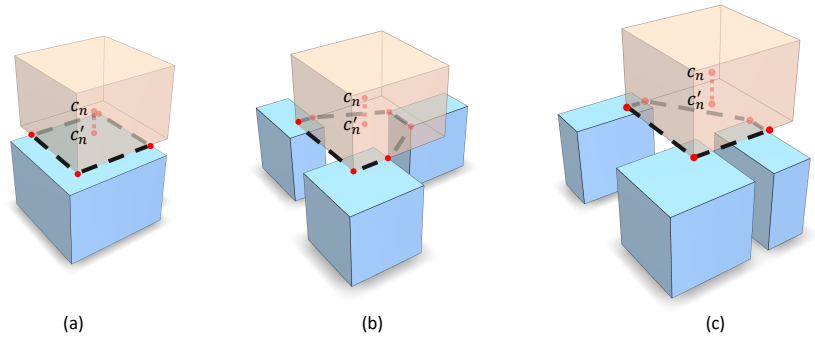

Figure 3 (a): The centroid of n is supported by a packed item directly therefore stable. (b): The vertical projection c'_n of centroid c_n is inside the convex hull which is constructed by contact points of $\mathcal{I}_{\text{support}}$. (c): c'_n falls outside the supported convex hull, n is unstable in this situation.

Adaptive Stacking Tree However, we not only need to estimate the stability of the current LP but also calculate the stability changes of $\mathcal{I}_{\text{packed}}$ which is introduced by the new contact. The mass of n changes the centroids of items which support it, and these changes may alter their stability. In other words, a top-down traverse of stability update is necessary. However, the amount of work done in this fashion is $\mathcal{O}(N^2)$ since a whole traverse is needed for each placed item. We propose an adaptive stacking tree structure to update the mass distribution flow of $\mathcal{I}_{\text{packed}}$ efficiently in $\mathcal{O}(N \log N)$. The key idea here is that the mass of n at LP would only distribute to the items $\mathcal{I}_{\text{active}} \subset \mathcal{I}_{\text{packed}}$ which support it, directly

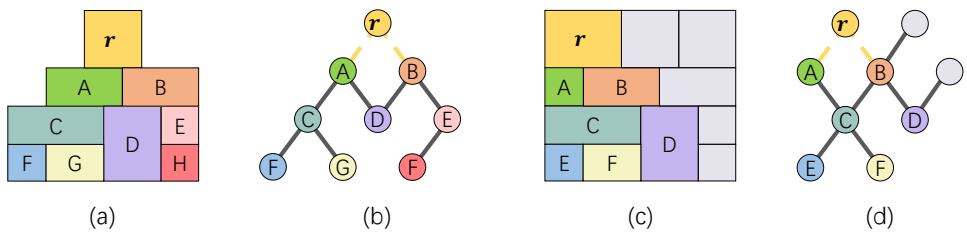

Figure 4 (a): r is the current item to be packed which will trigger an adaptive stacking tree to be updated. (b): The adaptive stacking tree of (a), the mass of r flows along each graph edge. (c): Placement of r does not change the mass flow over items with gray color. (d) The adaptive stacking tree of (c), the mass distribution of colored nodes needs to be updated. Note that the mass of inactive items (gray ones) will distribute to their adjacent active items while the update.

or indirectly. The stability of $\mathcal{I}_{\text{packed}} - \mathcal{I}_{\text{active}}$ would not change, and the stability update should only be done with $\mathcal{I}_{\text{active}}$ but not the whole packed bin.

In Figure 4, the mass distribution flow of $\mathcal{I}_{\text{packed}}$ can be formed as a graph $\mathcal{G}$, the nodes represent the items in the bin and the edges indicate the amount of mass distribution between two adjacent items. As we discussed above, only a subgraph $\mathcal{G}_n$ of $\mathcal{G}$ needs to be updated while n is placed at LP, where $\mathcal{G}_n$ only contains items in $\mathcal{I}_{\text{active}}$. Obviously, $\mathcal{G}_n$ is a tree which we denote as the adaptive stacking tree. $\mathcal{G}_n$ can be constructed from n with a top-down fashion, and only edges of $\mathcal{G}_n$ need to be updated while n is placed. The mass distribution update is processed from the root n to the bottom in $\mathcal{G}_n$ based on the calculated mass distribution, while the edges in $\mathcal{G} - \mathcal{G}_n$ should stay unchanging. This process can be easily implemented during packing in an incremental fashion based on a linked list data structure, which is both memory and time efficient.

Stability Update Updating mass distribution in $\mathcal{G}_n$ is critical for stability estimation. However, it would be highly time-consuming if we operate a force simulation here. We formulate a simple but effective approach based on the principle of leverage, which can achieve 99.9% accuracy with 100 times efficiency. While n is placed at LP, three types of mass distribution would be discussed next. If n is only supported by a single item n_0, the whole mass m_n of n and items above it would transport to n_0. The new centroid of the group of $\{n, n_0\}$ would change to,

$$c_{\{n,n_0\}} = \frac{c_n \cdot m_n + c_{n_0} \cdot m_{n_0}}{m_n + m_{n_0}} \tag{1}$$

Note that we only calculate centroid in the XY plane since the value along the Z axis would not influence the stability. If n is supported by two items n_0 and n_1 safely, we assume the mass distribution of m_n would obey the principle of leverage. As shown in right, n_0 needs to undertake F_0 to make the whole system stable.

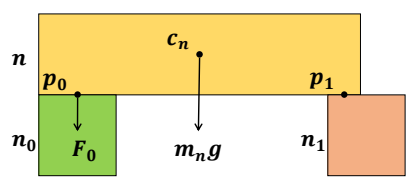

$$F_0 = g \cdot \|(c_n - p_1) \times m_n\| / \|p_0 - p_1\| \tag{2}$$

$$M\{n, n_0\} = F_0/g + m_{n_0} \tag{3}$$

where p_0 and p_1 are two contact points of n_0 and n_1. We can also calculate F_1 similarly. Then the centroid of group $\{n, n_0\}$ and $\{n, n_1\}$ can be updated with Equation(1) as well as the mass flow $M\{n, n_0\}$ and $M\{n, n_1\}$ which distributed to n_0 and n_1. If n is safely supported by more than two items $\{n_k\}$, the mass distribution of m_n would be calculated with Equation(2) for each two items in $\{n_k\}$. Then we will have $C_k^2 + 1$ constraints for mass distribution, a least squares optimization would be adopted to find the optimal F_i for each item in $\{n_k\}$. The centroid of group $\{n, n_i\}$ can be updated as well. Then each group $\{n, n_i\}$ would trigger the next iteration of update as well as n, this process would stop if the update reaches the bin bottom or instability is detected.

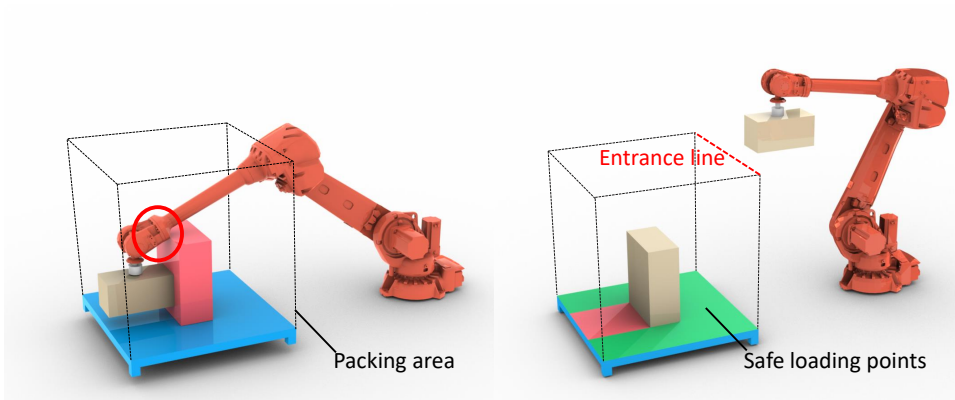

Figure 5 Left: The red box packed in the middle of the bin first introduces potential collisions for the following packing. Right: If the robot always enters the packing area at the same entrance line, packing on the green area may introduce fewer potential collisions.

3.3 Network Architecture and Training Method

We adopt ACKTR [48], which is a state-of-the-art on-policy framework. It iteratively updates an actor and a critic using Kronecker-factored approximate curvature (K-FAC) [33] with trust region. The actor is trained to learn a policy network that outputs the probability of choosing each action (i.e., placing n at each LP). The critic trains a state-value network predicting the state value $V(s_n)$ to indicate how much reward is earned from s_n, it is adopted to train our actor network. [50] has demonstrated that ACKTR has a surprising superiority over other DRL algorithms like SAC [18].

The BPP state consists of three components, the height map $\mathbf{H}_n$, the coming items to be placed $\mathbf{D}_n$, and the feasibility mask $\mathbf{M}_{n,o}$. We use a Convolutional Neural Network (CNN) to encode the raw BPP state. For calculation convenience, we "stretch" $\mathbf{d}_n$ into a three-channel tensor $\mathbf{D}_n \in \mathbb{Z}^{L\times W\times 3}$ so that each channel of $\mathbf{d}_n$ spans an $L \times W$ matrix with all of its elements being l_n, w_n, or h_n, respectively (also see Figure 2 (left)). Consequently, state $s_n = (\mathbf{H}_n, \mathbf{D}_n, \mathbf{M}_{n,o})$ becomes an $L \times W \times 6$ array (Figure 2 (right)).

Reward Shaping We want the agent to learn practical skills which meet the demands of both safety and efficiency without influenced by various robot type. The efficiency need is satisfied by the reward term which is about the volumetric occupancy introduced by the current item: $10 \times l_n \cdot w_n \cdot h_n/(L \cdot W \cdot H)$ for item n. [50] reports that the step-wise reward is superior to a termination one (e.g. the final space utilization).

We found that if the agent is only trained with r_n, the actor network would intend to place items uniformly on the XY plane of the bin. However, the robot arm is usually equipped aside rather than above the bin in practice, packing items near the robot arm first would potentially introduce collisions as shown in Figure 5. To avoid it, [32] propose that the packing should start from the farther corner as possible. Nonetheless, the performance would drop significantly if we force the actor to follow this principle strictly. To balance the two factors, we introduce a side reward to give our agent a soft constraint.

We assume that the robot arm would always enter the packing area at the same entrance line (EL) (Figure 5). If there is no obstacle on the straight line from EL to a feasible LP, the packing item at this LP can hardly introduce collisions and this type of LP is a safe LP. In other words, the more LPs which satisfy this principle the more likely our next packing is collision-free. We denote the sum of upon volume of all safe LPs as V_{safe}(Figure 5), our side reward is formulated as $r'_n = V_{\text{safe}}/(L \cdot W \cdot H)$ to ensure the maximum amount of safe LPs. The final reward of the system is then revised as $r_n = \alpha \times l_n \cdot w_n \cdot h_n/(L \cdot W \cdot H) + \beta V_{\text{safe}}/(L \cdot W \cdot H)$, where $\alpha = 10$ and $\beta = 0.1$ since the packing performance is still the primary. The packing episode ends when the current item is not placeable and the r_n is zero.

Action Space Decomposition The resolution choice of $\mathbf{D}_n$ would significantly influence the range of action space $\mathcal{A}$. However, we need a high resolution of $\mathbf{D}_n$ to ensure the packing accuracy in practical applications which would dramatically enlarge $\mathcal{A}$. To solve this problem, we propose a multi-task training method. The action prediction task is decomposed into three subtasks, predict x_n, y_n and o_n respectively (Figure 2). The decomposition would reduce the action space from $O(n^3)$ to $O(3n)$. Different from

previous work [41] which learns state-action value functions separately for each task, we train different actors for each task and use unified critic function $V(s_n)$ for the three individual actors to predict P_x, P_y, P_o. Note that the training of three actors is coupled as well, the advantage is that the x_n, y_n, and o_n can be predicted in a correlated fashion which is reasonable in our problem configuration. The policy gradient for actor network parameters training to predict P_x, P_y, P_o is formulated as:

$$\nabla\theta_{\text{actor}} = (r_n + \gamma V(s_{n+1}) - V(s_n))\nabla \log P_{\text{actor}}(a_n|s_n) \tag{4}$$

where $\theta_{\text{actor}}, \text{actor} \in \{x, y, o\}$ is the tunable parameters for different actors, $\gamma \in [0, 1]$ is the discount factor and we set γ as 1 so that $V(s_n)$ can directly present how much reward can agent obtain from s_n on.

Loss function The whole network is trained via a composite loss. The actor loss L_{actor} is inferred from Equation(4) and the critic loss is constructed based on our reward function r_n, which are the loss functions used for training the actor and the critic, respectively. Next, we use the feasibility mask $\mathbf{M}_n$ given by stability estimation to modulate outputs of three actors in order, i.e., the probability distribution of the actions. In theory, if the LP at (x, y) is infeasible for n with orientation adjustment o, the corresponding probability $P(a_n^o = o|s_n), P(a_n^x = x|s_n), P(a_n^y = y|s_n)$ should be set to 0 respectively. However, we find that setting P to a small positive quantity like $\epsilon = 10^{-20}$ works better in practice – it provides a strong penalty to an invalid action but a smoother transformation beneficial to the network training. Our loss function is defined as:

$$L = \alpha \cdot L_{\text{actor}} + \beta \cdot L_{\text{critic}} + \omega \cdot E_{\text{inf}} - \psi \cdot E_{\text{entropy}}. \tag{5}$$

$$\begin{cases} L_{\text{actor}} &= (r_n + \gamma V(s_{n+1}) - V(s_n)) \log P_{\text{actor}}(a_n|s_n) \\ E_{\text{inf}} &= \sum\limits_{\mathbf{M}_{n,o}(x,y)=0} P_{\text{actor}}(a_n|s_n) \\ E_{\text{entropy}} &= \sum\limits_{\mathbf{M}_{n,o}(x,y)=1} -P_{\text{actor}}(a_n|s_n) \cdot \log\left(P_{\text{actor}}(a_n|s_n)\right) \\ L_{\text{critic}} &= (r_n + \gamma V(s_{n+1}) - V(s_n))^2 \end{cases} \tag{6}$$

where actor $\in \{x, y, o\}$. To further discourage infeasible actions, we explicitly minimize the summed probability at all infeasible LPs with E_{inf}. Meanwhile, the action entropy loss E_{entropy} is adopted to push the agent to explore more LPs. In this way, we stipulate the agent to explore only feasible actions. We recommend the following parameters which lead to consistently good performance throughout our tests: $\alpha = 1$, $\beta = 0.5$, and $\omega = \psi = 0.01$.

3.4 BPP-k with $k = |\mathcal{I}_o| > 1$

In a more general case, the agent receives the information of $k > 1$ *lookahead* items (i.e., from n to $n + k - 1$). With additional information embedded into the environment state, the agent is expected to learn the policy $\pi(a_n|\mathbf{H}_n, \mathbf{d}_n, ..., \mathbf{d}_{n+k-1})$ and have better performance. [50] claims that simply encoding the lookahead information into the network input cannot help. We adopt the search-based solution to enable the agent to leverage lookahead information explicitly.

Virtual Placement Order [50] claims that the current item n's placement should be conditioned on the next $k - 1$ one. They enumerate different permutations of the sequence $(\mathbf{d}_n, ..., \mathbf{d}_{n+k-1})$ and drive the actor network to give related plans. To evaluate each sequence, they sum up the accumulated reward and the critic value of the end state after the k-th item is placed. The most promising a_n can be found in the sequence with the highest evaluation score. Note that only n's placement is determined in one permutation search and the actual placement of the k items still follows the order of arrival. To make the search scalable when k is large, [50] adapts the Monte Carlo tree search (MCTS) [39] to this problem and scales down the computational complexity from $O(k!)$ to $O(km)$ where m is the number of permutations sampled.

Adaptions for MCTS We have made some adjustments to [50]'s MCTS method so that it can be employed in practical scenarios. Firstly, we multiply each virtual item's mass with an extremely small coefficient (10^{-6}) during the search so that the stability estimation for n will not be influenced by nonexistent items. The mass coefficient cannot be zero otherwise the virtual items will lose physical constraint with other virtual ones. We also noticed that calculating the feasibility mask based on height map update and stability estimation in MCTS is inevitably time-consuming for real-time since this process would be invoked thousands of times in a whole MCTS. To improve the efficiency, we adpot a root-based parallelization for the MCTS. Different from leaf-based parallelization and tree-based parallelization [5], there is no communication between clones during the parallelization. All the root children of the separate Monte Carlo trees are merged with their corresponding clones after every parallelization iteration. After the search, we choose the action a_n corresponding to the permutation with the highest path value. MCTS allows a scalable lookahead for BPP-k with a complexity of $O(km)$ where m is the number of permutations sampled.

However, parallelizing MCTS will sacrifice the sampling frequency per thread and harm the performance. We modify the MCTS node expansion strategy to ensure that MCTS focuses more on meaningful sequences. The items coming soon in actual order will be more likely to be sampled during the permutation search since a partial sequence given by a virtual order is sometimes meaningless. The sample possibility of $item(v_i)$ is $P_{\text{sample}}(x) \sim \mathcal{N}(0,1)$ in our implementation where x means $item(v_i)$'s sorted arrival index among unselected items. This attention-based fashion would give MCTS a soft constraint to sample more items with the actual order.

3.5 Implementation

We implement our system in a practical industry environment to validate the actual packing performance. Figure 6 illustrates the online autonomous bin packing system with an RGB-D sensor and a robot arm.

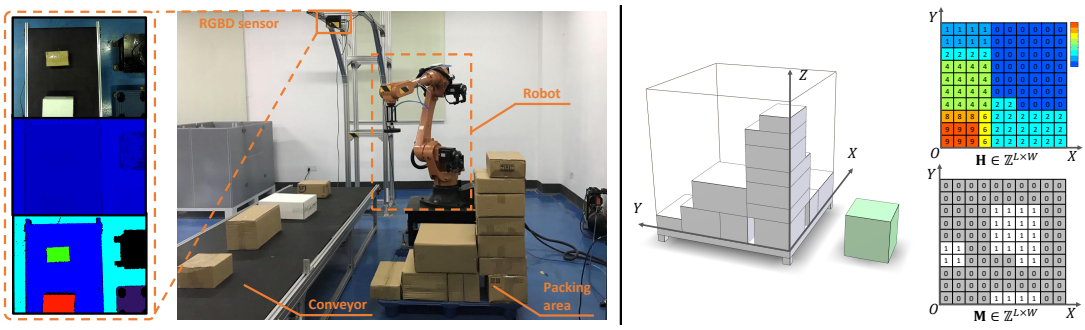

Figure 6 We implement our system in a practical industry environment. Left: The online autonomous bin packing system with an RGB-D sensor and a robot arm. Right: The digital twin which constructed by our method based on the captured data.

Environment Configuration The model of the RGB-D sensor is Percipio®FM811-GIX-E1 whose depth capture resolution is 1280 × 960. The robot arm is STEP®SR20E which can pick and place box-like items (up to 10kg) with an air pump driven suction cup. Our packing control system is implemented on a desktop computer (`ubuntu 16.04`), which equips with an `Intel Xeon Gold 5115` CPU @ 2.40 GHz, 64G memory, and an `Nvidia Titan V` GPU with 12G memory. The size of box items we adopted in our test meets the transportation standard of S.F.Express®and Taobao®, which ranges from 20cm to 50cm on each dimension. The items would come continuously and randomly with the conveyor belt at a fixed speed and will be packed by the robot arm to the bin placed on the floor which is driven by our system.

Terminals Our system contains three terminals. The robot terminal drives the robot arm with Codesys, which is a development environment for programming controller applications. The planning terminal is Python-based, which runs the DRL network we proposed to find the optimal packing strategy. The detection terminal is working with the RGB-D sensor, which segments the input depth image and recognizes the size and location of the coming items. These three terminals are communicated through a TCP protocol which centered on the detection terminal.

Executive Procedure Before running the whole system, an alignment between our three terminals is required. A hand-eye calibration [12] is performed between the detection terminal and robot terminal, thus the coordinate system of the object recognized by the RGB-D sensor can be aligned with the robot coordinate system. Meanwhile, we measure the relative position between the pallet and the robot and align the pallet coordinate system with the robot coordinate system as well.

The detection terminal would capture the depth image of items on the conveyor belt, and the PEAC method [15] is performed to judge whether there is an item in the field and measure its size. If an item is detected, the detection terminal would stop the conveyor belt and send the item size $\mathbf{d}_n = [l_n, w_n, h_n]^\top \in \mathbb{Z}^3$ to the planning terminal for packing strategy search.

While the planning terminal is calculating the optimal LP for the input item n, the robot terminal would drive the robot arm to pick this item up. It would take few seconds for the robot to complete the picking and the planning terminal can finish the calculating during this period. Once the robot arm leaves the camera's field of view, the detection terminal will restart the conveyor belt and the robot terminal would pack the picked item to the LP given by our planning terminal. At normal speed, it takes 8-9 seconds for our system to pack a box item from the conveyor belt to the pallet while human palletizers need to spend over 11 seconds without counting the rest time.

4 Experiments

Our DRL agent is trained in `PyTorch` [36]. Training on a spatial resolution of 100×100 takes about 12 hours and a single decision time is less than 10 ms. We compare our method with some state-of-the-art online bin packing methods to demonstrate the superiority. OnlineBPH [17] and OnlineDRL [50] are two representatives, where OnlineBPH is non-learning based and OnlineDRL learns how to pack from the data. [50] also proposes a heuristic baseline called *boundary rule* method. It replicates human's behavior by trying to place a new item side-by-side with the existing packed items and keep the packing volume as regular as possible.

4.1 Training and Test Set

We set $L = W = H = 100$ in our experiments with 125 pre-defined item dimensions ($|\mathcal{I}| = 125$). The 100×100 resolution is large enough for most of the packing applications in the real world. To avoid over-simplified scenarios, we limit $l_i \leqslant L/2$, $w_i \leqslant W/2$, and $h_i \leqslant H/2$. The training and test sequence is synthesized by generating items out of $\mathcal{I}$, and the total volume of items should be equal to or bigger than the bin's volume. We employ three types of data proposed by [50] to train and evaluate our method. One of them is called **RS** where the sequences are generated by sampling items out of $\mathcal{I}$ randomly. Due to the fact that the optimality of an RS sequence is unknown, the other two types of sequence are generated via *cutting stock* [16]. Specifically, items in a sequence are created by sequentially "cutting" the bin into items of the pre-defined 125 types so that we understand the sequence may be perfectly packed and restored to the bin. **CUT-1** sorts the cut items into the sequence based on Z coordinates of their FLBs, from bottom to top. **CUT-2** sorts the cut items on their stacking dependency. We generate 2000 sequences on RS, CUT-1, and CUT-2 respectively for testing purposes. The performance of the packing algorithm is quantitated with space utilization (*space uti.*) and the total number of items packed in the bin (*# items*).

4.2 Ablation Study and Evaluation

Table 1 reports an ablation study about different adoptions. The feasibility mask $\mathbf{M}_n$ saves the efforts of exploring invalid actions during the training and guarantees the basic performance of the algorithm. The performance is impaired if the infeasibility loss E_{inf} is not contributed to the final loss since some exploration would be wasted for infeasible actions. E_{entropy} encourages the agent to find better solutions and benefits the final performance. The decision condition input (Figure 2 (right)) for each actor-head also facilitates more stable training and better performance. Figure 7 visualizes the effect of different algorithm settings.

$\mathbf{M}_n$	E_{inf}	E_{entropy}	*Cond.*	Space uti.	# items
✗	✗	✗	✗	39.5%	15.2
✗	✓	✓	✓	60.0%	23.2
✓	✗	✓	✓	67.9%	26.1
✓	✓	✗	✓	68.3%	26.4
✓	✓	✓	✗	68.4%	26.5
✓	✓	✓	✓	**71.3**%	**27.6**

Table 1 Effect of different adoptions on the RS dataset. *Cond.* means the decision condition input for each actor-head.

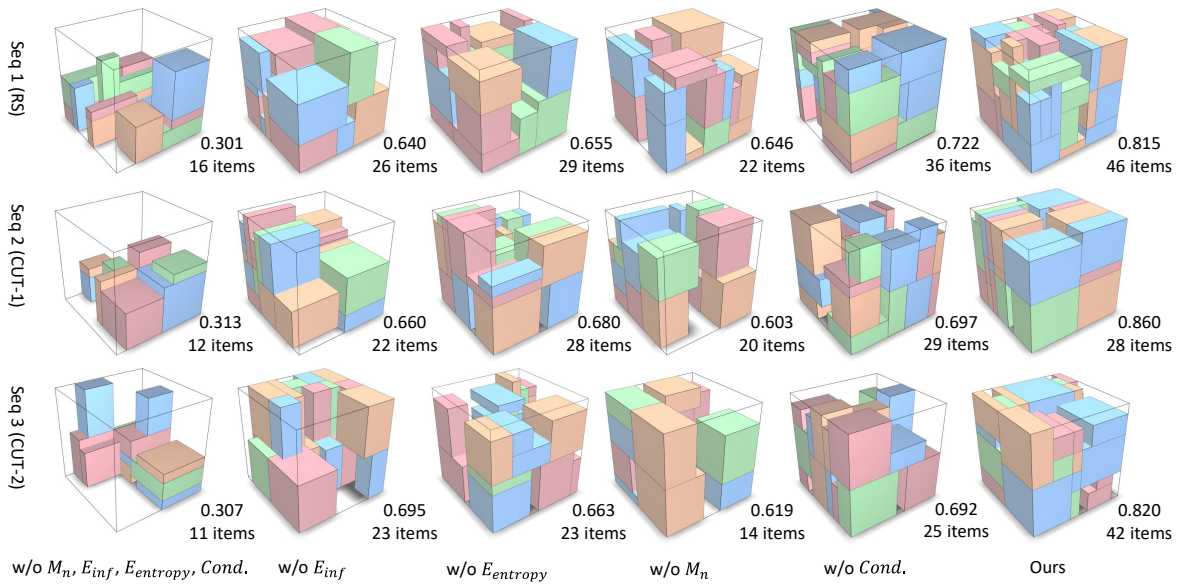

Figure 7 Visualization of the ablation study. The numbers beside each bin are *space uti.* and *# items*.

4.3 Stacking Tree based Stability Estimation

Next, we demonstrate that the proposed stacking tree based stability estimation is necessary for both efficiency and performance. Table 2 reports a quantitated comparison between the proposed method and four alternatives on the RS benchmark. The simulator reports the stability estimation given by Bullet [8] according to a numerical force analysis. Bullet is able to present precise stability estimation while it is time-consuming. Secondly, [50] introduced a heuristic-based approach to discriminate stability of n depends on whether it satisfies any of the following conditions: 1) over 60% of n's bottom area and all of its four bottom corners are supported by existing items; or 2) over 80% of n's bottom area and three out of four bottom corners are supported; or 3) over 95% of n's bottom area is supported. This approach can also promise 100% accuracy for stability estimation. However, the heuristic constraints are too strict for our network, the packing performance would be limited in this scenario.

To demonstrate that our stacking tree structure is necessary for stability estimation, we also evaluate the accuracy and performance if only the current packing item n's stability is checked based on our supported centroid approach. It will fail in some cases, which can only achieve 93.7% accuracy. The whole implementation of our stability estimation gives the best packing performance with a 99.9% estimation accuracy while OnlineBPH [17] can only achieve 29.1% space utilization with 11.7% estimation accuracy. Note that our high-speed stability estimation method is for the convenience of DRL training. In practice, we can perform a slower but more accurate stability estimation thread in parallel to avoid misestimation hazards which account for only 0.1%.

4.4 Action Space Decomposition

Benefit from the design of action space decomposition, our model is able to handle the input states with higher resolution. Table 3 reports the packing performance change as the resolution grows. Note that

	Space uti.	# items	stability	time
Bullet simulator [8]	− − −	− − −	− − −	$5.8 \times 10^{-}2$ s
Heuristic rule [50]	58.3%	22.8	100.0%	$1.8 \times 10^{-}4$ s
OnlineBPH [17]	29.1%	11.0	11.7%	− − −
Only current packing	66.3%	25.68	93.7%	$3.8 \times 10^{-}4$ s
Ours	**71.3**%	**27.6**	99.9%	$5.0 \times 10^{-}4$ s

Table 2 Quantitated comparison between the proposed method and four alternatives on RS benchmark. Our method can achieve the best packing performance with the high accuracy of stability estimation. Note that the counting stops when unstable placement occurs.

the small fluctuations of performance are normal since different sizes of convolution kernels are adopted to encode features with different resolutions. The packing performance of our method remains stable with different resolutions while the performance of [50] drops significantly. The results demonstrate the superiority of the action space decomposition when dealing with high resolution input states.

		10×10	30×30	50×50	100×100	200×200
Ours	space uti.	70.1%	71.7%	72.6%	71.3%	70.2%
	# items	27.1	27.7	28.1	27.6	27.1
[50]	space uti.	72.3%	72.4%	51.7%	− − %	− − %
	# items	27.9	27.9	20.6	−−	−−

Table 3 Performance comparison with different resolutions between method with decomposed action space (Ours) and method with uniform action space [50]. Note that [50] fails with 100×100 and 200×200 while our method can still maintain a good performance.

We compared our design of action space decomposition with other forms of action space: unified action space and continuous action space. Unified action space refers to the uniform encoding of all actions like $a_n = x_n + L \cdot y_n + L \cdot W \cdot o_n$ and only one actor is trained. Another alternative is to map a unified action space to a continuous domain in the range of [0,1]. As the agent predicts in this continuous domain, we project the result back to the discrete domain. We evaluate these three methods with two different resolutions, 10×10 and 100×100. The result is reported in Table 4.

The unified action space can achieve good performance with high computational efficiency when the state resolution is low. However, the performance drops significantly and the computing overhead has increased dramatically while the action space resolution is high. The continuous form of action space is computational friendly, but cannot achieve a good packing performance. Meanwhile, our action space decomposition method can not only maintain similar performance in different spatial dimensions but also maintain a reasonable computational overhead.

High-resolution action space provides more flexibility to practical packing. It is very common that the drift exists between the excepted position and actual placement due to robot manipulation error or item

		Uniform	Continuous	Ours
10×10	space uti.	72.3%	54.6%	70.2%
	# items	27.9	21.5	27.1
	GPU Memory	$948MB$	$942MB$	$944MB$
	Network update time	$2.0 \times 10^{-}2$ s	$1.4 \times 10^{-}2$ s	$2.7 \times 10^{-}2$ s
100×100	space uti.	− − %	51.2%	71.3%
	# items	−−	20.3	27.6
	GPU memory	$8302MB$	$1352MB$	$1356MB$
	Network update time	84.7 s	$1.6 \times 10^{-}2$ s	$5.1 \times 10^{-}2$ s

Table 4 Comparison between different methods. The action decomposition method has reached a good level both in performance and training overhead. Unified action space would fail while the resolution increasing to 100×100.

misdetection. Here we demonstrate that the high-resolution action space can tolerate item drift more via a toy example in Figure 8(a)(b). We also provide quantitative experimental results here. Assuming that the size of the pallet is 100cm × 100cm, we will impose a random non-zero drift from -5cm to 5cm to the placed item with a probability of 20%. If we pack items at a high resolution of 100 × 100, 1cm per cell, the packing utilization is 68.2% while the utilization is 71.3% if no drift is imposed. If we execute the same task at a low resolution of 10 × 10, which means 10cm per cell, the packing utilization drops to 63.1% from 70.1%.

If an item's dimension is not an integer with the resolution of the action space, an intuitive approach is filling in this non-unit dimension upwards, e.g., from 15cm to 20cm, which also leads to space waste shown in Figure 8(c). If the resolution cell is small, the agent can place items in a more compact way, shown in Figure 8(d). When the item set $\mathcal{I}$ contains non-unit dimension items, our method's performance at 100 × 100 is 70.9%, while 67.0% at 10 × 10.

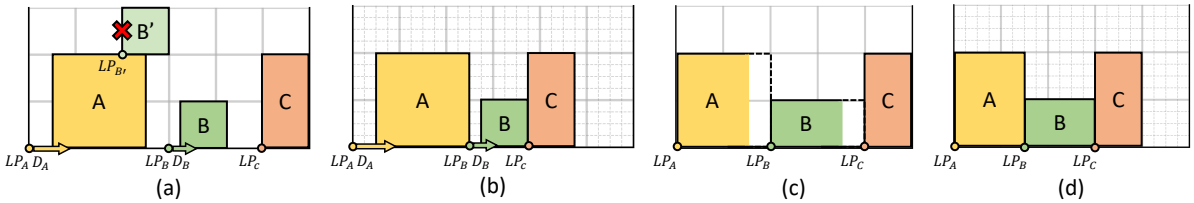

Figure 8 A, B, and C arrive in order. The grey grids represent the resolution of the action space, items can only be placed on the corners of the grids. (a): Packing in low-resolution action space, B cannot be placed at $LP_{B'}$ due to drift D_A and placing B at LP_B introduces additional space occupancy. The drift D_B will affect the placement of C as well. (b): High-resolution packing tolerates item drift and desires less space. (c): Filling up the non-unit dimension items A and B in low-resolution action space takes up five units while (d) packing at a high resolution takes less.

4.5 Scalability of BPP-k

Once the agent has mastered the capability of lookahead, it should better exploit the remaining space and deliver a more compact packing. The environment space increases exponentially as k value increases due to the factorial complexity, the agent should make its decision in a reasonable period. In Figure 9(a), we compared the time overhead and algorithm performance of the parallel MCTS compared to the serial approach. In Figure 9(b), we show that the performance of MCTS improves with the number of lookahead. Note that the root parallelization of MCTS not only improves the efficiency but also enables a larger search space, while the serial MCTS would be easily stuck in a local optimum with a large lookahead size.

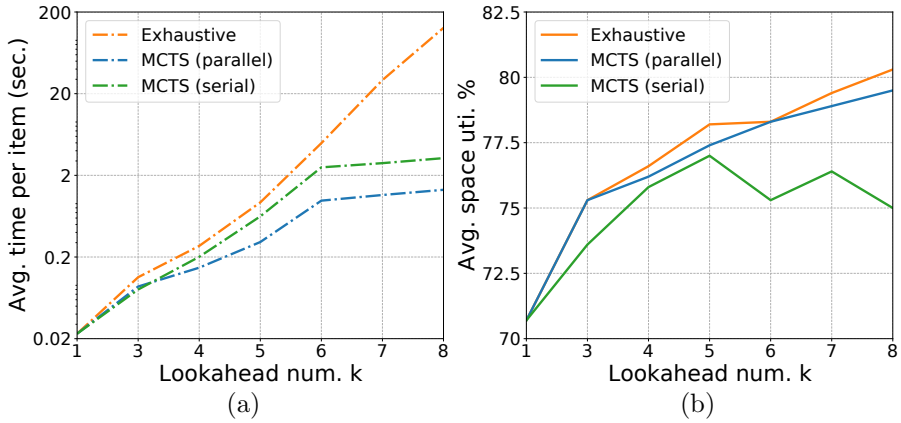

Figure 9 (a): Our MCTS implementation avoids the factorial computational cost of exhaustive permutation. The parallelization of MCTS also reduces the execution time by nearly half. (b) Our parallel MCTS achieves similar performance (avg. space utility) as the brute-force search over permutation tree. The performance of the serial MCTS drops when the lookahead size is larger than 5 due to the ever-increasing search space.

In the process of MCTS simulation, only legally sampled sequences would be adopted in the training. Sampling quality would directly impact search efficiency. We evaluate 6 different node sampling strategies

with $k = 10$ and the result is reported in Table 5. Here we can see, if we first sample the nodes that actually arrive late with strategy 5^x, the performance drops significantly since too many illegal sampled sequences are generated in this scenario. If we first sample the nodes that actually arrive early, MCTS begin to be able to utilize lookahead information and achieve better performance. $\mathcal{N}(0,1)$ is the best sampling strategy among them. Fixed sampling order performs worse since it can not explore the search space well.

	5^x	*random*	$1/x$	$\mathcal{N}(0,1)$	$(1/5)^x$	Fixed sampling order
space uti.	72.9%	79.1%	80.6%	82.5%	81.2%	77.9%
# items	27.7	31.1	31.6	32.3	31.6	30.0

Table 5 The node sample strategy will influence the performance of MCTS, $\mathcal{N}(0,1)$ is the best sampling strategy among them.

4.6 Comparison with Existing Methods

We compare to the state-of-the-art non-learning online bin packing method OnlineBPH [17] to demonstrate that our method can achieve the best performance on this problem. Note that [17] allows the agent to select the next best item from k lookahead ones (i.e., BPP-k with re-ordering). In Figure 10(a-c), we report the comparison under the setting of BPP-k with 100×100 state resolution on the three benchmarks. Our method can surpass OnlineBPH [17] significantly even re-ordering is allowed in this method while our method does not. Note that OnlineDRL [50] which adopts the unified action space would fail in the training with such a high state resolution.

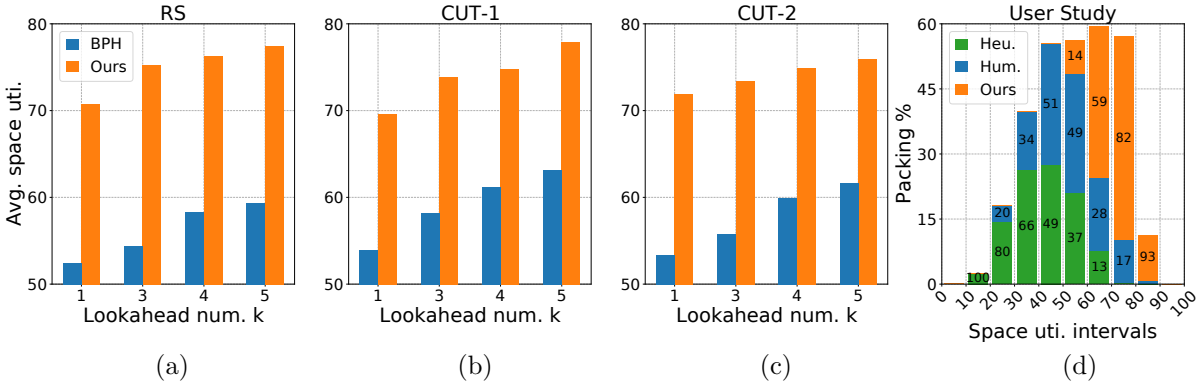

Figure 10 (a)-(c): Comparison with the online BPH method [17] on BPP-k with 100×100 state resolution. Note that BPH allows selecting any one of the lookahead k items while ours must place them in the order of arrival. (d): The distribution of space utilization using boundary rule (Heu.), human intelligence (Hum.), and our BPP-1 method (Ours). The more to the right of the distribution, the better the effect of the algorithm.

The most concerning issue should be the comparison between our algorithm and human intuition. To get the conclusion, we asked 50 human users to pack items manually with a Sokoban-like app proposed by [50]. The same sequence is collected and used to test AI (our method). The one with higher space utilization wins. 15 of the users are palletizing workers and the rest are CS-majored undergraduate/graduate students. No time limit is given. We conduct $2,104$ comparisons and the statistics are plotted in Figure 10 (d). Our method outperforms human players in general ($1,772$ AI wins vs. 289 human wins and 43 evens): it achieves 70.4% average space utilization while human players only have 56.3% (palletizing workers achieve 57.6% while CS-majored students achieve 55.2%).

4.7 Robot Implementation in Industry Environment

We tested our BPP-1 and BPP-k method in a practical industry environment, the implementation details can be found in Section 3.5. For BPP-1, the RGB-D camera recognizes the coming item by segment the captured depth map (Figure 11). If multiple items are observed by the camera at the same time, we sort their positions and select the item at the front of the queue to pack.

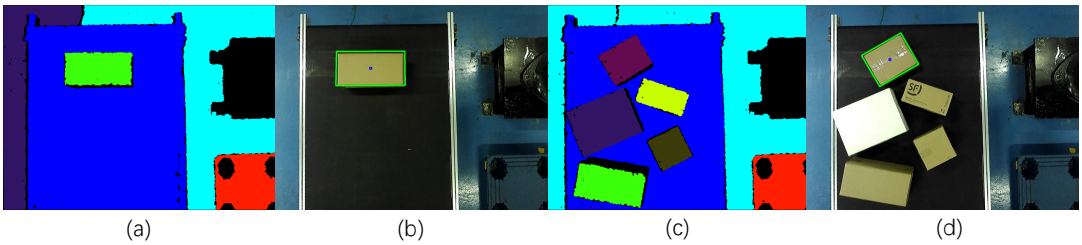

Figure 11 (a): Segmentation result under BPP-1 view. (b): The item n is labeled with an anchor, its dimension and location are recognized. (c): Segmentation result under BPP-k (k is a variable) view. (d) If multiple items are observed by the camera, we sort their positions and select the first one in the queue.

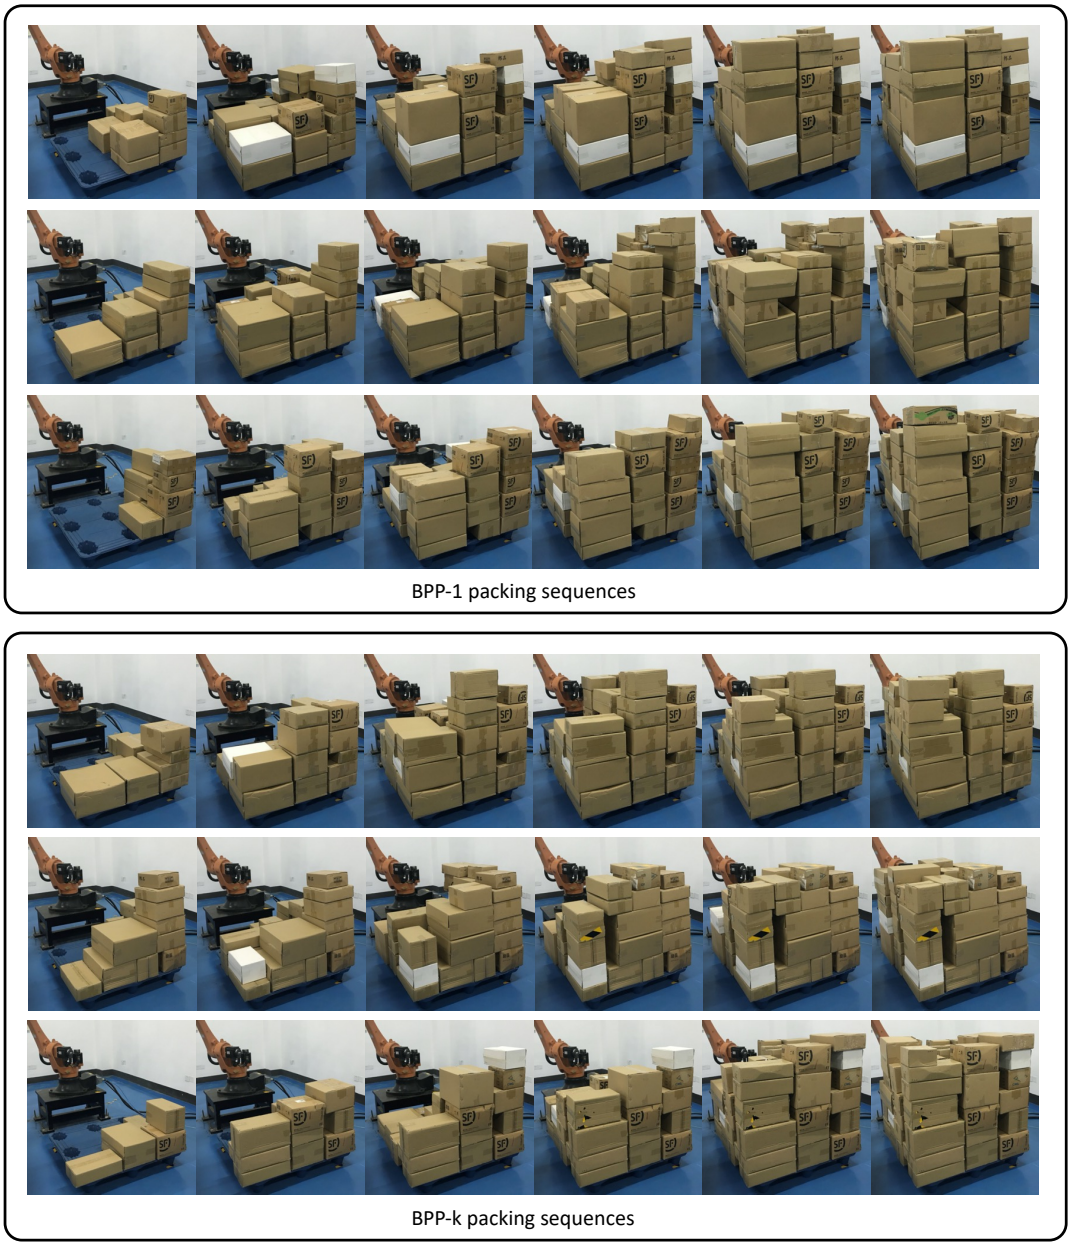

Figure 12 Visual examples are given by our robot implementation. The robot places items in a far-to-near order and reduces collisions with packed items.

Next, we evaluate our BPP-1 and BPP-k in this real demo as well in Table 6. Figure 12 shows some visual captures of the packing process given by our robot implementation. Note that, as shown in Figure 11 (d), the number of observed items is uncertain, k is changing while the packing. This is even more challenging than k is fixed. To fully test our BPP-k method, we choose Express®box standard as our test data since its box size is relatively small which ranges from 20cm to 50cm on each dimension. The input would include more numbers of observed items in this configuration. We also test the packing stability and robot arm collision in this experiment. The packing given by our method is 100% stable in all the 50 test sequences with BPP-1 and BPP-k. With the help of our design of collision-free reward, the robot places items in a far-to-near order and only 1 sequence encountered a collision between the robot arm and the packed item with BPP-1. If the collision-free reward is not included, the agent would place items uniformly on the XY plane of the bin and the number of collision sequences would significantly increase.

	space uti.	#items	stability	collision rate	
				ours	w/o collision-free reward
BPP-1	71.2%	49.0	100%	2%	35%
BPP-k	77.9%	53.5	100%	0%	30%

Table 6 Performance evaluation of BPP-1 and BPP-k in a practical industry environment. Our method can achieve good performance while keeping the stack stable and avoiding the collision.

5 Conclusion

We have provided a practical solution to online 3D-BPP with only partial sequence observation and deployed it on a real robot. To learn a feasible packing policy, we propose three new designs. First, we propose an online analysis of packing stability with a novel stacking tree. Second, we propose a decoupled packing policy learning for different dimensions of placement which enables high-resolution discretization and thus high packing precision. Third, we introduce a reward function that dictates the robot to place items in a far-to-near order and therefore simplifies the collision detection and movement planning of the robotic arm. As future work, we would like to investigate more on the problem of domain transfer of learned packing policies. In addition, we would like to investigate more relaxations of the problem such as stability estimation of items with irregular shapes or uneven mass distribution and studying combined solutions of palletizing and depalletizing.

References

1 Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 22–31. JMLR. org, 2017.

2 Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

3 Gabriel Barth-Maron, Matthew W Hoffman, David Budden, Will Dabney, Dan Horgan, Alistair Muldal, Nicolas Heess, and Timothy Lillicrap. Distributed distributional deterministic policy gradients. *arXiv preprint arXiv:1804.08617*, 2018.

4 Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.

5 Guillaume MJ-B Chaslot, Mark HM Winands, and H Jaap van Den Herik. Parallel monte-carlo tree search. In *International Conference on Computers and Games*, pages 60–71. Springer, 2008.

6 Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017.

7 Edward G Coffman, Michael R Garey, and David S Johnson. Approximation algorithms for bin-packing—an updated survey. In *Algorithm design for computer system design*, pages 49–106. Springer, 1984.

8 Erwin Coumans. Bullet physics simulation. In *ACM SIGGRAPH 2015 Courses*, page 1. 2015.

9 Teodor Gabriel Crainic, Guido Perboli, and Roberto Tadei. Extreme point-based heuristics for three-dimensional bin packing. *Informs Journal on computing*, 20(3):368–384, 2008.

10 Teodor Gabriel Crainic, Guido Perboli, and Roberto Tadei. TS2PACK: A two-level tabu search for the three-dimensional bin packing problem. *European Journal of Operational Research*, 195(3):744–760, 2009.

11 JL De Castro Silva, NY Soma, and N Maculan. A greedy search for the three-dimensional bin packing problem: the packing static stability case. *International Transactions in Operational Research*, 10(2):141–153, 2003.

12 Amit Dekel, Linus Harenstam-Nielsen, and Sergio Caccamo. Optimal least-squares solution to the hand-eye calibration problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13598–13606, 2020.

13 Kenny Erleben. Velocity-based shock propagation for multibody dynamics animation. *ACM Transactions on Graphics (TOG)*, 26(2):12–es, 2007.

14 Oluf Faroe, David Pisinger, and Martin Zachariasen. Guided local search for the three-dimensional bin-packing problem. *Informs journal on computing*, 15(3):267–283, 2003.

15 Chen Feng, Yuichi Taguchi, and Vineet R Kamat. Fast plane extraction in organized point clouds using agglomerative hierarchical clustering. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6218–6225. IEEE, 2014.

16 Paul C Gilmore and Ralph E Gomory. A linear programming approach to the cutting-stock problem. *Operations research*, 9(6):849–859, 1961.

17 Chi Trung Ha, Trung Thanh Nguyen, Lam Thu Bui, and Ran Wang. An online packing heuristic for the three-dimensional container loading problem in dynamic environments and the physical internet. In *European Conference on the Applications of Evolutionary Computation*, pages 140–155. Springer, 2017.

18 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

19 Donghui Han, Shu-wei Hsu, Ann McNamara, and John Keyser. Believability in simplifications of large scale physically based simulation. In *Proceedings of the ACM Symposium on Applied Perception*, pages 99–106, 2013.

20 Young-Dae Hong, Young-Joo Kim, and Ki-Baek Lee. Smart Pack: Online autonomous object-packing system using rgb-d sensor data. *Sensors*, 20(16):4448, 2020.

21 Shu-Wei Hsu and John Keyser. Automated constraint placement to maintain pile shape. *ACM Transactions on Graphics (TOG)*, 31(6):1–6, 2012.

22 Haoyuan Hu, Xiaodong Zhang, Xiaowei Yan, Longfei Wang, and Yinghui Xu. Solving a new 3D bin packing problem with deep reinforcement learning method. *arXiv preprint arXiv:1708.05930*, 2017.

23 Leonid V Kantorovich. Mathematical methods of organizing and planning production. *Management science*, 6(4):366–422, 1960.

24 Korhan Karabulut and Mustafa Murat İnceoğlu. A hybrid genetic algorithm for packing in 3D with deepest bottom left with fill method. In *International Conference on Advances in Information Systems*, pages 441–450. Springer, 2004.

25 Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

26 Bernhard Korte and Jens Vygen. Bin-packing. In *Kombinatorische Optimierung*, pages 499–516. Springer, 2012.

27 Alexandre Laterre, Yunguan Fu, Mohamed Khalil Jabri, Alain-Sam Cohen, David Kas, Karl Hajjar, Torbjorn S Dahl, Amine Kerkeni, and Karim Beguir. Ranked reward: Enabling self-play reinforcement learning for combinatorial optimization. *arXiv preprint arXiv:1807.01672*, 2018.

28 Xueping Li, Zhaoxia Zhao, and Kaike Zhang. A genetic algorithm for the three-dimensional bin packing problem with heterogeneous bins. In *IIE Annual Conference. Proceedings*, page 2039. Institute of Industrial and Systems Engineers (IISE), 2014.

29 Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

30 Andrea Lodi, Silvano Martello, and Daniele Vigo. Approximation algorithms for the oriented two-dimensional bin packing problem. *European Journal of Operational Research*, 112(1):158–166, 1999.

31 Silvano Martello, David Pisinger, and Daniele Vigo. The three-dimensional bin packing problem. *Operations research*, 48(2):256–267, 2000.

32 Silvano Martello, David Pisinger, Daniele Vigo, Edgar Den Boef, and Jan Korst. Algorithm 864: General and robot-packable variants of the three-dimensional bin packing problem. *ACM Transactions on Mathematical Software (TOMS)*, 33(1):7–es, 2007.

33 James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR, 2015.

34 Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.

35 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

36 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.

37 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

38 David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. 2014.

39 David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.

40 Shigeyuki Takahara and Sadaaki Miyamoto. An evolutionary approach for the multiple container loading problem. In *Fifth International Conference on Hybrid Intelligent Systems (HIS'05)*, pages 6–pp. IEEE, 2005.

41 Arash Tavakoli, Fabio Pardo, and Petar Kormushev. Action branching architectures for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

42 Kasper Kronborg Thomsen and Martin Kraus. Simulating small-scale object stacking using stack stability. In *23rd International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision WSCG 2015*, pages 5–8. Vaclav Skala-UNION Agency, 2015.

43 Eiji Uchibe and Kenji Doya. Constrained reinforcement learning from intrinsic and extrinsic rewards. In *2007 IEEE 6th International Conference on Development and Learning*, pages 163–168. IEEE, 2007.

44 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.

45 Hanqing Wang, Wei Liang, and Lap-Fai Yu. Scene mover: automatic move planning for scene arrangement by deep reinforcement learning. *ACM Trans. Graph.*, 39(6):233:1–233:15, 2020.

46 Ran Wang, Trung Thanh Nguyen, Shayan Kavakeb, Zaili Yang, and Changhe Li. Benchmarking dynamic three-dimensional bin packing problems using discrete-event simulation. In *European Conference on the Applications of Evolutionary Computation*, pages 266–279. Springer, 2016.

47 Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.

48 Yuhuai Wu, Elman Mansimov, Roger B Grosse, Shun Liao, and Jimmy Ba. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In *Advances in neural information processing systems*, pages 5279–5288, 2017.

49 Cong Zhang, Wen Song, Zhiguang Cao, Jie Zhang, Puay Siew Tan, and Chi Xu. Learning to dispatch for job shop scheduling via deep reinforcement learning. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

50 Hang Zhao, Qijin She, Chenyang Zhu, Yin Yang, and Kai Xu. Online 3D bin packing with constrained deep reinforcement learning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 741–749. AAAI Press, 2021.