

Title	An Efficient Deep Reinforcement Learning Model for Online 3D Bin Packing Combining Object Rearrangement and Stable Placement
Author(s)	Zhou, Peiwen; Gao, Ziyang; Li, Chenghao; Chong, Nak Young
Citation	2024 24th International Conference on Control, Automation and Systems (ICCAS): 964-969
Issue Date	2024-10-29
Type	Conference Paper
Text version	author
URL	http://hdl.handle.net/10119/19405
Rights	This is the author's version of the work. Copyright (C) ICROS. 2024 24th International Conference on Control, Automation and Systems (ICCAS 2024), 2024, pp. 964-969. DOI: 10.23919/ICCAS63016.2024.10773090. Personal use of this material is permitted. This material is posted here with permission of Institute of Control, Robotics and Systems (ICROS).
Description	2024 24th International Conference on Control, Automation and Systems (ICCAS), Jeju, Korea, October 29-November 1, 2024



An Efficient Deep Reinforcement Learning Model for Online 3D Bin Packing Combining Object Rearrangement and Stable Placement

Peiwen Zhou*, Ziyang Gao*, Chenghao Li, and Nak Young Chong

School of Information Science, Japan Advanced Institute of Science and Technology
1-1 Asahidai, Nomi, Ishikawa 923-1292, Japan
{s2310071, ziyang-g, chenghao.li, nakyoung}@jaist.ac.jp

Abstract: This paper presents an efficient deep reinforcement learning (DRL) framework for online 3D bin packing (3D-BPP). The 3D-BPP is an NP-hard problem significant in logistics, warehousing, and transportation, involving the optimal arrangement of objects inside a bin. Traditional heuristic algorithms often fail to address dynamic and physical constraints in real-time scenarios. We introduce a novel DRL framework that integrates a reliable physics heuristic algorithm and object rearrangement and stable placement. Our experiment show that the proposed framework achieves higher space utilization rates effectively minimizing the amount of wasted space with fewer training epochs.

Keywords: 3D Bin Packing, Object Rearrangement, Placement Stability, Deep Reinforcement Learning

1. INTRODUCTION

Robotic bin packing has many applications in the fields of logistics, warehousing, and transportation. The 3D Bin Packing Problem (3D-BPP), a well-known NP-hard problems [3], is referred to as an optimization problem of packing multiple objects into a bin(s), while satisfying the bin capacity constraint [10]. The 3D-BPP can be tackled offline or online depending on whether all objects can be accessible or not. In terms of offline bin packing task, this setting assumes the prior knowledge of all objects, usually, finding the optimal packing sequence and optimal placement are involved in this setting. Typically, meta-heuristic algorithms have been employed to determine the optimal order sequence in previous studies [6], thereafter, heuristic algorithms, such as DBLF proposed by Korha and Mustaf [8] or HM proposed by Wang and Kris [16], are leveraged to determine where to place the object into the bin.

Compared with offline bin packing, online bin packing is more challenging. Basically, the packing order is random, and the agent can only observe the upcoming objects (either single or multiple objects) as illustrated in Fig. 1. In this context, relying exclusively on heuristics results in a considerable decline in bin utilization [6]. Under these constraints, Yang *et al.* [18] employed unpacking-heuristics to improve the utilization. Nonetheless, this method raises the time cost, thereby diminishing the overall efficiency of the packing process.

Recent progress in DRL has shown promising results in various domains by enabling models to learn optimal policies through trial and error [12]. Compared with heuristic algorithms, DRL excels in addressing optimization problems effectiveness in complex environments. However, real-world physical law damages the training efficiency as learning the physics in complex environment takes many trial-and-error iterations, and the stable placement cannot be guaranteed. Zhao *et al.* [20] and




Fig. 1.: Online 3D-BPP, where the agent can only observe an upcoming object and pack it on-the-fly.

Yang *et al.* [18] leveraged neural network to predict the physical feasibility map, enabling the agent to learn feasible packing strategies. Although these methods have achieved promising results in 3D-BPP, object stability is not guaranteed. To address these challenges, we propose an efficient and effective DRL framework using a highly reliable physics heuristic algorithm for online 3D-BPP. The main contributions of this paper are as follows.

- We proposed a highly reliable physics heuristic algorithm that guarantees the stability of object placement in complex multi-stack environments, while retaining as many placement positions as possible.
- We incorporated an object rearrangement process into the proposed framework which allows the robot manipulator to change the orientation of the upcoming object. It is also an efficient action that directly enhances space utilization without requiring additional time costs.

2. RELATED WORK

2.1 Heuristics in Bin Packing Problem

The bin packing problem is a key challenge in combinatorial optimization, aiming to arrange multiple objects

This work was supported by JSPS KAKENHI Grant Number JP23K03756.

*These authors contributed equally to this work.

efficiently within the larger container. However, 3D-BPP become unsolvable within a reasonable time frame using exact algorithms [11] when involving a large number of objects. Over the years, various heuristic and meta-heuristic methods have been developed to address this problem [1][7][8][17]. Heuristic algorithms critically depend on the sequence of object placement, and current research often employs meta-heuristic techniques such as simulated annealing [19] and genetic algorithms [8].

Consequently, if complete information on all objects to be packed is unavailable, the effectiveness of heuristic algorithms drops significantly. Moreover, in real-world logistics warehouse, gathering detailed information about all objects can be challenging and time-consuming, reducing operational efficiency. Therefore, We propose using the object rearrangement method to change the orientation of objects in order to improve bin utilization, under the constraints of unchangeable order sequence.

2.2 DRL in 3D-BPP

DRL combines the decision-making capabilities of reinforcement learning with the powerful representation learning of deep neural networks. Furthermore, it can be adaptable to changing conditions and provide feasible solutions with highly efficient [12], where traditional methods may struggle to find efficient solutions. DRL has recently demonstrated strong performance across various robotics tasks [15][13], showcasing its ability to handle complex spatial and dynamic challenges effectively.

Thus applying DRL to the 3D-BPP could indeed be a highly efficient approach. For example, Zhao *et al.* [20] introduced a prediction-and-projection scheme where the agent first generates a physical stability mask for placement actions as an auxiliary task, then using this mask to adjust the action probabilities output by the actor during training. However, DRL models can suffer from instability and sensitivity to hyperparameters, making them difficult to tune and sometimes resulting in unpredictable performance. Moreover, most work focuses only on sample constraints, without considering real-world physical properties of objects, including the object CoM and its deviation in a complete stack. These factors can result in solutions that are impractical for real-world applications where physical stability and balance are essential.

Thus we propose the DRL framework integrated with a physics heuristics. This not only guarantees the stability of object placement but also enhances the training efficiency of the model, allowing for faster convergence.

2.3 Stability check in 3D-BPP

Stable stacking is a critical factor when designing an online 3D bin packing pipeline. Learning the rules of real-world physics is a very difficult process for DRL. This not only lengthens the training time for the model but also causes fluctuations in model convergence.

Therefore, for 3D-BPP, it is necessary to design a reliable and efficient physics heuristics for feasible action detection to quickly rule out incorrect actions in the current state. Zhao *et al.* [20] and Yang *et al.* [18] use the sim-




Fig. 2.: The main idea of convexHull-1. The left image depicts a sliding window that matches the size of the incoming object, along with portions of the scene objects contained within the sliding window. The right figure shows the zoom-in version of the content inside the sliding window. To determine the stability of the object, we calculate the largest convex hull of the highest points within the window. Next, we verify whether the center of the window lies within the convex hull. The object is deemed stable when positioned at the center of the sliding window if the convex hull includes the window’s center.

ilar scheme that combines the ratio of overlapping parts between the placed item and its contact items with a neural network model for prediction. But this is not a reliable method, since the model is a black box, there are always parts that are inexplicable and unpredictable. On the other hand, Wang and Kris [16] proposed a mathematical model that using a linear programming method solves for the torque balance and force balance of the object for all contact forces. Although this is a very reliable method, it is too complex for regular objects and usually takes a long time to evaluate all the candidates actions.

Thus we propose a new physics heuristic algorithm for rectilinear objects, which can guarantee the stability of object placement in an efficient and effective way, under real-world physical constraint.

3. METHOD

We describe our method in two parts. First, we present our stability check method, which is a highlight of our work. Second, we introduce a DRL framework that integrates physical heuristics and object rearrangement.

3.1 Stability Checking via Physics Heuristics

In our research, we assume that the object $i \in I$ for bin packing are rigid body and have a uniform mass distribution, so that the center of mass (CoM) is the geometric center of the object. But our method is not limited by the mass distribution, here just for simplify questions. For

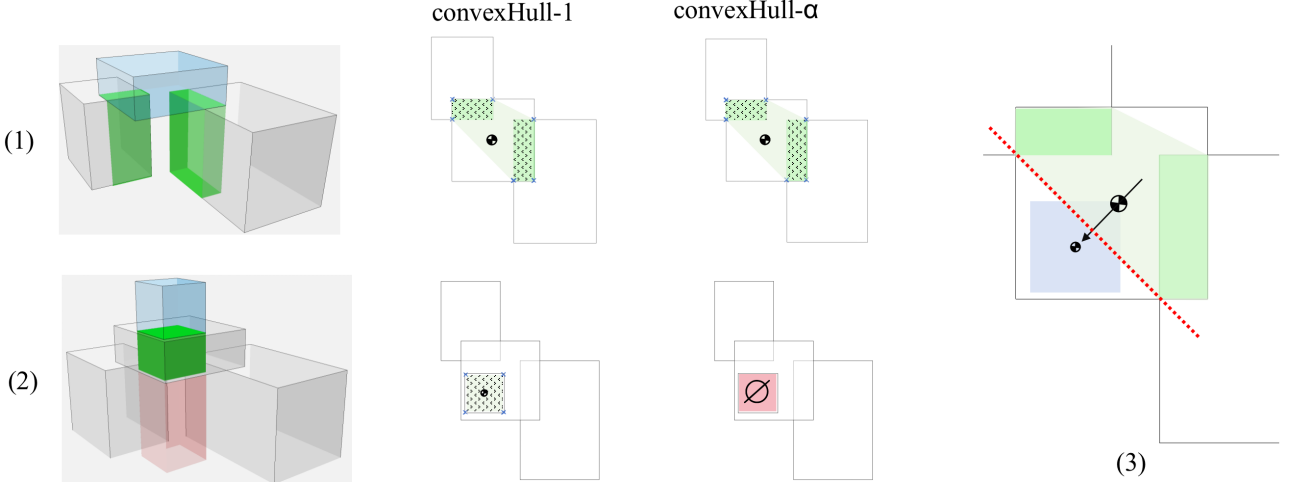


Fig. 3.: Multi-layer packing scenarios showcasing the difference between `convexHull-1` and `convexHull-α` algorithms for checking the stability of the placement. (1) Both `convexHull-1` and `convexHull-α` consider the arrangement to be stable. (2) Conversely, `convexHull-1` might incorrectly assess the stability if the incoming object is significantly heavier than the object in the middle layer, as detailed in (3).

uneven objects, we can use Gao *et al.* [4][5] to estimate as close as possible the CoM.

For the current state bin (W, D, H) , we generate a bottom-to-top depth heightmap with a resolution of 0.005, where each voxel represents a 5 mm^3 vertical column of 3D space in the agent’s workspace. The object to be placed i is defined by its dimensions (w_i, d_i, h_i) . We employ a $w_i \times d_i$ sliding window to traverse the height map to check the stability of each placement.

Based on physics’ principles, we introduce the `convexHull-1` method, as shown in Fig. 2. The upward support force, denoted as $p_f^i = \{p_1, p_2, \dots\}$ is defined as the set of highest points in the window, obtained by object i under currently placement. We utilize OpenCV [2] to calculate the largest convex hull formed by p_f . Then, we evaluate the placement stability by verifying if the center of the sliding window is within the convex hull or not. During our experiments, we observed that relying only on a single layer of the convex hull cannot ensure the stability of object placement. Fig. 3 shows an example using `convexHull-1` for stability check and fail.

To address the aforementioned issue, we introduce `convexHull-α`, for managing multiple stacks of objects in complex environments. Throughout the object packing procedure, we maintain an empty map with the same size as the action map. The main concept of `convexHull-α` is that the supporting force must be vertical and originate from the ground. Basically, for each position inside the sliding window, we check the number of wasted voxels along the z axis. We consider that only no wasted voxels can be the reliable support force, which corresponds to the empty map value is zero, denoted as p_f' . After each placement, we update the empty map outlines in Algorithm 2. Similarly, we use the new set of points p_f' to calculate the convex hull and determine whether the window’s CoM is within it or not. Fig. 3 illustrates an example of stability check using `convexHull-α`. Algorithm 1

outlines our algorithm in detail.

Algorithm 1 `convexHull-α`

Input: Bin depth heightmap B , empty map E , object i (w_i, d_i, h_i), bin height H

Output: Stable action map A

for $b_{x,y} \in B$ **do**

$window \leftarrow B[x: x + w_i][y: y + d_i]$

if $window$ out of boundary **or** $window.max + h_i > H$ **then**

$A_{x,y} \leftarrow 0$

continue

end if

if $window.max == 0$ **then**

$A_{x,y} \leftarrow 1$

continue

end if

Find p_f in $window$

Use E update p_f to p_f'

$center \leftarrow Point(\frac{w_i}{2}, \frac{d_i}{2})$

$hull \leftarrow convexHull(p_f')$

$inside \leftarrow Boolean center.within(hull)$

$A_{x,y} \leftarrow inside$

end for

return A

3.2 DRL for Bin Packing

3.2.1 Problem Formulation

Formally, online 3D bin packing involves packing a number of object $i \in I$, each with arbitrary dimensions (w_i, d_i, h_i) and cuboid shapes, into a bin of arbitrary dimensions (W, D, H) . The process is constrained by the visibility of only the immediately upcoming object could be packed into the bin. Once the bin is filled or can not pack upcoming object the process will stop.

To solve this task, we formulate it as a Markov De-




Fig. 4.: The pipeline of the DRL framework combined with object rearrangement and physics heuristics.

Algorithm 2 Empty map update

Input: Bin depth heightmap B , empty map E , action point (X, Y) , object $i = (w_i, d_i, h_i)$

Output: Updated empty map E

$window^B \leftarrow B[X : X + w_i, Y : Y + d_i]$

$window^E \leftarrow E[X : X + w_i, Y : Y + d_i]$

$h \leftarrow window^B.max$

for $e_{x,y}$ **in** $window^E$ **do**

$window_{x,y}^E \leftarrow e_{x,y} + h - window_{x,y}^B$

end for

$E[X : X + w_i, Y : Y + d_i] \leftarrow window^E$

return E

cision Processes (MDPs), which can be denoted as a tuple $M = \langle S, A, P, R, \gamma \rangle$. Specifically, we employ two agents with policies π_o and π_p to independently predict placement orientation and position.

The whole process is described as follow: At the time step t , the agent observes the environment and takes a state representation, denoted as s_t . Then the agent π_o predicts the action o and pass to agent π_p to predict action p . Execute the action tuple (o_t, p_t) , causing the environment to transition to s_{t+1} , then immediately obtains a reward r_t . The process aims to achieve the maximal cumulative rewards with discount γ , as shown in Eq. (1) and (2), by jointly optimizing two policies.

$$J_{\pi_o^*} = \max E_{\pi_o, s_t, o_t = \pi_o(s_t)} \left[\sum_t \gamma^t r_t \right] \quad (1)$$

$$J_{\pi_p^*} = \max E_{\pi_p, s_t, p_t = \pi_p(s_t | o_t)} \left[\sum_t \gamma^t r_t \right] \quad (2)$$

3.2.2 State Definition

We define state s_t as the configuration of the bin along with the object that is about to be packed. Use the depth image of the bin to generate a bottom-to-top $W \times D$ depth heightmap dh_t [20]. Following the work conducted by Yang *et al.* [18], given the object i with dimensions

(w_i, d_i, h_i) , we create a three channel map with the dimension $W \times D \times 3$. Each channel corresponds to one of the object's dimensions and is fully populated with the respective dimension values. Then combine them as $(dh_i, w_i^*, d_i^*, h_i^*)$ to represent the State.




Fig. 5.: Six possible orientations of the packing object.

3.2.3 Action Definition

In this work, we propose to arrange object orientation in order to achieve better placement. Therefore, the action is defined as the conjunction of object rearrangement and placement, which is represented by (o, p) , where o represents the target object orientation and p represents a specific position on top layer of the bin. To simplify the packing procedure, both o and p are discretized.

As illustrated in Fig. 5, there are six different orientations. The number of positions for possible placement is the same as the number of pixels inside the heightmap. Given (o, p) , the agent firstly uses object rearrangement operation to achieve the object orientation o , and then place the object to the position p .

3.2.4 Reward Function

Following the idea mentioned in [20], at the time step t , the immediate reward is the weighted subtraction of increased utilization space and the wasted space given by Eq. (3) to (5). Please note that the wasted space can be calculated efficiently by comparing the summation of the empty map before and after the placement. In addition, both α and β are set to be one in our experiment.

$$R^i = \alpha \cdot r_v^i - \beta \cdot r_{waste}^i \quad (3)$$

$$r_v^i = \frac{w_i \cdot d_i \cdot h_i}{W \cdot D \cdot H} \quad (4)$$

$$r_{waste}^i = \frac{v_{waste}^i}{W \cdot D \cdot H} \quad (5)$$

3.2.5 Physics Heuristics DRL Framework

Distinct from other works [18], we proposed a two-agents DRL framework integrated with physics heuristics as shown in Fig. 4. Based on Proximal Policy Optimization (PPO) [14], we develop two actor networks: *Actor_o* dedicated to predicting the object’s orientation and *Actor_p* to determining the packing position. Both actor networks takes input as the 4-channels maps, the output of *Actor_o* is a six-dimensional vector where each element dedicates one specific object orientation, the output of *Actor_p* is the action map for placement with the same size as the heightmap.

The training pipeline is as follows: Given the object i and configuration of bin dh_i , firstly, the Phy-Heu module generates stable action maps for all potential object orientations. Using these stable action maps, we construct an orientation mask to exclude orientations that do not allow for any feasible stable placement. Meanwhile, *Actor_o* will predict the probability distribution of the object orientations. Using the orientation mask and the predicted distribution of orientations, the orientation is sampled. Next, based on the sampled orientation, the agent *Actor_p* takes the dh_i and shuffled (w_i^*, d_i^*, h_i^*) to predict the placement score map. Lastly, we sample the action P_i from the intersected map of the corresponding stable action map and the predicted action score map to ensure the placement stability.

4. EXPERIMENT AND RESULT

Our experiments were performed with CoppeliaSim using Newton physical engine. The experiments include: (1) the validation of the physical heuristic algorithms; (2) the training and testing of the DRL framework.

4.1 Physics Heuristics Validation

We compare the physics heuristic with algorithms convexHull-1 and convexHull- α on CoppeliaSim. The bin dimensions $W = D = H = 0.6m$. Objects are randomly generated with dimensions $w_i, d_i, h_i \in [0.03, 0.3]m$. In this experiment, based on the stable action map computed by convexHull-1 or convexHull- α , a random position considered to be stable for placement is selected at each time step. The stability of the bin objects are checked after each placement. The runtime of the two algorithms and the number of un-stable placement are reported in Table 1. Based on Table 1, We find that convexHull- α significantly surpasses convexHull-1 *w.r.t.* the accuracy of stability check. There was only one instance where convexHull- α incorrectly assessed the stability. We suspect this is due to the stable issue

of the physical engine. In addition, convexHull-1 and convexHull- α have similar runtime which indicate the efficiency of convexHull- α .

Table 1.: Comparison of physics heuristics.

	convexHull-1	convexHull- α
Object number	3000	3000
Fall number	153	1
Time cost(s)	4203.3	4452.6
Per cost(s)	1.40 - 1.00	1.48 - 1.00
Fall rate	5.1%	0.03%

4.2 DRL framework result

RS [20] bin packing dataset is leveraged to train and test the proposed DRL framework. To evaluate the effectiveness of our proposed method, the result reported in Zhao *et al.* [20] is our baseline as we share the same setting. Consistent with previous studies, we employed space utilization (Uti.) as the metric to evaluate the bin packing policy, where a higher value indicates better performance. We test on the dataset RS, CUT-1, and CUT-2 [20], which is summarized in Table 2.

Table 2.: Comparison of packing performance.

	RS	CUT-1	CUT-2	epoch
Ours	61.2%	63.3%	62.5%	18.6k
[20]	50.5%	60.8%	60.9%	100k

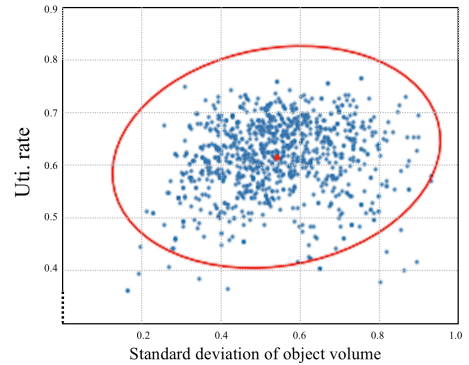


Fig. 6.: Space utilization of our model independent of the standard deviation in object volume.

The results show that our method achieves higher Uti. with fewer training epochs. Additionally, we analyzed the RS test results, comparing each test’s Uti with the standard deviation of the object volumes in the object sequence. Specifically, larger standard deviation indicates greater volume difference among the objects. As shown in Fig. 6, we found that the our model trained on the RS dataset is not affected by the differences in object volume within the sequence.

5. CONCLUSION

In this paper, we proposed an efficient DRL framework for tackling the online 3D-BPP, integrating a re-

liable physics heuristic algorithm and object rearrangement technique to enhance the stability and utilization of object placements within bins. Our experiments demonstrated that the proposed method achieves higher utilization rates with fewer training epochs compared to the baseline. The integration of the physics heuristic ensures the stability of object placements, significantly reducing the occurrence of object falls during both training and testing phases. The object rearrangement technique further improves the utilization of bin space without additional time costs, making our framework highly efficient for on-the-fly applications.

In the future, we aim to make our physics heuristic algorithm more accurate by precisely predicting each stable placement position, and to improve the training efficiency of the DRL model. Additionally, we will incorporate the method proposed by Gao *et al* [4][5] and Li *et al* [9] to grasp and pack irregular objects in the real world and attempt to propose a strategy for packing unknown and uneven objects in complex real-world environments.

REFERENCES

- [1] Brenda S Baker, Edward G Coffman Jr, and Ronald L Rivest. "Orthogonal packings in two dimensions". In: *SIAM Journal on computing* 9.4 (1980), pp. 846–855.
- [2] G. Bradski. "The OpenCV Library". In: *Dr. Dobbs's Journal of Software Tools* (2000).
- [3] Robert J Fowler, Michael S Paterson, and Steven L Tanimoto. "Optimal packing and covering in the plane are NP-complete". In: *Information Processing Letters* 12.3 (1981), pp. 133–137.
- [4] Ziyang Gao, Armagan Elibol, and Nak Young Chong. "On the generality and application of Mason's voting theorem to center of mass estimation for pure translational motion". In: *IEEE Transactions on Robotics* 40 (2024), pp. 2656–2671.
- [5] Ziyang Gao, Armagan Elibol, and Nak Young Chong. "Zero moment two edge pushing of novel objects with center of mass estimation". In: *IEEE Transactions on Automation Science and Engineering* 20.3 (2023), pp. 1487–1499.
- [6] EBCH Hopper and Brian CH Turton. "An empirical investigation of meta-heuristic and heuristic algorithms for a 2D packing problem". In: *European Journal of Operational Research* 128.1 (2001), pp. 34–57.
- [7] David S. Johnson, Alan Demers, Jeffrey D. Ullman, Michael R Garey, and Ronald L. Graham. "Worst-case performance bounds for simple one-dimensional packing algorithms". In: *SIAM Journal on computing* 3.4 (1974), pp. 299–325.
- [8] 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*. Springer. 2004, pp. 441–450.
- [9] Chenghao Li, Peiwen Zhou, and Nak Young Chong. "Safety-optimized Strategy for Grasp Detection in High-clutter Scenarios". In: *2024 21st International Conference on Ubiquitous Robots (UR)*. IEEE. 2024, pp. 192–197.
- [10] Silvano Martello, David Pisinger, and Daniele Vigo. "The three-dimensional bin packing problem". In: *Operations Research* 48.2 (2000), pp. 256–267.
- [11] Silvano Martello and Daniele Vigo. "Exact solution of the two-dimensional finite bin packing problem". In: *Management Science* 44.3 (1998), pp. 388–399.
- [12] Karl Pertsch, Youngwoon Lee, Yue Wu, and Joseph J Lim. "Guided reinforcement learning with learned skills". In: *arXiv preprint arXiv:2107.10253* (2021).
- [13] Zhenhuan Rao, Yuechen Wu, Zifei Yang, Wei Zhang, Shijian Lu, Weizhi Lu, and ZhengJun Zha. "Visual navigation with multiple goals based on deep reinforcement learning". In: *IEEE Transactions on Neural Networks and Learning Systems* 32.12 (2021), pp. 5445–5455.
- [14] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. "Proximal policy optimization algorithms". In: *arXiv preprint arXiv:1707.06347* (2017).
- [15] Shuran Song, Andy Zeng, Johnny Lee, and Thomas Funkhouser. "Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations". In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4978–4985.
- [16] Fan Wang and Kris Hauser. "Dense robotic packing of irregular and novel 3D objects". In: *IEEE Transactions on Robotics* 38.2 (2021), pp. 1160–1173.
- [17] Lei Wang, Songshan Guo, Shi Chen, Wenbin Zhu, and Andrew Lim. "Two natural heuristics for 3D packing with practical loading constraints". In: *PRICAI 2010: Trends in Artificial Intelligence: 11th Pacific Rim International Conference on Artificial Intelligence, Proceedings 11*. Springer. 2010, pp. 256–267.
- [18] Shuo Yang, Shuai Song, Shilei Chu, Ran Song, Jiyu Cheng, Yibin Li, and Wei Zhang. "Heuristics integrated deep reinforcement learning for online 3d bin packing". In: *IEEE Transactions on Automation Science and Engineering* (2023).
- [19] Defu Zhang and Wenqi Huang. "A simulated annealing algorithm for the circles packing problem". In: *International Conference on Computational Science*. Springer. 2004, pp. 206–214.
- [20] Hang Zhao, Qijin She, Chenyang Zhu, Yin Yang, and Kai Xu. "Online 3D bin packing with constrained deep reinforcement learning". In: *AAAI Conference on Artificial Intelligence*. Vol. 35. 1. 2021, pp. 741–749.