

Variable Grasp Pose and Commitment for Trajectory Optimization

Jiahe Pan

University of Melbourne

Melbourne, Australia

jmpa@student.unimelb.edu.au

Kerry He

Monash University

Melbourne, Australia

kerry.he@monash.edu

Jia Ming Ong

Monash University

Melbourne, Australia

jong0052@student.monash.edu

Akansel Cosgun

Deakin University

Melbourne, Australia

akansel.cosgun@monash.edu

Abstract—We propose enhancing trajectory optimization methods through the incorporation of two key ideas: variable-grasp pose sampling and trajectory commitment. Our iterative approach samples multiple grasp poses, increasing the likelihood of finding a solution while gradually narrowing the optimization horizon towards the goal region for improved computational efficiency. We conduct experiments comparing our approach with sampling-based planning and fixed-goal optimization. In simulated experiments featuring 4 different task scenes, our approach consistently outperforms baselines by generating lower-cost trajectories and achieving higher success rates in challenging constrained and cluttered environments, at the trade-off of longer computation times. Real-world experiments further validate the superiority of our approach in generating lower-cost trajectories and exhibiting enhanced robustness. While we acknowledge the limitations of our experimental design, our proposed approach holds significant potential for enhancing trajectory optimization methods and offers a promising solution for achieving consistent and reliable robotic manipulation.

Index Terms—Robots, Grasping, Planning, Trajectory

I. INTRODUCTION

Robotic manipulation techniques have traditionally followed a two-step process [1]. Firstly, a grasp pose is determined based on the likelihood of grasp success using either model-based or deep learning-based approaches. Subsequently, the arm trajectory is computed by employing sampling-based planning methods such as Rapidly-exploring Random Trees (RRT) [2] or Probabilistic Roadmaps (PRM) [3], or through trajectory optimization algorithms such as Covariant Hamiltonian Optimization for Motion Planning (CHOMP) [4] or Sequential Convex Optimization (TrajOpt) [5]. However, solely selecting the end-effector goal pose based on predicted grasp success without considering the arm motion can lead to suboptimal arm trajectories in terms of factors like completion time or energy efficiency. This limitation presents an opportunity for improvement, particularly when grasp poses have similar probabilities of success or when the robot operates in cluttered environments where different grasp poses may yield vastly different trajectories.

Previous works have explored various approaches for trajectory planning with variable grasp poses. Dragan et al. [6] extended CHOMP [4] to include goal sets as end-point




Fig. 1: Trajectory planning for various grasp poses for an occluded cuboid target (top). The robot arm starts in a highly constrained position (arm-tuck), and plans a trajectory towards grasping a cylindrical target (bottom).

constraints. Murooka et al. [7] combined sampling-based planning and gradient-based optimization to solve for optimal trajectories that reach towards sampled grasp poses. Horowitz et al. [8] explored grasp and manipulation planning as a single optimal control problem. Ichnowski et al. [9] used Dex-Net 4.0 [10] and an SQP-based algorithm for optimizing minimum-time trajectories in warehouse picking tasks.

Vahrenkamp et al. [11] evaluated trajectories towards different grasp poses using an online quality measurement module. Jetchev et al. [12] used 3D laser point cloud information to construct an occupancy probability map for trajectory generation in cluttered environments. Dogar et al. [13] analyzed physical contacts with cluttered scenes to improve planner efficiency and success rates. Berenson et al. [14] used goal regions in the manipulator’s workspace for grasp and trajectory planning. Bergman et al. [15] proposed Receding Horizon Planning (RHP), and Mandalika et al. [16] introduced Lazy Receding Horizon A* for path planning in cluttered scenes.

In this paper, we propose an approach that incorporates the selection of grasp poses into trajectory optimization. The ability to vary the grasp pose relative to the target object becomes particularly advantageous in cluttered environments, as the chosen grasp pose directly affects the quality and feasibility of the robot arm’s trajectory during the reach and grasp actions. Therefore, our approach aims to explore the space of all possible grasp poses for a given target object and select one that leads to a more efficient trajectory for performing the grasping action. Additionally, we explore trajectory

commitment in this paper in which we gradually shrink the optimization horizon towards the goal region to increase the efficiency of trajectory optimization. To validate our approach, we implement several baseline methods based on TrajOpt and RRT-Connect. We compare these baselines against our approach in four distinct task environments featuring primitive targets and obstacles arranged in different configurations, with varying initial arm poses. Finally, we evaluate our method by analyzing planning time, trajectory cost, and success rates of the different approaches in simulated and real-world experiments on a 7-DoF serial manipulator.

Our approach differs from existing methods by directly sampling grasp poses during planning, without pre-computing or ruling out infeasible grasps. We treat TrajOpt as a black-box optimizer, refining segments of the trajectory near the target object. Our refinement step ensures optimality of the entire trajectory. Notably, our method is accessible to researchers without extensive knowledge of grasp-pose synthesis or trajectory planning with high-DOF robotic manipulators, and despite treating TrajOpt as a black-box optimizer, users can still control its high-level input parameters which directly relate to the trajectories generated.

II. PROPOSED APPROACH

We build upon the original TrajOpt algorithm [5], treating it as a black-box optimizer. Given information about the task environment, an initial robot arm pose, and a goal pose, TrajOpt outputs the optimal trajectory from the initial to the final goal pose. We augment TrajOpt with an outer optimization loop that randomly samples different grasp poses relative to the target object. For each sampled pose, we use TrajOpt to solve for an optimal trajectory.

A key hypothesis behind our commitment idea is that since only the grasp pose varies between iterations, the waypoints near the target object are more affected than those near the starting pose. Therefore, it is unnecessary to recompute the entire trajectory in each iteration. Instead, we gradually shrink the optimization horizon towards the goal region to increase TrajOpt's efficiency. This reduction in the optimization horizon helps the manipulator escape from constrained configurations and navigate clustered environments. We define the waypoints outside TrajOpt's optimization horizon as committed waypoints, which are stored in memory. As the outer-loop iterations progress, we add new committed waypoints and further shrink the optimization horizon towards the goal region.

Sec. II-A provides a brief overview of the TrajOpt algorithm and illustrates how it is embedded as a black-box optimizer in our algorithm. Sec. II-B presents how we parameterize the variable grasp pose for a primitive object shape. In Sec II-C, we provide a description of our algorithm and illustrate the high-level logic using a flowchart.

A. Black-box Optimizer

Trajectory optimization for robotic manipulation is often posed as a non-convex, constrained optimization problem. In the TrajOpt implementation, the cost function is defined as

the squared sum of the accumulative velocity, acceleration and jerk costs across consecutive joint-space waypoints of the trajectory, which relates to the length of the physical trajectory. Equality and inequality constraints arise from the robot's joint limits, avoiding collision objects, and reaching the goal pose. The optimization problem is solved iteratively through a series of convex optimization problems using SQP. The penalty coefficients for the constraints vary according to improvement in the cost, until it converges and a valid solution trajectory is found, or an iteration limit has been reached.

Our algorithm interacts directly with the following 5 parameters of the TrajOpt algorithm:

- **Start pose** - The initial state of the robot.
- **Goal pose** - The target pose that TrajOpt will plan the trajectory towards.
- **InitTraj** - The initial trajectory given to TrajOpt to initialize the SQP algorithm.
- **Tolerance** - The tolerance for constraint-violation. Values greater than this parameter leads to a trajectory being considered invalid by TrajOpt.
- **Waypoints** - The number of joint-space waypoints to plan for given some initial and final poses.

Overall, the TrajOpt algorithm is treated as a black-box optimizer which takes inputs of planning queries and outputs feasible solution trajectories if found.

B. Grasp Pose Parametrization

To introduce how we parameterize the set of possible ways a robot can grasp a primitive object, let us first consider a cuboid. When the robot faces this object directly, it has 3 feasible grasp direction - {front, left, right}, and the extra grasp pose from the top face and all the grasp angles in between. First analyzing a single side (see Fig. 2), we define the parameters (θ, s) , where θ governs the grasp angle from horizontal, and s is the location of the end-effector's centre point along an arc which varies with θ . This parametrization model maintains a high overlap-region between the cross-sections of the end-effector and the object, hence ensuring stable grasps for all feasible poses. The bounds for (θ, s) are set relative to the object's dimensions, allowing our model to scale automatically to different-sized objects. Our model can easily generalize to cylindrical and spherical objects by adding a rotational degree of freedom in the parametrization, using the parameter α (see Fig. 2).




Fig. 2: (Left) Side view of grasp angles of 0° , 45° and 90° (left to right). Red dot is the center point position of the end-effector, purple line shows the possible positions of s , overlap regions are highlighted. (Right) α represents the approach angle of the grasp (Top view).

Symbol	Description
μ_{\max}	A predefined tolerance threshold ensuring that the final solution trajectory satisfies goal constraints and is collision-free.
μ_0	The initial value of the penalty coefficient, initialized at two orders of magnitude higher than μ_{\max} . This relaxes the problem initially, aiding TrajOpt in generating trajectories efficiently, especially for highly constrained environments.
MaxIters	The iteration limit for the outer optimization loop, independent of any iteration limits within the TrajOpt algorithm.
\mathbf{q}_s	Initial robot configuration from the problem description
N_0	A fixed parameter defining the total number of waypoints to plan for over the entire horizon.
\mathbf{q}_0	The current starting robot configuration (Start pose)
\mathbf{G}	The randomly-sampled Cartesian grasp pose in each iteration. (Goal pose)
\mathbf{x}_0	The linearly-interpolated trajectory between \mathbf{q}_0 and \mathbf{G} , unless in the refinement step, where TrajOpt is initialized with the concatenated trajectory. (InitTraj)
μ	The current constraint-tolerance (Tolerance)
N	The current number of waypoints inside the optimization horizon, initially set to N_0 . (Waypoints)
\mathcal{C}	The list of committed trajectory waypoints, updated in each commitment step.
s_{thres}	A threshold variable that determines when commitment steps are triggered. Initialized as 1 and incremented at each commitment step, encouraging exploration before making further commitments.
s	The count of successful TrajOpt runs for the current optimization horizon. Resets to 0 after each commitment step.

TABLE I: Summary of parameters and variables. Highlighted variables are used as inputs to the TrajOpt algorithm.

C. The Algorithm

The algorithm’s parameters and variables are shown in Table I, and a flowchart representation is shown in Fig. 3. A summary of key aspects of the algorithm are as follows:

- 1) **Initialization:** In the first iteration, a grasp pose, \mathbf{G} , is randomly sampled with respect to the target object and converted to the corresponding pose in the robot’s task-space frame. Then, TrajOpt is executed with the initial parameters: $\mathbf{q}_0 = \mathbf{q}_s$, $N = N_0$, and $\mu = \mu_0$.
- 2) **Commitment Step:** In any outer iteration, if TrajOpt successfully generates a solution trajectory, we store it in a temporary list of maximum size s_{thres} , and increment s by 1. If $s = s_{thres}$, a commitment step is triggered, where we pick the lowest-cost trajectory from the temporary list and add a fraction of its waypoints to the committed list, \mathcal{C} . Then, we decrease the penalty coefficient, μ (unless it has reached μ_{thres}), reduce N by the number of newly-committed waypoints, increase s_{thres} , reset s to 0, and empty the temporary list. If TrajOpt fails to find a solution, the algorithm proceeds to the next iteration without updating any variables.
- 3) **Grasp Pose Sampling:** Using our parametrization model from Sec. II-B, a grasp pose is randomly sampled in each outer-loop iteration, and planning is performed for the current optimization horizon. This process continues until MaxIters is reached.
- 4) **Upper Bound on Committed Waypoints:** To prevent an excessively long commitment list, an upper bound

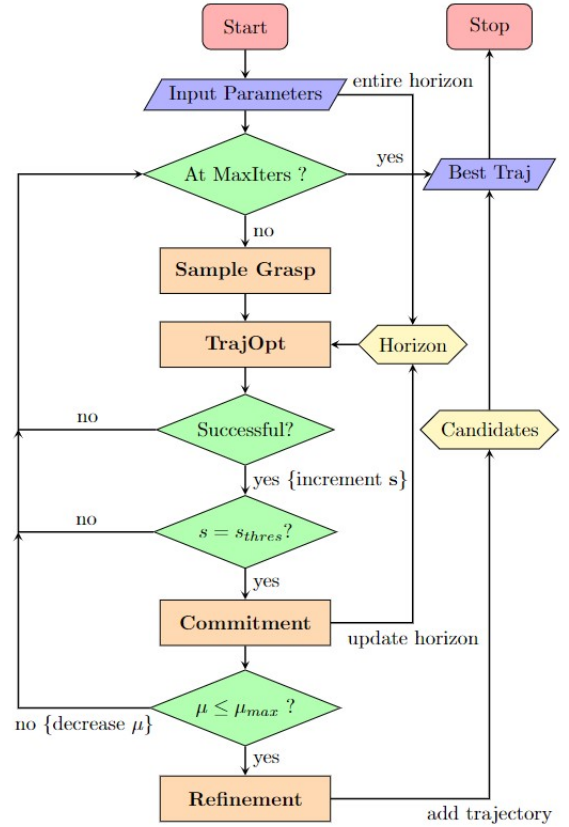


Fig. 3: Flowchart representing the logic of the algorithm.

is imposed on the size of \mathcal{C} . Once reached, no further commitment steps take place. The fraction of waypoints added at each commitment step affects how quickly this bound is reached and is a hyper-parameter that requires tuning.

- 5) **Refinement Step:** If TrajOpt successfully solves in any outer-loop iteration and the condition $\mu \leq \mu_{\max}$ is satisfied, a refinement step is performed. In this step, the current solution trajectory is concatenated with the list of committed waypoints and passed to TrajOpt for further optimization. The resulting optimal trajectory spans the entire horizon and is added to a list of candidate solution trajectories.
- 6) **Final Solution:** When MaxIters is reached, the final solution trajectory is selected from the candidate list based on the lowest cost. If the list is empty, indicating that no valid candidate trajectories were generated, the algorithm returns with failure.

III. EXPERIMENTS & RESULTS

In this section, we present a comprehensive evaluation of our proposed approach by comparing it to several baseline methods. We conduct performance comparisons in four distinct task environments, which are pre-constructed and include simulations as well as real-world experiments using the Fetch robot. By systematically analyzing the results across these environments, we aim to assess the effectiveness and efficiency of our approach in different scenarios.




Fig. 4: Visualisation of our algorithm planning a grasp towards a target (highlighted) in simulation, for 4 different scenes based on a shelf or tabletop workspace, and with multiple primitive objects in different poses. Each scene also plans from different initial arm configurations, including the constrained arm-tuck pose.

A. Baseline Approaches

For analysis of both the variable grasp pose planning and trajectory commitment ideas, we have designed the following 4 baseline approaches. Note that the baselines do not include any aspect of our trajectory commitment approach.

- **RRT-Connect:** Randomly sample a grasp pose and run RRT-Connect once.
- **Fixed Goal:** Run the TrajOpt optimizer once, for a fixed hand-chosen grasp pose relative to the target object.
- **Variable Goal, Single Iteration:** Randomly sample a grasp pose relative to target object, and run TrajOpt once.
- **Variable Goal, Multi-Iteration:** Run TrajOpt 5 separate times, each with a different randomly sampled grasp pose, then return trajectory with lowest cost.

B. Experimental Scenes

For experimentation, we have constructed 4 different task scenes. They are mostly distinct in the initial arm configuration, the objects' shapes, dimensions and poses, and whether the workspace is a shelf-like or table-top scenario. Figure 4 shows the following scenes, **Box Shelf**, **Box Table**, **Cylinder 1** and **Cylinder 2**, in the order of left to right, top to bottom. The target object in each scene is highlighted.

- **Box Shelf** - A shelf-like scene, with 3 large boxes ($25\text{cm} \times 25\text{cm} \times 8\text{cm}$) arranged in a row. Initially the arm is inside the shelf.
- **Box Table** - A table-top scene, with 3 medium-sized boxes ($15\text{cm} \times 15\text{cm} \times 6\text{cm}$) arranged such that the target is occluded by another collision object. Initially in arm-tuck position.
- **Cylinder 1** - A table-top scene, with 3 small cylinders arranged in a row. Initially in arm-tuck position.
- **Cylinder 2** - Same as the **Cylinder 1** scene, apart from the arm initially being in one of our randomly-generated

configurations, which from previous experimentation is a very easy pose to plan from.

C. Metrics

For both the simulated and real world experiments, we base our discussion on the following three metrics:

- Planning Time (s)
- Trajectory Cost (using the cost definition from TrajOpt)
- Success Rate (%)

D. Simulation Results

We ran 300 independent trials for each of the 5 approaches (4 baselines + Ours), for each of the 4 different task scenes shown in Fig 4. The results are shown in Table II.

In the simulated scenes, our approach by far had the highest success rate in all scenarios. Our approach also consistently generated the lowest-cost trajectories compared to the baselines in three out of four scenarios. For the **Cylinder 2** scene, where the initial arm configuration is unconstrained, the **Variable Goal, Multi-Iteration** baseline generated slightly lower-cost trajectories on average, while matching the 100% success rate of our approach. For the **Fixed Goal** baseline, we chose to not include its success rate, as the fixed grasp poses it plans towards are hand-picked for each scene to guarantee feasibility. Since all 4 other approaches adopt the automated random grasp sampling, the success rate is not a meaningful metric for comparison here.

Regarding computing time, the **RRT-Connect** baseline consistently exhibited the shortest computation time for generating a trajectory. However, it's important to note that the trajectories generated by this baseline were significantly less efficient in terms of cost compared to trajectory optimization methods. It's also worth noting that both our approach and the **Variable Goal, Multi-Iteration** baseline exhibited significantly longer planning times compared to the single-iteration trajectory optimization baselines, as a consequence of running TrajOpt for multiple grasp poses.

Overall, the results demonstrate that our approach outperforms the baselines in terms of generating lower-cost trajectories and achieving higher success rates, especially for constrained tasks involving cluttered objects and challenging initial configurations. However, the improvements were less significant in less-constrained task environments.

During the simulated experiments, we made an observation that confirmed one of our previous hypotheses. As our approach committed more trajectory waypoints, resulting in a shrinking optimization horizon, TrajOpt's speed in generating a solution trajectory or returning failure significantly increased. This speed-up is attributed to the fact that the dimension of the optimization problem is proportional to the number of trajectory waypoints to plan for. Consequently, our approach was able to plan for approximately double the number of outer-loop iterations (the number of times the TrajOpt optimizer is run) compared to the **Variable Goal, Multi-Iteration** baseline, within a similar total planning time. This improvement directly contributed to the higher success rate achieved by our approach.

	Box Shelf			Box Table			Cylinder 1			Cylinder 2		
	%	—	sec	%	—	sec	%	—	sec	%	—	sec
Method	Succ.	Cost	Time	Succ.	Cost	Time	Succ.	Cost	Time	Succ.	Cost	Time
RRT-Connect [2]	22	0.540	4.3	28	1.99	1.3	36	1.77	1.3	37	9.76	1.3
Fixed Goal	-	0.130	12.9	-	0.69	25.6	-	0.72	20.6	-	4.81	17.3
Variable Goal, Single Iteration	19	0.130	15.6	20	0.69	30.8	33	0.72	22.0	34	4.83	16.4
Variable Goal, Multi-Iteration	82	0.128	77.8	74	0.7	158.5	93	0.71	211.5	100	4.74	81.8
Variable Goal, Multi-Iteration, Commitment (Ours)	95	0.122	79.8	97	0.61	123.3	96	0.66	208.2	100	4.88	96.9

TABLE II: Simulation results from 300 trials for each scenario per method. Best results are displayed in **bold**.

	%	—	sec
Method	Succ.	Cost	Time
Variable Goal, Single Iteration	40	0.76	23.6
Variable Goal, Multi-Iteration	80	0.68	80.7
Variable Goal, Multi-Iteration, Commitment (Ours)	80	0.42	80.1

TABLE III: Real robot results from 10 trials for each method. Best results are displayed in **bold**.

E. Real Robot Results



Fig. 5: Top view of our real-life experimental scene. The cylinder (highlighted) is fixed to be the target object across all trials.

In our real-world experiments, we employed the Fetch robot within a table-top environment featuring a fixed cylindrical target object and several cuboids of varying dimensions and poses (see Fig. 5). To simplify the pose determination of all objects, including the table, we utilized Alvar markers. We conducted 10 independent trials for each of the three approaches, and the results are shown in Table III.

The real-world experiment results demonstrate that our approach consistently generated lower-cost trajectories compared to the other two baselines. Moreover, both our approach and the **Variable Goal, Multi-Iteration** baseline exhibited a remarkable improvement in success rate, succeeding in 8 out of the 10 trials. This significant enhancement in robustness is accompanied by a trade-off of increased planning time. However, it is important to note that the planning times for our approach and the **Variable Goal, Multi-Iteration** baseline were comparable.

For a visual comparison of the solution trajectories generated by our approach and the **Variable Goal, Multi-Iteration** baseline, Figure 6 illustrates the qualitative difference between the solutions of the two approaches, highlighting the superior performance of our approach in terms of trajectory quality.

IV. DISCUSSION

Of the four simulation scenes we designed, our algorithm consistently generated the lowest-cost trajectories and exhibited the highest success rate in three of them. However, in the less constrained environment of the **Cylinder 2** scene, the **Variable Goal, Multi-Iteration** baseline outperformed our

approach in trajectory cost while maintaining a similar level of consistency. Consequently, we hypothesize that our approach may not offer a significant advantage over existing motion planners in easier task environments. However, it excels in highly constrained tasks by efficiently selecting better grasp poses and consistently generating corresponding lower-cost trajectories.

Regarding our experimental design, a key limitation is the lack of variation in the task scenes. The **Box Table** and **Cylinder 1** scenes feature a table-top environment with three objects of the same type scattered on the table, while the **Box Shelf** and **Cylinder 2** scenes have similar arrangements. Throughout the 300 simulation trials conducted for each scene, all aspects of the environments remained unchanged. This design choice enables consistent calculation of success rates for each approach. Additionally, since most of the approaches, including our method, incorporate random sampling, conducting repeated trials in consistent task scenes allows for more accurate interpretation of performance metrics using mean and standard deviation values.

However, this design also limits the ability to draw broader conclusions about the performance of our method. To achieve more generalizable conclusions, it may be necessary to introduce higher variation in experimental scenes by incorporating different collision object arrangements across trials. One approach is to randomly generate a large number of distinct, highly constrained task scenes and conduct slightly fewer repeated trials for each scene. This method provides the distribution of outcomes for each scene while still utilizing mean and standard deviation as valid metrics to represent the performance of each approach. Moreover, this approach introduces greater scene variation, which helps reduce systematic bias towards scenes with specific characteristics.

Another potential improvement in scene design involves the choice of initial arm configurations. Generating random initial configurations that avoid immediate collisions with objects in the highly constrained task space can be challenging. One solution is to position the robot directly in a feasible pose within the task environment and then randomly sample a large number of task space goal poses for the end-effector. An efficient trajectory planning algorithm such as RRT-Connect can then be used to move the arm towards each task space goal pose. Successful trajectories from this process yield valid joint-space arm configurations that can be used as initializations. Overall, these adjustments to our experimental design could lead to more powerful conclusions regarding the performance




Fig. 6: Trajectories generated by the **Variable Goal, Multi-Iteration** baseline (top-row) and our approach (bottom-row). Our approach produced a lower-cost trajectory.

and robustness of our method, which would generalize well to a wider range of scenes.

V. CONCLUSIONS

In this paper, we introduced the concept of variable grasp pose planning and trajectory commitment, and the key motivations behind these ideas. Through both simulated and real experiments, we demonstrated that our proposed algorithm exhibits more consistent and robust performance across a series of four task scenes compared to various baseline approaches based on existing motion planning algorithms. However, we also discussed the limitations of our experimental design, which provides valuable insights for future improvements.

One immediate focus for future work is the design of a method to generate randomized task scenes that incorporate greater variation in object types, dimensions, positions, orientations relative to the robot, and the initial configuration of the robot itself. By testing the proposed method in a broader range of settings, we can obtain more meaningful results and evaluate its performance under diverse conditions. Additionally, improving the computation time for the proposed method, such as by parallelizing computing trajectories of multiple grasp poses [7], will allow our method to be even more competitive with existing approaches.

For robotic manipulators with a mobile base, such as the Fetch robot, an intriguing aspect to explore is incorporating the base pose of the robot as additional degrees of freedom in grasp planning and trajectory optimization. Preliminary results suggest that the base pose can directly impact the optimal grasp pose and trajectory cost for moving the arm manipulator toward the target object. Therefore, investigating the influence of the base pose on the overall performance and cost of the algorithm holds promise for further advancements.

In summary, this paper presents a novel approach for variable grasp pose planning and trajectory commitment, showcasing its superiority over baseline approaches in terms of consistency and robustness. While acknowledging the limitations of our experimental design, we propose future directions that can build upon the concepts and methods introduced here. These future endeavors have the potential to address current limitations, enhance the overall performance, and increase the robustness of the algorithm in real-world scenarios.

REFERENCES

- [1] Rhys Newbury, Morris Gu, Lachlan Chumbley, Arsalan Mousavian, Clemens Eppner, Jürgen Leitner, Jeannette Bohg, Antonio Morales, Tamim Asfour, Danica Kragic, et al. Deep learning approaches to grasp synthesis: A review. *IEEE Transactions on Robotics*, 2023.
- [2] J.J. Kuffner and S.M. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *ICRA*, 2000.
- [3] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 1996.
- [4] Nathan Ratliff, Matt Zucker, J. Andrew Bagnell, and Siddhartha Srinivasa. Chomp: Gradient optimization techniques for efficient motion planning. In *IEEE International Conference on Robotics and Automation*, 2009.
- [5] John Schulman, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, and Pieter Abbeel. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 2014.
- [6] Anca D. Dragan, Nathan D. Ratliff, and Siddhartha S. Srinivasa. Manipulation planning with goal sets using constrained trajectory optimization. In *IEEE International Conference on Robotics and Automation*, 2011.
- [7] Takayuki Murooka, Artur Istvan Karoly, Felix von Drigalski, and Yoshihisa Ijiri. Simultaneous planning of grasp and motion using sample regions and gradient-based optimization. In *IEEE International Conference on Automation Science and Engineering (CASE)*, 2020.
- [8] Matanya B. Horowitz and Joel W. Burdick. Combined grasp and manipulation planning as a trajectory optimization problem. In *IEEE International Conference on Robotics and Automation*, 2012.
- [9] Jeffrey Ichnowski, Michael Danielczuk, Jingyi Xu, Vishal Satish, and Ken Goldberg. Gomp: Grasp-optimized motion planning for bin picking. In *IEEE International Conference on Robotics and Automation*, 2020.
- [10] Jeffrey Mahler, Matthew Matl, Vishal Satish, Michael Danielczuk, Bill DeRose, Stephen McKinley, and Ken Goldberg. Learning ambidextrous robot grasping policies. *Science Robotics*, 2019.
- [11] Nikolaus Vahrenkamp, Martin Do, Tamim Asfour, and Rüdiger Dillmann. Integrated grasp and motion planning. In *IEEE International Conference on Robotics and Automation*, 2010.
- [12] Nikolay Jetchev and Marc Toussaint. Trajectory prediction in cluttered voxel environments. In *IEEE International Conference on Robotics and Automation*, 2010.
- [13] Mehmet R Dogar, Kaijen Hsiao, Matei Ciocarlie, and Siddhartha Srinivasa. *Physics-based grasp planning through clutter*. MIT Press: Cambridge, MA, USA, 2012.
- [14] Dmitry Berenson, Siddhartha S. Srinivasa, Dave Ferguson, Alvaro Collet, and James J. Kuffner. Manipulation planning with workspace goal regions. In *IEEE International Conference on Robotics and Automation*, 2009.
- [15] Kristoffer Bergman, Oskar Ljungqvist, Torkel Glad, and Daniel Axehill. An optimization-based receding horizon trajectory planning algorithm. *IFAC-PapersOnLine*, 2020.
- [16] Aditya Mandalika, Oren Salzman, and Siddhartha Srinivasa. Lazy receding horizon a* for efficient path planning in graphs with expensive-to-evaluate edges. In *International Conference on Automated Planning and Scheduling*, 2018.