

GOMP: Grasp-Optimized Motion Planning for Bin Picking

Jeffrey Ichnowski*, Michael Danielczuk*, Jingyi Xu*, Vishal Satish*, and Ken Goldberg*

Abstract—Rapid and reliable robot bin picking is a critical challenge in automating warehouses, often measured in picks-per-hour (PPH). We explore increasing PPH using faster motions based on optimizing over a set of candidate grasps. The source of this set of grasps is two-fold: (1) grasp-analysis tools such as Dex-Net generate multiple candidate grasps, and (2) each of these grasps has a degree of freedom about which a robot gripper can rotate. In this paper, we present Grasp-Optimized Motion Planning (GOMP), an algorithm that speeds up the execution of a bin-picking robot’s operations by incorporating robot dynamics and a set of candidate grasps produced by a grasp planner into an optimizing motion planner. We compute motions by optimizing with sequential quadratic programming (SQP) and iteratively updating trust regions to account for the non-convex nature of the problem. In our formulation, we constrain the motion to remain within the mechanical limits of the robot while avoiding obstacles. We further convert the problem to a time-minimization by repeatedly shorting a time horizon of a trajectory until the SQP is infeasible. In experiments with a UR5, GOMP achieves a speedup of 9x over a baseline planner.

I. INTRODUCTION

With increasing e-commerce, robots are in demand for bin picking. This “pick-and-place” operation is ostensibly simple, as humans can do it with little effort. However, the same is not true for robots. Having robots take on this task requires tackling many computational challenges in order to make it reliable and fast, including sensing, grasp analysis, motion planning, and executing the actual robot arm’s motion. With recent advances speeding up sensing and grasp analysis, motion planning and robot arm movement are rapidly becoming the bottleneck [1], [2]. In this paper, we present Grasp-Optimized Motion Planner (*GOMP*), an algorithm that speeds up the execution of a bin-picking robot’s operations by incorporating the dynamics of a robot arm and a *set* of candidate grasps produced by grasp analysis into an *optimizing motion planner*.

The pipeline for GOMP starts with the observation that grasp analysis tools can provide a set of candidate grasps. A motion planner can then select a grasp from the set that leads to the fastest motion between picking up and placing the object. The variation in speed of motion comes from both location and angle of grasp. Grasping locations and angles closer to the placement configuration, or that present alternate routes around obstacles can allow for faster bin picking motions.

When generating optimized trajectories between pick and place, GOMP allows the grasp point to vary over one degree

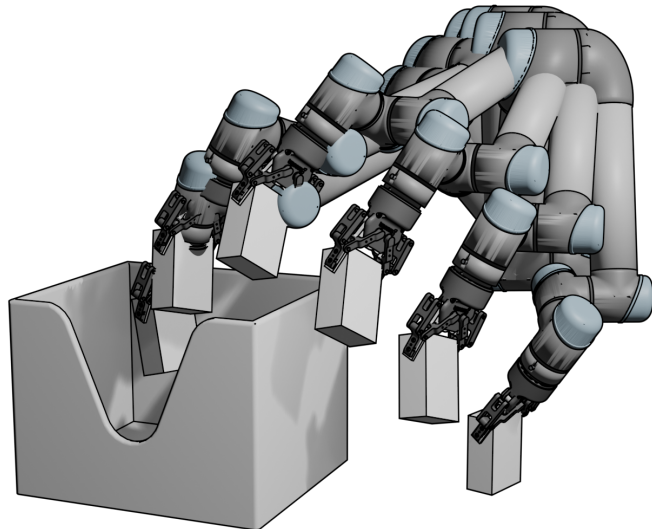


Fig. 1. GOMP computes minimum-time trajectories that avoid obstacles, while constraining the start and end configurations to a *set* computed by grasp analysis. Although many bin-picking systems limit the robot to picking from a top-down vector, the start and end configurations of a trajectory generated by GOMP are allowed to vary over the degree of freedom as defined by the grasp axis. This figure shows the optimal trajectory computed by GOMP for the UR5 robot with obstacles as illustrated.

of freedom. This degree of freedom for parallel-jaw grippers is around the grasp axis. For suction, the degree of freedom is around the suction cup contact normal. While traditional approaches limit the grasp to a single direction of approach (e.g., top-down) to a grasp-analysis frame, we propose allowing an optimization to select an approach angle that results in the same pair of points in contact with the gripper jaws, while allowing the robot to make faster motions.

GOMP incorporates the mechanical limits (e.g., each joint’s maximum torque and velocity) of the robot arm to minimize the resulting motions’ execution time. This minimization is non-convex, as it must avoid obstacles in the robot’s workspace (e.g., the bin sides) and allow for degrees of freedom in the forward-kinematic space of the start and goal poses while taking into account the mechanical limits. To address this problem, GOMP uses a sequential convex optimization on a discretization of a trajectory into a fixed number of waypoints at a fixed time interval that: (1) iteratively updates trust regions between convex optimization steps to account for the constraints of the start and goal poses and obstacles, (2) explicitly constrains joint motions to remain within their velocity and torque limits, and (3) minimizes the acceleration to produce smooth trajectories.

To minimize motion time, the sequential convex opti-

*Jeffrey Ichnowski, Mike Danielczuk, Jingyi Xu, Vishal Satish, and Ken Goldberg are with the AUTOLAB, University of California at Berkeley, Berkeley, CA 94720, U.S.A. {jeffi, mdanielczuk, jingyi-xu, vsatish, goldberg}@berkeley.edu

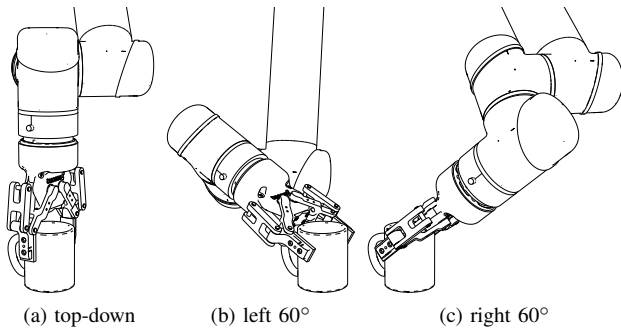


Fig. 2. Grasp analysis produces a pair of grasp points that a parallel jaw gripper is able to grasp from any angle rotated around the grasp axis created connecting the two grasp points. The optimizing motion planner starts with the traditional top-down grasp from (a) and optimizes around the axis to find the fastest motion.

mization above is repeated with a decreasing number of waypoints until the minimization is no longer feasible. This repeated optimization is sped up through the use of warm starting the subsequent optimization based upon the previous optimization.

In experiments, we integrate grasp analysis with the optimizing motion planner and deploy it on a UR5 robot arm [3] with a Robotiq gripper [4]. For the grasp analysis, we use a modified Dex-Net 4.0 [5] which can produce multiple grasp candidates in sub-second times. The result of these is an optimized motion with no observed loss of reliability, resulting in faster PPH than prior work.

This paper provides 3 main contributions:

- 1) The GOMP problem formalization and algorithm for minimizing the execution time of a robot’s motion while accounting for obstacles, robot’s mechanical limits, and grasp constraints.
- 2) Grasp-analysis-based pose constraints that allow an optimization to speed up motions by varying over a degree of freedom implied by the gripper’s design.
- 3) Experimental results suggesting GOMP motions can significantly speed up the bin picking pipeline.

II. RELATED WORK

Data-driven grasp-analysis or grasp planning algorithms for parallel jaw grippers, such as Dex-Net [6], [7], [8], [5], GG-CNN [1], GPD [9], or FC-GQ-CNN [2], typically take sensor input (e.g., an object mesh, a depth-camera image), perform some pre-processing (e.g., image inpainting), and produce either a grasp or grasp quality score for a presampled grasp candidate [10]. The majority of these algorithms are based on convolutional neural networks (CNNs) and may be learned from human annotations [11], simulated training data [12], [13], human or self-supervised labels from grasps attempted on a physical system [14], [15], [16] or a combination of the above [17]. These data-driven methods often represent grasps using a center-axis [5] or rectangle formulation [11] in the image plane, resulting in 4 degrees of freedom (a 3D translation, plus a rotation about the camera z-axis). Recent work has also explored introducing additional

degrees of freedom for grasps in cluttered environments [18], [19], [20], [21], noting that top-down grasps leave out a wide range of feasible high quality grasps on many objects [22]. In this paper, we build upon the 4 degree of freedom representation commonly used with CNNs and propose using the output coordinate frame as a basis for an infinite set of grasps around the grasp axis. We note that this additional degree of freedom applies also to suction cup grasps, which can be rotated around the surface normal of the grasp point. While we propose that many grasp-analysis algorithms could be incorporated into GOMP, we use FC-GQ-CNN as it can produce multiple grasp candidates in sub-second times.

Generating motion plans from a multiple start to multiple goal configurations (e.g., by varying degrees of freedom on the configurations), is well-suited for sampling-based motion planners such as PRM [23], RRG [24], and bi-directional RRTs [25]. While asymptotically-optimal variants of these planners [24] are guaranteed to produce optimal motions eventually, the slow convergence rate of these planners in 6+ degree-of-freedom robot arms does not lend itself to producing time-optimal motion plans fast enough to improve PPH. While sampling-based motion planners lend themselves to parallelization [26], [27], the associated speed up in convergence rate may still be insufficient. In any case, the generated motions often still require subsequent smoothing [28], time-parameterization [29], or both [30] in order to run smoothly and quickly on a robot arm.

Optimizing motion planners for robot such as Trajopt [31], CHOMP [32], STOMP [33], ITOMP [34], based on interior point optimization [35], based on gradients [36], and based on Gaussian Process [37], through various formulations of a minimization problem work to produce a locally minimized trajectory from a start to goal pose while avoiding obstacles. In prior formulations, the trajectory is discretized into a sequence of waypoints, and the constrained minimization objective is cast either as sum-of-squared distances between waypoints, or sum-of-distances between waypoints, without taking into account the mechanical limits of the robot arm (i.e., requiring a subsequent time-parameterization step). In this paper, like prior work, we discretize the trajectory into a sequence of a fixed number of waypoints. Unlike prior work, we minimize the sum-of-squared accelerations while incorporating mechanical limits and dynamics between waypoints (to produce smooth motions), and then repeat the process with fewer waypoints until we find a minimum-time trajectory. With the observation that our obstacle-avoidance problem is relatively simple compared to those addressed by prior work, we use constraints similar to those of Trajopt, though with a faster implementation using a different obstacle model.

Integrated grasp and motion planning algorithms find collision-free trajectories to grasps or grasp sets that are precomputed or synthesized during the planning process. Dragan et al. [38] use CHOMP to compute an optimized motion plan to a discretized goal set. Wang et al. [39] extend on Dragan’s approach to include learning and online synthesis of grasps that increase the goal set during optimiza-

tion. Vahrenkamp et al. [40] propose a planner (GraspRRT) that integrates finding a feasible grasp, inverse kinematics, and finding a collision-free motion plan using a sampling-based motion planner. Fontanals et al. [41] and Gravidahl et al. [42] similarly integrates RRT with grasp-biased samples to find collision-free motions to a grasp goal. Deng et al. [43] sequence learned grasp planning into trajectory generation. Pardi et al. [44] propose selecting grasps that enable a motion planner to avoid obstacles. Berenson et al. [45] propose task space regions (TSR) for constraining robot motions. GOMP integrates a variation of TSR constraints on the start and end configuration to compute an optimized motion that allows for degrees of freedom about pre-computed pick and place frames, and computing GOMP for multiple candidate grasps, allows one to find grasps that minimizes motion time while avoiding obstacles.

III. PROBLEM DEFINITION

Let \mathbf{g}^{z-} be a top-down grasp produced by grasp analysis for a parallel-jaw gripper. This grasp has a rotational degree of freedom associated with it based on the parallel-jaw gripper axis. We note that this formulation can be extended to the rotational axis of a suction gripper's contact normal. Let $\mathbf{q} \in \mathcal{C}$ be the complete specification of a robot's n degrees of freedom (e.g., the angles of all n joints in the robot), where $\mathcal{C} \subseteq \mathbb{R}^n$ is the set of all possible configurations (valid or not). For robots with joint rotation limits, let $\mathbf{q}_{\min} \in \mathbb{R}^n$ and $\mathbf{q}_{\max} \in \mathbb{R}^n$ specify those limits. Thus $\mathbf{q} \in \mathcal{C}$ implies that $\mathbf{q}_{\min} \leq \mathbf{q} \leq \mathbf{q}_{\max}$. Let $\mathcal{C}_{\text{obstacle}} \subset \mathcal{C}$ be the set of configurations that are in collision with an obstacle. Let $\mathcal{C}_{\text{free}} \subseteq \mathcal{C} \setminus \mathcal{C}_{\text{obstacle}}$ be the set of configurations that are valid. Let $p : \mathcal{C} \rightarrow \text{SE}(3)$ be the forward kinematic function for a robot. The start configuration for a motion is a grasp $\mathbf{g}_{\text{start}}^{z-} \in \text{SE}(3)$ produced by grasp analysis and corresponding to a top-down grasp. By adding a 1 degree of freedom in rotation $R_{\mathbf{a}}(\cdot)$ about axis \mathbf{a} corresponding to the vector between the two grasp contact points, we get a set of starting grasps,

$$G_{\text{start}} = \left\{ \mathbf{g}_i \mid \mathbf{g}_i = R_{\mathbf{a}}(\theta) \mathbf{g}_{\text{start}}^{z-}, \theta \in \left[-\frac{\pi}{2}, \frac{\pi}{2} \right] \right\},$$

where here we limit the rotation to $\pm \frac{\pi}{2}$. We note that the rotation limit may be set based on application needs. We define G_{goal} similarly based on the target placement frame of the grasped object, though we note that the goal may have a different axis for its rotational degree of freedom, or may have no constraints on its rotation.

Let $\mathbf{v}_{\max} \in \mathbb{R}_+^n$ be the maximum velocity of each joint, and $\mathbf{v}_{\min} = -\mathbf{v}_{\max}$ be each joints minimum velocity. Similarly, let $\mathbf{a}_{\max} \in \mathbb{R}_+^n$ be the maximum acceleration of each joint, and $\mathbf{a}_{\min} = -\mathbf{a}_{\max}$. Let $\tau \in \mathcal{T}$ be a sequence of joint configurations $\tau = (\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{\text{end}})$, and $f : \mathcal{T} \rightarrow \mathbb{R}_+$ be the time required to traverse the τ from \mathbf{q}_0 to \mathbf{q}_{end} .

Algorithm 1 Grasp-Optimized Motion Planning

Require: $\mathbf{g}_{\text{start}}^{z-}$ and $\mathbf{g}_{\text{goal}}^{z-}$ are start and goal gripper poses, H and t_{step} are the discretization parameters, \mathbf{q}_{\min} , \mathbf{q}_{\max} , \mathbf{v}_{\max} , and \mathbf{a}_{\max} are the mechanical limits, \mathbf{D} is the collision model (depth map)

- 1: $\mathbf{q}_0 \leftarrow p^{-1}(\mathbf{g}_{\text{start}}^{z-})$
- 2: $\mathbf{q}_H \leftarrow p^{-1}(\mathbf{g}_{\text{goal}}^{z-})$
- 3: $\mathbf{x}_{H+1} \leftarrow \emptyset$
- 4: $\mathbf{x}_H \leftarrow$ spline interpolation from \mathbf{q}_0 to \mathbf{q}_H
- 5: $\mathbf{P} \leftarrow$ initialize matrix
- 6: $\mathbf{A}, \mathbf{b} \leftarrow$ linearize and initialize constraints from $\mathbf{q}_{\min}, \mathbf{q}_{\max}, \mathbf{v}_{\max}, \mathbf{a}_{\max}, \mathbf{g}_{\text{start}}^{z-}, \mathbf{g}_{\text{goal}}^{z-}, \mathbf{D}$
- 7: **for** $h \leftarrow H$ down to 2 **do**
- 8: **for** $i \leftarrow 1, 2, \dots$ **do**
- 9: $\mathbf{x}_h \leftarrow \text{argmin}_{\mathbf{x}} \frac{1}{2} \mathbf{x}^T \mathbf{P} \mathbf{x}$, s.t. $\mathbf{A} \mathbf{x} \leq \mathbf{b}$
- 10: **if** QP is infeasible **or** $i > i_{\max}$ **then**
- 11: **return** \mathbf{x}_{h+1}
- 12: $\mathbf{A}, \mathbf{b} \leftarrow$ update linearization and trust region from $\mathbf{x}_h[\mathbf{q}_0], \mathbf{x}_h[\mathbf{q}_H], \mathbf{D}$
- 13: **if** constraints and trust region in tolerance **then**
- 14: **break**
- 15: $\mathbf{x}_{h-1} \leftarrow$ interpolate \mathbf{x}_h to fit shorter trajectory
- 16: $\mathbf{A}, \mathbf{b} \leftarrow$ add constraint to reflect $\mathbf{v}_h = 0$
- 17: **return** \mathbf{x}_2

The objective of GOMP is to compute:

$$\text{argmin}_{\tau} f(\tau), \quad (1)$$

$$\text{s.t. } \mathbf{q}_{\min} \leq \mathbf{q}_t \leq \mathbf{q}_{\max}, \quad \forall t \in [0, f(\tau)] \quad (2)$$

$$\mathbf{v}_{\min} \leq \mathbf{v}_t \leq \mathbf{v}_{\max}, \quad \forall t \in [0, f(\tau)] \quad (3)$$

$$\mathbf{a}_{\min} \leq \mathbf{a}_t \leq \mathbf{a}_{\max}, \quad \forall t \in [0, f(\tau)] \quad (4)$$

$$\mathbf{q}_t \in \mathcal{C}_{\text{free}}, \quad \forall t \in [0, f(\tau)] \quad (5)$$

$$p(\mathbf{q}_0) \in G_{\text{start}} \quad (6)$$

$$p(\mathbf{q}_{\text{end}}) \in G_{\text{goal}} \quad (7)$$

where \mathbf{q}_t , \mathbf{v}_t , and \mathbf{a}_t are the configuration, velocity, and accelerations at time t respectively. Thus, equations (2), (3), and (4) constrain the motion to the robot's mechanical limits, equation (5) constrains the motion to avoid obstacles, and equations (6) and (7) constrain the motion to start at a pick location, and end at a placement location.

IV. GOMP ALGORITHM

In this section we present a method for computing the constrained minimum-time trajectory of equation (1) using sequential quadratic programming (SQP). The SQP method solves a sequence of quadratic program (QP) subproblems by establishing constraints based on trust regions. We propose using QP solvers because fast implementations are readily available (e.g., [46]), and when solving a sequence of QP subproblems, each QP can be warm started using the solution to the previous subproblem.

A QP solver minimizes a quadratic objective in the form:

$$\begin{aligned} \underset{\mathbf{x}}{\operatorname{argmin}} \quad & \frac{1}{2} \mathbf{x}^T \mathbf{P} \mathbf{x} + \mathbf{p}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A} \mathbf{x} \leq \mathbf{b}, \end{aligned}$$

where $\mathbf{P} \in \mathbf{S}_+$ is a positive semidefinite matrix, \mathbf{A} , and \mathbf{b} represent the linear (and linearized) constraints. For completeness, we include \mathbf{p}^T , but set it to 0 in this paper.

To set up the QP, we first recast the problem so that τ is discretized over a finite horizon of H steps, with each step separated by a fixed time interval t_{step} . We define \mathbf{x} to match this definition as:

$$\mathbf{x} = [\mathbf{q}_0^T \quad \mathbf{q}_1^T \quad \cdots \quad \mathbf{q}_H^T \quad \mathbf{v}_0^T \quad \mathbf{v}_1^T \quad \cdots \quad \mathbf{v}_H^T]^T,$$

thus containing the configuration and velocity of each joint at each waypoint.

The recasting to a discretized fixed-interval sequence of waypoints serves several purposes: the objective is now compatible with a QP solver, the constrained dynamics can be readily specified as linear constraints, and t_{step} can be set to match the real-time cycle frequency of the robot's control loop. In section IV-E, this SQP formulation is then converted to a time-minimization.

The overall algorithm is presented in Alg. 1. The details of the algorithm are expanded on in the subsections that follow.

A. Smooth trajectory objective

After discretizing the trajectory, the next step is to set up the minimization objective for the QP. Since the discretization means the trajectory has a fixed time interval, we instead have the QP encourage smooth trajectories by minimizing a the sum-of-squared accelerations as approximated by the difference in velocities. Thus,

$$\mathbf{P} = \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{P}_v \end{bmatrix},$$

where, \mathbf{P}_v is a $(H+1) \times (H+1)$ matrix in the form:

$$\mathbf{P}_v = \begin{bmatrix} 2 & -1 & 0 & 0 & \cdots \\ -1 & 2 & -1 & 0 & \\ 0 & -1 & 2 & -1 & \\ \vdots & & \ddots & \ddots & \ddots \end{bmatrix}. \quad (8)$$

B. Dynamics constraints

To enforce that the trajectory remains within the mechanical limits of the robot, we place linear constraints on \mathbf{q}_i , \mathbf{v}_i , and \mathbf{a}_i to match the problem definition. The position and velocity constraints can be directly applied to the coefficients of \mathbf{x} . The accelerations constraint are set to

$$-\mathbf{a}_{\max} \leq \frac{1}{t_{\text{step}}} (\mathbf{v}_{i+1} - \mathbf{v}_i) \leq \mathbf{a}_{\max}.$$

To enforce that the waypoints follow the dynamics implied by the velocity, we apply linear constraints using a common model-predictive control technique. These constraints take the form:

$$\mathbf{q}_{i+1} = \mathbf{q}_i + \mathbf{v}_i t_{\text{step}}.$$

With these constraints in place, the QP will enforce that the trajectory follows the dynamics implied by each waypoint while not exceeding the robot's mechanical limits.

Finally, to insure that the trajectory starts and ends with zero velocity, we add the constraints:

$$\mathbf{v}_0 = 0 \quad \mathbf{v}_H = 0$$

C. Obstacles

To enable obstacle avoidance, we add a linearization of the constraints based on the robot's forward kinematics and the obstacle. The method we apply follows directly from that of Trajopt [31]; we establish trust regions as linearized constraints in a QP. In implementation, we observe that pick-and-place robots can operate with a simplified task-specific obstacle model—specifically a depth map. As such, we can avoid the complexities associated with the GJK [47] and EPA [48] methods for collision detection and instead incorporate a constraint based on each waypoint's position over the depth map. We observe that an interior point optimization using boundary functions [35] would also serve to avoid obstacles.

Assuming we define the positive z -axis as the up vector, obstacle avoidance constraint takes the form:

$$z_{\text{obstacle}} - p(\mathbf{q}_i^{(k)}) + \mathbf{J}_z^{(k)} \mathbf{q}_i^{(k)} \leq \mathbf{J}_z^{(k)} \mathbf{q}_i^{(k+1)},$$

where, z_{obstacle} is the height of the obstacle that waypoint i must be above, $\mathbf{J}_z^{(k)}$ is the z -axis translation row of the Jacobian of robot arm after SQP iteration k .

D. Start and goal constraints

The pick-and-place operation has a set of start grasps G_{start} and a set of goal positions G_{goal} for the gripper. To enforce the constraint that the trajectory starts with $p(\mathbf{q}_0) \in G_{\text{start}}$ and ends at $p(\mathbf{q}_H) \in G_{\text{goal}}$, we start with an initial QP constraint of having the $\mathbf{q}_0 = p^{-1}(\mathbf{g}_{\text{start}}^{z-})$ and $\mathbf{q}_H = p^{-1}(\mathbf{g}_{\text{goal}}^{z-})$, where $p^{-1}(\cdot)$ produces an inverse kinematic solution. As described in the problem definition, these positions have at least one degree of freedom in the form of a rotation about the axis of the grasp.

The inequality constraint on the start configuration \mathbf{q}_0 that allows the rotation about a free axis is:

$$\mathbf{b}_0^{(k+1)} - \epsilon_0 \leq R_0 \mathbf{J}^{(k)} \mathbf{q}_0^{(k+1)} \leq \mathbf{b}_0^{(k+1)} + \epsilon_0,$$

where R_0 rotates the coordinate frame so that a single component of the Jacobian corresponds to the degree of freedom. Thus, ϵ_0 is a vector in which coefficient corresponding to that degree of freedom is large, and the remaining values in the vector are small. The bounds of the inequality for the $k+1$ -th QP is based on the solution to the k -th QP,

$$\mathbf{b}_0^{(k+1)} = R_0 \left(\mathbf{g}_{\text{start}}^{z-} - p(\mathbf{q}_0^{(k)}) + \mathbf{J}^{(k)} \mathbf{q}_0^{(k)} \right).$$

The constraint on the goal configuration follows a similar formulation as the bounds on the start configuration, with potentially a different set of degrees of freedom.




Fig. 3. In experimental setup, we take a top-down image of the objects to grasp (left) using a high-resolution depth camera. We then feed the image through a modified version of Dex-Net 4.0, which produces a diverse set of 28 candidate grasps (right). From these grasps, we then compute pick-to-place trajectories.

E. Time minimization

After solving the SQP with the constraints specified as above, \mathbf{x} will correspond to a trajectory that satisfies dynamic constraints, avoids obstacles, and minimizes the sum-of-squared accelerations between waypoints. To convert this formulation into a time minimization, we next repeatedly solve the SQP, reducing H until the SQP is infeasible [49], and taking the \mathbf{x} associated with the smallest value of H for which the SQP is feasible. By finding the minimum H we explicitly minimize the number of discrete time steps, and thus the overall time for the trajectory.

To implement the reduced horizon H , we convert the velocity constraint on the last waypoint from the inequality bound by \mathbf{v}_{\max} to the equality constraint $\mathbf{v}_H = 0$.

We note that this process can be stopped at any time after the solution with the first feasible H , in which case the shortest time solution found so far would be usable as a trajectory for the robot.

F. Warm starting

The convergence time of QP solvers can benefit from providing a good initial estimate of the solution or *warm starting* the solver [46]. To speed up the SQP, we warm start several phases of the optimization process.

We warm start the first QP with a spline that interpolates between \mathbf{q}_0 and \mathbf{q}_H , and starting and stopping with zero velocity without regarding obstacles.

When reducing H to find a minimum-time feasible trajectory, the solution from the previous optimization does not align with the optimization of the subsequent optimization. As such we interpolate the solution from the previous H to smaller set of waypoints.

V. RESULTS

To test the performance and compatibility of the trajectories GOMP computes, we apply them to a physical setup resembling what one may find at a pick-and-place station in a warehouse. In this setup, a high-resolution depth camera produces an image of the objects to grasp (Fig. 3 (a)). This image is then sent to our modified version of Dex-Net 4.0 to generate a diverse set of grasp candidate poses (Fig. 3 (b)).

TABLE I
TIMING COMPARISON FOR 28 PICK-TO-PLACE MOTIONS

	baseline	GOMP	speedup
mean	5.042 s	0.544 s	9.2×
stdev	0.440 s	0.172 s	
min	4.144 s	0.256 s	16.2×
max	5.624 s	0.976 s	5.8×

TABLE II
PER-OBJECT GOMP MOTION TIMES

object (color)	grasps	minimum	mean	maximum
Car (orange)	5	0.400 s	0.659 s	0.736 s
Castle (blue, right)	2	0.432 s	0.464 s	0.496 s
Clamp (dark blue)	2	0.480 s	0.512 s	0.544 s
Mug (black)	8	0.382 s	0.578 s	0.976 s
Nozzle (blue, left)	2	0.400 s	0.440 s	0.480 s
Pipe connector (purple)	3	0.384 s	0.501 s	0.384 s
Turbine (green)	6	0.256 s	0.496 s	0.896 s

From the grasp candidates, we compute a trajectory for the UR5 robot that takes the robot from pick-point to placement, while avoiding obstacles including the side of the placement bin. This physical setup, along with the execution of an example trajectory, is shown in Fig. 4.

We modify the original Dex-Net 4.0 to produce a diverse set of grasp candidates for GOMP by changing the sampling method from argmax to rejection-sampling with a minimum distance constraint between grasps in image space.

In the physical experiment, each grasp is initially computed as a top-down grasp. We set the SQP constraints to allow rotation about the grasp point by up to 45 degrees in either direction. In Fig. 4 (a), the SQP computed an initial grasp that aligns to the computed grasp points, but is rotated to allow for a faster motion from pick-point to placement.

At any point in time after the initial computation of a trajectory, the robot has a valid trajectory. With additional compute time, successive reduction of H allows for a faster trajectory. An example of this process of computing faster trajectories is shown in Fig. 5. In this figure, we can see an initial smooth trajectory that is successively made faster until it cannot be shortened any further due to mechanical limits of the robot.

We compare the motion execution time for the 28 grasps shown in Fig. 4, to a baseline trajectory that uses top-down grasps and performs 3 steps: lift to safe height, move over bin, then lower to placement height. The baseline motion does not incorporate obstacle avoidance, instead, it avoids obstacles by lifting the object sufficiently high to avoid foreseeable collisions in the workspace. This baseline is run at the UR5’s default speed. The comparison between the baseline and the optimized motions are shown in table I. From the table, we observe that GOMP performs faster in all tested cases, and on average provides speedup that is 9.2 faster than the baseline. These savings can be passed on to the complete grasping pipeline to speed up the pick and place process.

We also test the ability of the trajectory optimization to be integrated with grasp selection. In Fig. 3, each object has multiple candidate grasp points. For each object, we




Fig. 4. UR5 robot moves quickly to transport picked object to placement over bin. The UR5 moves within its joint limits, and with both pick and placement orientation optimized as part of the motion planning process. This motion takes 0.43 seconds.

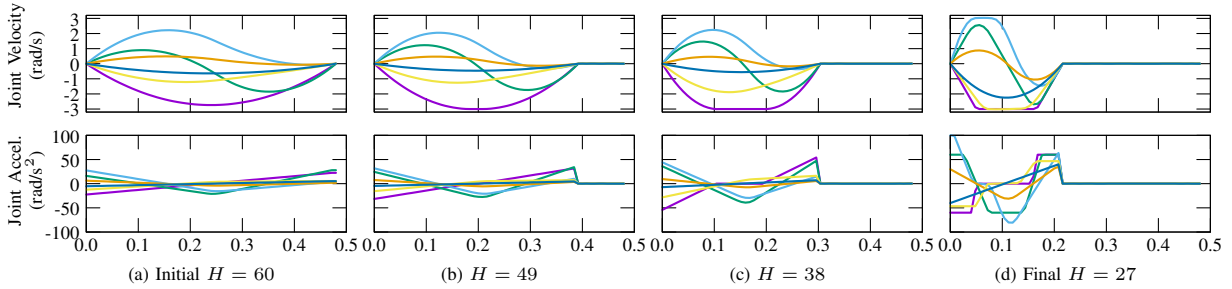


Fig. 5. **Joint velocities and accelerations computed during time optimization of the trajectory.** In these plots, the velocity and acceleration of each of the 6 joints on the UR5 robot is plotted in a separate color. Each column represents the result of trajectory optimization for a successively smaller value of H . The time optimization of the trajectory starts with a value of H and subsequently shrinks this value by constraining the velocity to zero at the end of the trajectory. Trajectory optimization stops once the QP is infeasible. In this case, the shortest time is $H = 27$, which, at $t_{step} = 0.008$, results in a trajectory that requires 0.216 seconds to execute.

compute trajectories to all of its candidate grasp points in parallel, and then select the grasp that results in the shortest execution time. In table II we observe that selecting the grasp that produces the shortest motion can sometime produce an additional performance benefit.

The bin-picking pipeline in these experiments requires the following steps: (a) imaging (1.200 ± 0.051 s), (b) grasp analysis (0.529 ± 0.041 s) (c) motion to grasp, (d) gripper closing (0.280 ± 0.140 s), (e) motion to placement, and (f) gripper opening (0.280 ± 0.140 s). Plugging the mean values for baseline motions and GOMP motions from table I, the mean total pick time for baseline motions is 12.373 s, and with GOMP motions is 3.377 s. This suggests that GOMP can speed up our pipeline’s mean total pick time by $3.6\times$. We caveat that this PPH has not been confirmed with a fully integrated system.

VI. CONCLUSION

We presented a time-optimizing motion planner that speeds up pick-and-place operation by incorporating joint limits, obstacle avoidance, and allowing a degree of freedom to be incorporated into grasping and placement positions. By first converting the problem from a time-minimization problem to minimization that encourages smooth velocities and linearizing constraints, we are able to recast the problem to a sequential quadratic program (SQP). After solving the SQP, we then repeatedly shorten the trajectory’s time horizon in the SQP—stopping at any time, or when the SQP is detected as infeasible, resulting in a fast motion plan. When comparing to a baseline motion, the optimized motion plan

runs significantly faster, creating an opportunity to speed up the entire pick-and-place operation.

In future work, we plan to test this motion planner on different robots with differing degrees of freedom, as well as integrate it into the Dex-Net pipeline. We also see opportunities to extend this work to account for additional constraints and trajectory objectives that one might encounter when deploying this to different robots and environments—including, for example, generating minimum-jerk trajectories, and avoiding dynamic obstacles.

ACKNOWLEDGMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, Berkeley Deep Drive (BDD), the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS “People and Robots” (CPAR) Initiative. Authors were also supported by the Scalable Collaborative Human-Robot Learning (SCHool) Project, a NSF National Robotics Initiative Award 1734633, and in part by donations from Google and Toyota Research Institute. We thank our colleagues who provided helpful feedback and suggestions, in particular Ashwin Balakrishna and Brijen Thananjeyan. This article solely reflects the opinions and conclusions of its authors and do not reflect the views of the sponsors or their associated entities.

REFERENCES

- [1] D. Morrison, P. Corke, and J. Leitner, “Learning robust, real-time, reactive robotic grasping,” *The International Journal of Robotics Research*, p. 0278364919859066, 2019.
- [2] V. Satish, J. Mahler, and K. Goldberg, “On-policy dataset synthesis for learning robot grasping policies using fully convolutional deep networks,” *IEEE Robotics & Automation Letters*, vol. 4, no. 2, pp. 1357–1364, 2019.
- [3] U. Robotics. UR5 collaborative robot arm. [Online]. Available: <https://web.archive.org/web/20190815054522/https://www.universal-robots.com/products/ur5-robot/>

- [4] Robotiq. 2F-85 and 2F-140 grippers. [Online]. Available: <https://web.archive.org/web/20190519030456/https://robotiq.com/products/2f85-140-adaptive-robot-gripper>
- [5] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, 2019.
- [6] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, "Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1957–1964.
- [7] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. Aparicio, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Proc. Robotics: Science and Systems (RSS)*, 2017.
- [8] J. Mahler and K. Goldberg, "Learning deep policies for robot bin picking by simulating robust grasping sequences," in *Conf. on Robot Learning (CoRL)*, 2017, pp. 515–524.
- [9] A. ten Pas, M. Gualtieri, K. Saenko, and R. Platt, "Grasp pose detection in point clouds," *Int. Journal of Robotics Research (IJRR)*, vol. 36, no. 13-14, pp. 1455–1473, 2017.
- [10] D. Kappler, J. Bohg, and S. Schaal, "Leveraging big data for grasp planning," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2015, pp. 4304–4311.
- [11] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *Int. Journal of Robotics Research (IJRR)*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [12] E. Johns, S. Leutenegger, and A. J. Davison, "Deep learning a grasp function for grasping under gripper pose uncertainty," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 4461–4468.
- [13] U. Viereck, A. t. Pas, K. Saenko, and R. Platt, "Learning a visuomotor controller for real world robotic grasping using simulated depth images," in *Conf. on Robot Learning (CoRL)*, 2017.
- [14] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *Conf. on Robot Learning (CoRL)*, 2018, pp. 651–673.
- [15] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3406–3413.
- [16] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning hand-eye coordination for robotic grasping with large-scale data collection." Springer, 2016, pp. 173–184.
- [17] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, et al., "Using simulation and domain adaptation to improve efficiency of deep robotic grasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4243–4250.
- [18] A. Mousavian, C. Eppner, and D. Fox, "6-dof grasplet: Variational grasp generation for object manipulation," in *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2019, pp. 2901–2910.
- [19] A. Murali, A. Mousavian, C. Eppner, C. Paxton, and D. Fox, "6-dof grasping for target-driven object manipulation in clutter," *arXiv preprint arXiv:1912.03628*, 2019.
- [20] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee, "Learning 6-dof grasping interaction via deep geometry-aware 3d representations," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–9.
- [21] M. Liu, Z. Pan, K. Xu, K. Ganguly, and D. Manocha, "Generating grasp poses for a high-dof gripper using neural networks," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [22] C. Eppner, A. Mousavian, and D. Fox, "A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set," in *Int. S. Robotics Research (ISRR)*, 2019.
- [23] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high dimensional configuration spaces," *IEEE Trans. Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [24] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. Journal of Robotics Research (IJRR)*, vol. 30, no. 7, pp. 846–894, June 2011.
- [25] J. Kuffner and S. LaValle, "An efficient approach to path planning using balanced bidirectional rrt search," *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep.*, 2005.
- [26] N. M. Amato and L. K. Dale, "Probabilistic roadmap methods are embarrassingly parallel," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, May 1999, pp. 688–694.
- [27] J. Ichnowski and R. Alterovitz, "Scalable multicore motion planning using lock-free concurrency," *IEEE Trans. Robotics*, vol. 30, no. 5, pp. 1123–1136, 2014.
- [28] J. Pan, L. Zhang, and D. Manocha, "Collision-free and smooth trajectory computation in cluttered environments," *The International Journal of Robotics Research*, vol. 31, no. 10, pp. 1155–1175, 2012.
- [29] Q.-C. Pham, "A general, fast, and robust implementation of the time-optimal path parameterization algorithm," *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1533–1540, 2014.
- [30] T. Kunz and M. Stilman, "Time-optimal trajectory generation for path following with bounded acceleration and velocity," *Robotics: Science and Systems VIII*, pp. 1–8, 2012.
- [31] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Robotics: science and systems*, vol. 9, no. 1. Citeseer, 2013, pp. 1–10.
- [32] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," 2009.
- [33] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 4569–4574.
- [34] C. Park, J. Pan, and D. Manocha, "ITOMP: Incremental trajectory optimization for real-time replanning in dynamic environments," in *Twenty-Second International Conference on Automated Planning and Scheduling*, 2012.
- [35] A. Kuntz, C. Bowen, and R. Alterovitz, "Fast anytime motion planning in point clouds by interleaving sampling and interior point optimization," *ISRR*, 2017, 2017.
- [36] M. Campana, F. Lamiroux, and J.-P. Laumond, "A gradient-based path optimization method for motion planning," *Advanced Robotics*, vol. 30, no. 17-18, pp. 1126–1144, 2016.
- [37] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuous-time gaussian process motion planning via probabilistic inference," *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.
- [38] A. D. Dragan, N. D. Ratliff, and S. S. Srinivasa, "Manipulation planning with goal sets using constrained trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 4582–4588.
- [39] L. Wang, Y. Xiang, and D. Fox, "Manipulation trajectory optimization with online grasp synthesis and selection," *arXiv preprint arXiv:1911.10280*, 2019.
- [40] N. Vahrenkamp, M. Do, T. Asfour, and R. Dillmann, "Integrated grasp and motion planning," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2883–2888.
- [41] J. Fontanals, B.-A. Dang-Vu, O. Porges, J. Rosell, and M. A. Roa, "Integrated grasp and motion planning using independent contact regions," in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 887–893.
- [42] I. Gravidahl, K. Seel, and E. I. Grøtli, "Robotic bin-picking under geometric end-effector constraints: Bin placement and grasp selection," in *2019 7th International Conference on Control, Mechatronics and Automation (ICCMA)*. IEEE, 2019, pp. 197–203.
- [43] Z. Deng, X. Zheng, L. Zhang, and J. Zhang, "A learning framework for semantic reach-to-grasp tasks integrating machine learning and optimization," *Robotics and Autonomous Systems*, vol. 108, pp. 140–152, 2018.
- [44] T. Pardi, R. Stolkin, et al., "Choosing grasps to enable collision-free post-grasp manipulations," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2018, pp. 299–305.
- [45] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [46] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *ArXiv e-prints*, Nov. 2017.
- [47] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional

space,” *IEEE Journal on Robotics and Automation*, vol. 4, no. 2, pp. 193–203, 1988.

- [48] S. Cameron and R. Culley, “Determining the minimum translational distance between two convex polyhedra,” in *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, vol. 3. IEEE, 1986, pp. 591–596.
- [49] G. Banjac, P. Goulart, B. Stellato, and S. Boyd, “Infeasibility detection in the alternating direction method of multipliers for convex optimization,” *Journal of Optimization Theory and Applications*, vol. 183, no. 2, pp. 490–519, 2019. [Online]. Available: <https://doi.org/10.1007/s10957-019-01575-y>