

Planning for Tabletop Object Rearrangement

Jiaming Hu

Jan Szczekulski

Sudhansh Peddabomma

Henrik I. Christensen

Abstract—Finding an high-quality solution for the tabletop object rearrangement planning is a challenging problem. Compared to determining a goal arrangement [1], rearrangement planning is challenging due to the dependencies between objects and the buffer capacity available to hold objects. Although [3] has proposed an A* based searching strategy with lazy evaluation for the high-quality solution, it is not scalable, with the success rate decreasing as the number of objects increases. To overcome this limitation, we propose an enhanced A*-based algorithm that improves state representation and employs incremental goal attempts with lazy evaluation at each iteration. This approach aims to enhance scalability while maintaining solution quality. Our evaluation demonstrates that our algorithm can provide superior solutions compared to [3], in a shorter time, for both stationary and mobile robots.

I. INTRODUCTION

Tabletop Object Rearrangement with Overhand Grasps (TORO) in bounded workspaces is a challenging manipulation task in robotics, with applications ranging from automated assembly lines to household assistance [2]. While many approaches can generate feasible rearrangement solutions [4], optimizing these solutions to minimize travel costs remains a challenging problem. This complexity arises particularly from the need to consider dependencies between objects to avoid collisions. Consequently, developing an algorithm that can efficiently search for high-quality plans with minimal travel cost is essential. This work aims to explore such an algorithm.

To find high-quality rearrangement plans with minimal cost, the Object Rearrangement with Lazy A* (ORLA*) algorithm [3] was introduced as a state-of-the-art method. ORLA* extends the A* algorithm for tabletop rearrangement planning by treating object arrangements as states and using pick-and-place actions to transition between them. To reduce computation times, ORLA* employs lazy evaluation, which avoids unnecessary buffer allocations and enhances efficiency. ORLA* still faces a challenge in terms of scalability. Like A*, it only terminates the search upon finding a solution or timing out, which implies exponential increases in computation time as the number of objects grows. This scalability issue makes it difficult for ORLA* to provide high-quality solutions in reasonable time for complex scenes.

This paper presents a revised A*-based algorithm, Scalable Tabletop Rearrangement A*-based Planning (STRAP), designed to address the scalability concern. Unlike ORLA [3], which performs lazy evaluation during the expansion of the search tree, STRAP conducts lazy evaluation after each exploration step. This enables rapid searches for feasible rearrangement plans from the current state to the goal in



Fig. 1. An example of mobile robot tabletop rearrangement, where black arrows point to the desired rearrangement goals.

each iteration. Consequently, while ORLA* may take longer to return a solution, STRAP quickly provides a high-quality solution. Additionally, STRAP incorporates the robot status to represent the system state, allowing exploration of more configurations and ultimately leading to improved planning outcomes compared to ORLA* [3]. The evaluation section presents experiments demonstrating STRAP's advantages for both stationary and mobile robots.

The paper is organized with related work in Section II, a problem statement in Section III, the improved algorithm in Section IV. Then, we evaluate the algorithm with both stationary and mobile manipulators in Section V.

II. RELATED WORKS

A. Object Rearrangement

Rearranging objects is a well-studied task and motion planning problem in robotics, including rearrangement planning for mobile robots [17], [18], [20] and planning in constrained spaces [19]. By considering the positions of the robot and objects, a solution is sought to ensure feasible object movements. Thus, these studies primarily focus on object accessibility from the robot's perspective.

B. Tabletop Rearrangement

Furthermore, tabletop rearrangement is a common rearrangement planning problem in robotics manipulation. Compared to other object rearrangement planning problem, the primary challenge arises from the dependencies between objects caused by potential collisions. To address this, a dependency graph is typically used, as described in [7], [8]. Utilizing this graph transforms the problem into the feedback vertex set problem and the traveling salesperson problem, both of which are NP-hard, as noted by [9], and causes the scalability problem. By incorporating an external buffer, [9] was able to search for the minimal cost solution to complete

the rearrangement task efficiently. However, it is common for the external buffer to be limited in size. To address this constraint, Running Buffer [11] introduced the concept of a running buffer, which restricts the size of the external buffer while still enabling the completion of the rearrangement task efficiently.

C. High-Quality Rearrangement planning

Nevertheless, in many cases, external buffers are not available for single manipulators. Allocation of an internal buffer that ensures optimality is a computationally intensive task. To address this challenge, [14] proposed the Neural Rearrangement Planner, which employs deep learning techniques to solve the problem recursively. However, this approach still does not guarantee a high success rate, even when dealing with a small number of objects. On the other hand, pre-identifying potential buffer candidates [12], [13] can be a viable solution. However, this approach suffers from scalability issues in dense environments. Consequently, TRLB [4] aims to solve the rearrangement task using only internal buffers, even when dealing with many objects. Despite its advantages, this method does not consider the travel cost. To further enhance this approach, Monte-Carlo Tree Search (MCTS) based methods [10], [16] have been proposed; However, these methods are specifically designed for stationary robots and require an exceptionally long time to achieve high-quality solutions. Subsequently, ORLA* [3] incorporates lazy buffer allocation within the A* search algorithm and prioritizes buffer poses based on various cost function optimizations for both stationary and mobile robots. Besides that, ORLA* [3] has demonstrated superior performance in terms of plan quality compared to MCTS-based methods.

To accomplish this, ORLA* classifies the states explored as deterministic (DS) and non-deterministic states (NDS). As the names suggest, deterministic states have well-defined object positions. These states have the objects either in the start states or the goal states. On the other hand, non-deterministic states occur as a result of an object being placed in a buffer position. Due to lazy buffer allocation, the buffer positions remain undefined until the exploration algorithm reaches a deterministic state. Upon reaching a deterministic state, the buffer positions are calculated by backtracking the actions. This way of calculation saves time by calculating buffer positions only when needed. However, this design requires a state representation that considers only object arrangements, excluding robot locations. As a result, the state representation is incomplete, leading to suboptimal solutions. Furthermore, although ORLA* incorporates actions to stack objects on top of one another using a neural network to predict stability, our focus is solely on the rearrangement planning aspect. We exclude stacking actions to avoid unnecessary complexity.

III. PROBLEM STATEMENT

Given a 2D table workspace with n objects, an arrangement A is defined as $\{p_1, p_2, \dots, p_n\}$, where each p_i represents the pose of an object i . An arrangement is considered

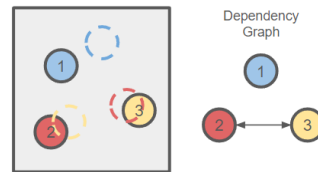


Fig. 2. Left: A rearrangement task involving a table as the boundary box. The colored circles represent objects in their initial positions, while the dotted circles show their target positions. Right: Its dependency graph.

valid only if no objects collide with each other or with the boundaries of the workspace. Unlike previous study [3], we assume that objects are placed directly on the table without stacking. A rearrangement plan consists of a sequence of pick-and-place actions $[a_1, a_2, \dots]$ that transition objects from one location to another. Each action a specifies which object to move, where to pick it up, and where to place it. An action is valid only if it does not cause any collisions upon placing the object.

Similar to [3], the quality of a rearrangement plan is evaluated based on the total travel cost required to move all objects from the initial arrangement, A_{init} , to the goal arrangement, A_{goal} . A plan with a lower total cost is considered to be of higher quality, indicating a more efficient rearrangement process. The travel cost is defined differently for stationary and mobile robots. For a stationary robot, the travel cost is the Euclidean distance traveled by the end-effector (EE). For a mobile robot, the travel cost is the Euclidean distance traveled by the mobile base (MB) around the table. Additionally, akin to [3], the mobile robot always moves to the closest point around the table to the pick or place location before performing the manipulation.

IV. METHODOLOGY

This section describes our new A*-based algorithm, named STRAP, aiming to overcome the scalability issue of ORLA* in TORO problem for both stationary and mobile robots.

In the A* algorithm, incomplete state representation can result in incorrect or suboptimal solutions. Therefore, unlike ORLA*, STRAP incorporates the robot's pose into the state representation to ensure higher completeness. That is, given n objects, a state s is defined as $s = \{p_r, p_1, p_2, \dots\}$ where p_r is the robot's pose, while each p_i is the pose of object i . On the other hand, each pick-and-place action a is defined as $a = (k, p_{pick}, p_{place})$ where k is the object, p_{pick} is the pick location, and p_{place} is the place location.

Given a state, the action decision policy involves selecting which object that is not in its goal position to be pick-and-placed. Similar to [3], [4], the placement location of an action is determined based on the object's condition as following:

- 1) If the object's goal position is available, then move it to its goal position.
- 2) If the object's goal position is not available, then move it to a valid buffer location different from where it was.

Starting from the initial state s_{start} , defined by both the initial arrangement and the robot's starting position, STRAP

explores one state at a time, adding all subsequent states derived from the explored state to the priority queue, often called open list. This queue is sorted based on the g and h values, where $g(s)$ is the cost from the initial state to s , and $h(s)$ is the estimated cost from s to the goal. In the next iteration, the priority queue pops the state with the lowest $f(s) = g(s) + h(s)$ and explores it until the explored state contains the goal arrangement or timeout.

To address scalability issues, STRAP is designed to deliver a solution within a short time frame. At the end of each iteration, STRAP performs a rapid goal-attempting procedure to identify a feasible, though not necessarily optimal, rearrangement plan from the current explored state to the goal arrangement. The best rearrangement plan, which minimizes the cost from the start state through the explored state to the goal state, is recorded. When a timeout occurs or the explored state matches the goal arrangement, the best plan available at that point is returned.

The following section covers state cost estimation, search tree exploration, goal attempting, and plan refinement.

A. Cost Estimation

To use the A* algorithm for the TORO problem, both $g(s)$ and $h(s)$ cost estimations must be precisely defined for each state s . Given a state s , $g(s)$ is the sum of the travel cost and manipulation cost used from s_{start} to s . Specifically, given a pick-and-place action sequence $[a_1, a_2, \dots, a_m]$ used to rearrange from s_{start} to s , we have

$$g(s) = MC \cdot m + TC(p_{robot}, p_{pick}^{a_1}) + TC(p_{pick}^{a_1}, p_{place}^{a_1}) + \sum_{i=2}^m [TC(p_{place}^{a_{i-1}}, p_{pick}^{a_i}) + TC(p_{pick}^{a_i}, p_{place}^{a_i})] \quad (\text{A.1})$$

In Eq. A.1, p_{robot} denotes the initial robot location, p_{pick}^a and p_{place}^a are pick-and-place location of action a respectively, m is the number of actions, and TC refers to the travel cost between two points. The MC represents the manipulation cost defined by user for each pick-and-place action. For instance, MC could account for the cost of opening and closing the gripper.

Furthermore, as shown in Eq. A.2, the $h(s)$ is the sum of manipulation cost times the number of objects not in goal, and the total pick-and-place travel cost as the lower bound cost. That is, given a state s with a set of objects $[o_1, o_2, \dots, o_l]$ which are not in their goal location,

$$h(s) = MC \cdot l + \sum_{i=1}^l TC(p_i, p_i^g) \quad (\text{A.2})$$

where p_i is the current location of o_i , while p_i^g is the goal location of o_i , and l is the number of missing placed objects.

Additionally, the function of travel cost is calculated differently for stationary and mobile robots, as described in [3] and illustrated in Fig. 3. Given two locations p_1 and p_2 , for a stationary robot, $TC(p_1, p_2)$ is the Euclidean travel distance of the end-effector in the x-y plane. On the other

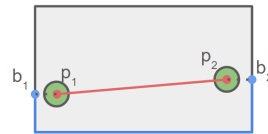


Fig. 3. Different definitions of travel cost between p_1 and p_2 . For stationary robots, the TC is the red line length. For a mobile robot, given b_1 and b_2 which are the closest points to p_1 and p_2 on the table side, the TC is the traveling distance moving along the table side (blue).

hand, for a mobile robot, $TC(p_1, p_2)$ is the Euclidean travel distance to move along the table side from b_1 to b_2 , the closest points on the table side to p_1 and p_2 respectively.

B. Exploration

Following the A* principle, each iteration explores the state s with the lowest f value from the priority queue. First, STRAP examines all states derived from s based on the decision policy, calculating their corresponding f values, then adds them to the priority queue. If a duplicate state is found in the queue, only the version with the lower f value is retained. Additionally, the explored state s is incorporated into the search tree as closed list in A*.

When exploring new states from s , each object can lead to one or more possible states. If the object can be moved to its goal location, the new state will reflect the arrangement with the robot location after the robot picks and places the object at its goal. If its goal location is occupied, we sample multiple buffer locations on the table for it. If a buffer location is successfully allocated, the new state will reflect the arrangement with the robot location after the object is placed in the buffer. The number of sampling attempts is based on the number of objects in our implementation. For instance, in a state with three objects, where one is already at its goal location, another can be moved directly to its goal, and the third's goal location is occupied, the outcomes are as follows: The object at its goal location does not lead to any new states. The object that can be moved directly to its goal location leads to one new state. For the object that cannot be moved to its goal location, three random positions are sampled on the table. If two of these sampled positions are collision-free with other objects in the current state and suitable as buffer locations, this object results in two new states. Thus, this scenario produces a total of three subsequent states from the given state.

However, this approach can result in redundant actions, such as moving the same object twice in succession or moving an object from one buffer location to another when it could have been placed there directly in the first instance. To address this, two additional rules are implemented during exploration. First, during exploring from a given state to others, STRAP avoids exploring actions that involve moving the same object that was used to reach the given state. Second, to avoid moving an object from one buffer to another where it could have been placed directly initially, during sampling buffer location for an object, STRAP does randomly sampling on table if that object is never moved at

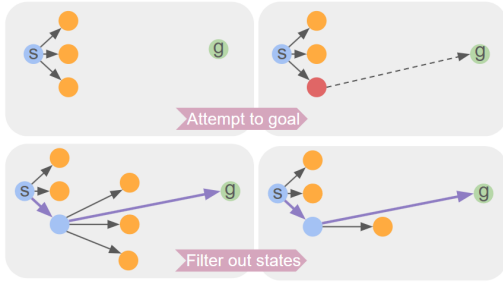


Fig. 4. Exploration process. The orange circles represent the states that are ready to be explored as in priority queue, whereas the blue circles indicate the states that have already been explored, and the green circle with g is the state whose arrangement is the goal arrangement. In the upper section, during the exploration phase, the node with the minimum f value is selected from the priority queue. The process of goal attempting involves searching for a feasible rearrangement solution from an explored state to the goal arrangement. In the lower section, if a plan exists that connects the start state to an explored state with a feasible solution to the goal, shown as purple path, and this plan is currently the best solution, its cost will be calculated. Subsequently, nodes in the priority queue with f values exceeding this cost will be filtered out.

all. If that object has been moved and is currently in a buffer location, STRAP only samples new buffer locations from areas where that object could not have been placed before being moved to its current buffer due to collision. That is, in a given state $s_{current}$, when sampling a new buffer location for an object O_{in_buffer} already in a buffer, STRAP identifies the ancestor state $s_{ancestor}$ of $s_{current}$ in the search tree where O_{in_buffer} was last moved to its current buffer. The new buffer location is then sampled in a way that O_{in_buffer} would collide with other objects in the arrangement of $s_{ancestor}$. If newly sampled buffer location is collision-free with other objects in the arrangement of $s_{current}$, it becomes the new valid buffer location for O_{in_buffer} .

C. Goal Attempting

To ensure a solution within a short time frame, STRAP implements a process called *goal attempting*, as shown in Fig. 4. After exploration in each iteration, it finds a feasible rearrangement plan from the current explored state to a state with goal arrangement. Since the process is executed at every iteration, it must be fast. To achieve this efficiency, STRAP leverages a state-of-the-art rearrangement local solver with lazy buffer allocation from [4], which is optimized for solving rearrangement tasks in cluttered environments in short time, even if the resulting plan is not optimal. After each goal attempting, the current best plan is re-evaluated by concatenating the action sequence leading to the re-explored state with the newly generated plan to the goal state. If this combined plan has a lower rearranging cost than the previous best plan, it is adopted as the new best plan. Consequently, states in the priority queue with an f value exceeding the cost of the new best plan are filtered out, thereby streamlining the process and reducing unnecessary sorting operations in the priority queue.

Algorithm 1 Plan Refinement

INPUT:

actionSequence - pick-and-place action sequence

arr_start - start arrangement: $\{1 : p_1, 2 : p_2, \dots\}$

```

1:  $B \leftarrow \{\}$ 
2:  $H \leftarrow \{0 : arr\_start\}$   $\triangleright$  Arrangement history
3:  $arr\_current = arr\_start$ 
4: for  $i, action \in Enum(actionSequence)$  do  $\triangleright$  start at 1
5:    $k, p_{pick}, p_{place} = action$ 
6:   if  $k \in B$  then  $\triangleright$  if object  $k$  was moved
7:      $ToBuf = B[k]$   $\triangleright$  previous action index on  $k$ 
8:      $C = getConstraint(H[ToBuf]/k)$ 
9:     for  $a \in actionSequence[ToBuf : i]$  do
10:       $C.addConstraint(a[2])$ 
11:     end for
12:      $p_1 = actionSequence[ToBuf][1]$ 
13:      $p_2 = actionSequence[i - 1][2]$ 
14:      $p_3 = actionSequence[ToBuf + 1][1]$ 
15:      $p_4 = actionSequence[i][2]$ 
16:      $P = BufGen(C)$ 
17:      $\triangleright$  Generate a buffer set under constraint  $C$ 
18:      $p^* \leftarrow \text{select from } P \min \sum_{j=1}^4 TravelCost(p^*, p_j)$ 
19:      $actionSequence[ToBuf][2] = p^*$ 
20:      $actionSequence[i][1] = p^*$ 
21:   end if
22:    $B[k] = i$ 
23:    $arr\_current[k] = p_{place}$   $\triangleright$  update arrangement
24:    $H[1] = arr\_current$   $\triangleright$  record arrangement
25: end for

```

D. Rearrangement Plan Refinement

After termination, the current best rearrangement plan should be refined further. Given a pick-and-place action sequence $[a_1, a_2, \dots, a_m]$, our refining process iteratively improves the plan until the cost no longer improves.

In the refinement process, the algorithm iterates over all the actions involving moving an object to a buffer location, called “to-buffer”, from the beginning of the sequence. Each iteration aims to find a better buffer location for “to-buffer” actions. Once an improved buffer location is identified, the object’s location is updated, and the subsequent “to-buffer” actions are optimized with the modified placement.

For a to-buffer action, the following locations are crucial:

- p_1 Last gripper position before placing at buffer.
- p_2 Last gripper position before picking at the buffer.
- p_3 First gripper position after placing at the buffer.
- p_4 First gripper position after picking at the buffer.

From the object’s entry into the buffer until its exit, we search for a collision-free buffer location during this interval and minimize the total travel cost from this buffer location to all four positions using sampling. This iteration process is demonstrated as the pseudocode in Alg. 1.

This paragraph provides more details of the refinement. Given an action sequence and initial arrangement, the goal

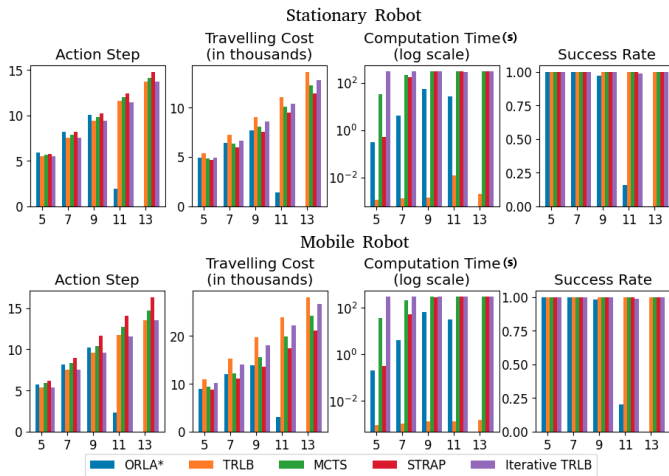


Fig. 5. The comparisons of different algorithms for stationary and mobile robots across a varying number of objects. The horizontal axis represents the number of objects.

is to optimize the buffer location. First, we initialize a dictionary B that contains the action index in the action sequence for each object moved previously. Additionally, H represents the arrangement history. Starting from line 4, each action in the sequence is processed. If the object has been moved before, it indicates it was placed in a buffer location since objects are not moved from their goal positions. From lines 6 to 7, if an object k was moved previously, we identify the action index $ToBuf$ where it was moved to a buffer location. From lines 8 to 11, we determine the obstacle constraints to ensure the buffer location does not collide from action $ToBuf$ to the current index. From lines 12 to 18, based on these constraints, we sample a set of possible buffer locations and evaluate the best one. Finally, from lines 19 to 20, the action sequence is updated accordingly.

V. EVALUATION

To ensure a fair comparison, we applied the disk test cases from TRLB [4] on both stationary and mobile settings across the following algorithms:

- 1) STRAP: Planner presented in this paper.
- 2) ORLA*: Planner [3] as the state-of-art.
- 3) TRLB: Planner [4] returns non-high-quality solutions.
- 4) Iterative TRLB: Runs the TRLB repeatedly on the task until timeout, returning the best result.
- 5) MCTS: Uses MCTS with reward function from [16], incorporating robot configuration in the state.

Given two random arrangements of disks of equal size (radius R), the objective is to rearrange them from one configuration to another on a 1000×1000 table. During the rearrangement, the distance between any two objects must remain greater than $2R$ to prevent collisions. In the simulation (upper part of Fig. 9), two setups are considered. For the stationary robot, represented as a box moving horizontally above the table, it picks and places objects by positioning itself directly above them. The total cost in this setup is the sum of its travel distance and the number of pick-and-place

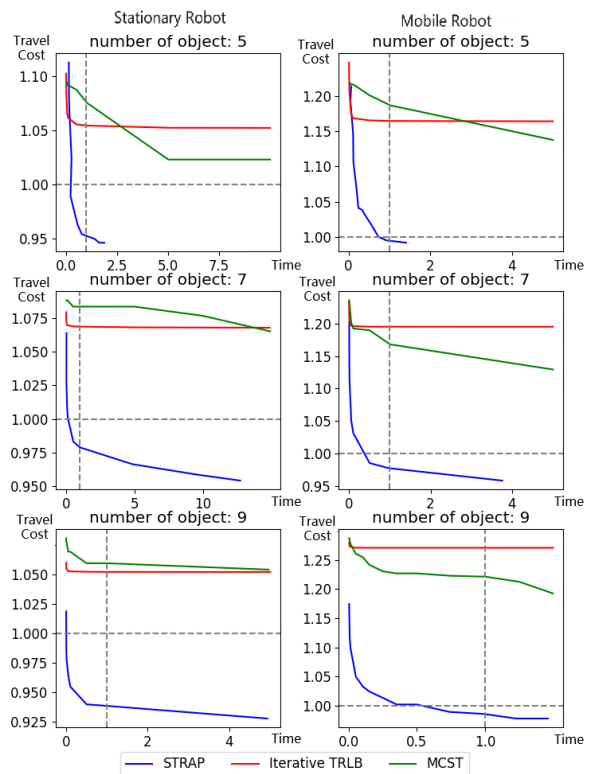


Fig. 6. The travel cost ratio compared to ORLA* over different timeout ratio of ORLA* on different number of objects for both the mobile and the stationary settings. The STRAP may stop early if the priority queue runs out of states.

actions. In contrast, the mobile robot, represented as a box that moves around the table, must navigate to the nearest location before executing a pick-and-place action. Here, the total cost is determined by the robot's travel distance and the number of pick-and-place actions.

A. Results

Our evaluation uses the same metrics as in [3]: average number of actions, average travel cost, average computation time, and success rate, while manipulation cost is one.

As illustrated in Fig. 5, we compare the general performance of all algorithms with 5 minutes timeout and 0.2 density, and it reveals the following trends: ORLA* struggles to scale effectively with increasing object number, regardless of whether the robots are stationary or mobile, whereas the other algorithms consistently maintain a high success rate. Regarding travel cost, both ORLA* and STRAP outperform others. ORLA* only occasionally shows very low cost in scenarios where its success rate is low.

To better understand the time consumption of ORLA*, we conducted tests with a one-hour timeout on 70 tasks in both mobile and stationary settings, each involving 11 objects. Our results show that in the mobile setting, 25% of the cases required more than 2500 seconds to execute, while in the stationary setting, 75% of the cases exceeded 1400 seconds.

Furthermore, we compare the travel costs of various algorithms over time against ORLA* for both mobile and

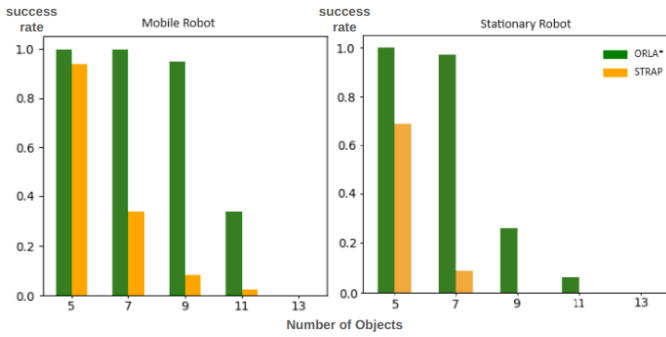


Fig. 7. Success rate of ORLA* and STRAP without goal attempting.

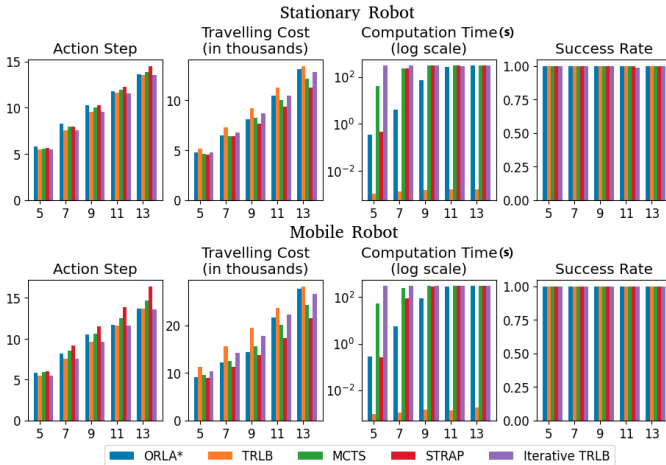


Fig. 8. The comparison of both ORLA* & STRAP with goal attempting. The horizontal axis represents the number of objects.

stationary settings, as shown in Fig. 6. This evaluation is restricted to scenarios (5, 7, 9 objects) where all algorithms achieve a high enough success rate. Given 30 tasks, we first determine the average computation time of ORLA* across different numbers of objects in both mobile and stationary settings. We then evaluate other algorithms on these same 30 tasks using different timeout ratios of ORLA*. As shown in Fig. 6, neither MCTS nor Iterative TRLB achieves a lower travel cost compared to ORLA* in reasonable time. Conversely, STRAP matches the travel cost of ORLA* while requiring significantly less computation time. Additionally, STRAP can lower travel costs more than ORLA* while using the same computation time. This is because STRAP includes the robot location in its state representation, making its state representation more complete than ORLA*, which does not consider the robot location.

B. Ablation Study

This section performs two ablation studies. First, we compare ORLA* and STRAP without goal attempting to verify whether lazy evaluation of ORLA* actually improves the performance as an A* based algorithm. Using 100 tasks with a timeout of 5 minutes, Fig. 7 demonstrates that ORLA* achieves a higher success rate than STRAP without goal attempting, indicating that lazy evaluation indeed benefits

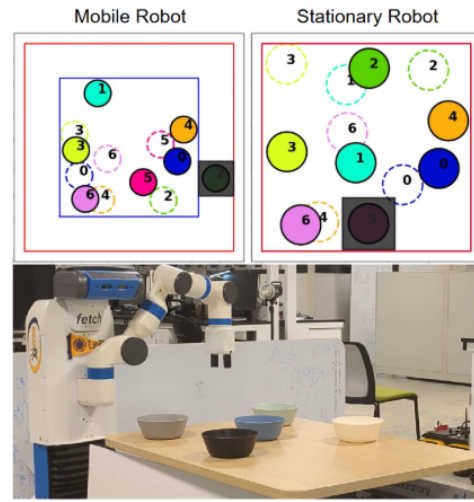


Fig. 9. Simulation and Real-World Experiments: In the simulation (upper), solid circles represent the object’s current positions, while semicircles indicate the desired locations. In the real-world experiment (lower), the mobile robot moves around the table to rearrange the bowls.

A*-based algorithms.

Next, we examine whether incorporating goal attempting into ORLA* can further improve its performance. In ORLA*, goal attempting can only be executed in the deterministic states due to the nature of the algorithm. As shown in Fig 8, the success rate of ORLA* is constantly high. However, ORLA*’s travel cost does not improve and be better than STRAP. The reason for this behavior is that there are very few deterministic states in ORLA*’s search tree.

C. Real-World Experiment

The accompanying video, as shown in the lower part of Fig. 9, demonstrates STRAP running on the Fetch robot [15] to rearrange bowls on the table. After the real-world experiment, we have multiple concerns. First, for the mobile robot, a lot of pick-and-place locations are close to each other, so certain mobile movements are not necessary. Therefore, in the future, this should be taken into account. Second, real-world rearrangement tasks are prone to failure if grasp estimation is inaccurate, highlighting the need for a recovery system to handle such failures.

VI. DISCUSSION

We present STRAP for high-quality tabletop object rearrangement, improving scalability for stationary and mobile robots. Evaluations show it outperforms ORLA*, achieving higher-quality solutions while maintaining feasibility. However, to avoid dead loops, STRAP limits state exploration per step, preventing optimality guarantees. Moreover, TRLB, ORLA*, and STRAP always position at the nearest location for single-object manipulations like pick-and-place, leading to inefficiencies. Future work should develop an algorithm that optimally selects locations for multiple manipulation actions.

REFERENCES

- [1] Jiammy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. "TidyBot: Personalized Robot Assistance with Large Language Models." *Autonomous Robots*, pp. 1087–1102, 47.8 (2023),
- [2] Dhruv Batra, Angel X. Chang, S. Chernova, Andrew J. Davison, Jia Deng, Vladlen Koltun, Sergey Levine, Jitendra Malik, Igor Mordatch, Roozbeh Mottaghi, Manolis Savva and Hao Su. "Rearrangement: A Challenge for Embodied AI." ArXiv abs/2011.01975 (2020)
- [3] Kai Gao, Yan Ding, Shiqi Zhang, and Jingjin Yu. "ORLA*: Mobile Manipulator-Based Object Rearrangement with Lazy A*." arXiv preprint arXiv:2309.13707 (2023)
- [4] Kai Gao, Darren Lau, Baichuan Huang, Kostas E. Bekris, and Jingjin Yu. "Fast high-quality tabletop rearrangement in bounded workspace." In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 1961-1967. IEEE, (2022)
- [5] Sertac Karaman, Matthew R. Walter, Alejandro Perez, Emilio Frazzoli, and Seth Teller. "Anytime motion planning using the RRT*." In *2011 IEEE international conference on robotics and automation*, pp. 1478-1483. IEEE, (2011).
- [6] Peter E. Hart, Nils J. Nilsson and Bertram Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," in *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100-107, July (1968).
- [7] Athanasios Krontiris and Kostas E. Bekris. "Dealing with Difficult Instances of Object Rearrangement." In *Robotics: Science and Systems*, vol. 1123. (2015).
- [8] Athanasios Krontiris and Kostas E. Bekris. "Efficiently solving general rearrangement tasks: A fast extension primitive for an incremental sampling-based planner." In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3924-3931. IEEE, (2016).
- [9] Shuai Han, Nicholas M. Stiffler, Athanasios Krontiris, Kostas E. Bekris, and Jingjin Yu. "Complexity results and fast methods for optimal tabletop rearrangement with overhand grasps." *The International Journal of Robotics Research* 37, no. 13-14 (2018): 1775-1795.
- [10] Yann Labbe, Sergey Zagoruyko, Igor Kalevtykh, Ivan Laptev, Justin Carpentier, Mathieu Aubry, and Josef Sivic. "Monte-Carlo Tree Search for Efficient Visually Guided Rearrangement Planning." *IEEE Robotics and Automation Letters*, vol.5, no.2, pp. 3715–3722, (2020).
- [11] Kai Gao, Si Wei Feng, and Jingjin Yu. "On running buffer minimization for tabletop rearrangement." In *17th Robotics: Science and Systems, RSS 2021*. MIT Press Journals, (2021).
- [12] Jur van Den Berg, Jack Snoeyink, Ming C. Lin, and Dinesh Manocha. "Centralized path planning for multiple robots: Optimal decoupling into sequential plans." In *Robotics: Science and systems*, vol. 2, no. 2.5, pp. 2-3. (2009).
- [13] Fan Wang and Kris Hauser. "Robot packing with known items and nondeterministic arrival order." *IEEE Transactions on Automation Science and Engineering*, 18, no. 4, pp. 1901-1915. (2020)
- [14] Ahmed Qureshi, Arsalan Mousavian, Chris Paxton, Michael C. Yip, and Dieter Fox. "Nerp: Neural rearrangement planning for unknown objects." In *Robotics: Science and systems*, (2021).
- [15] Melonee Wise, Michael Ferguson, Daniel King, Eric Diehr, and David Dymesich. *Fetch & freight : Standard platforms for service robot applications*. Technical report, Fetch Robotics, (2016).
- [16] Baichuan Huang, Xujia Zhang, and Jingjin Yu. "Toward Optimal Tabletop Rearrangement with Multiple Manipulation Primitives." *IEEE International Conference on Robotics and Automation*. (2024).
- [17] Mike Stilman and James Kuffner. "Planning Among Movable Obstacles with Artificial Constraints." *The International Journal of Robotics Research*. (2008).
- [18] Mike Stilman and James Kuffner. "Navigation among movable obstacles: Real-time reasoning in complex environments." *International Journal of Humanoid Robotics*. (2005).
- [19] Rui Wang, Kai Gao, Jingjin Yu, and Kostas Bekris. "Lazy Rearrangement Planning in Confined Spaces." *International Conference on Automated Planning and Scheduling*. (2022).
- [20] Martin Levihn, Jonathan Scholz, and Mike Stilman. "Hierarchical Decision Theoretic Planning for Navigation Among Movable Obstacles." *Algorithmic Foundations of Robotics X*. Springer. (2013)