

Simultaneous Trajectory Optimization and Contact Selection for Contact-rich Manipulation with High-Fidelity Geometry

Zhang, Mengchao; Jha, Devesh K.; Raghunathan, Arvind; Hauser, Kris

TR2025-047 April 17, 2025

Abstract

Contact-implicit trajectory optimization (CITO) is an effective method to plan complex trajectories for various contact-rich systems including manipulation and locomotion. CITO formulates a mathematical program with complementarity constraints (MPCC) that enforces that contact forces must be zero when points are not in contact. However, MPCC solve times increase steeply with the number of allowable points of contact, which limits CITO's applicability to problems in which only a few, simple geometries are allowed to make contact. This paper introduces simultaneous trajectory optimization and contact selection (STOCS), as an extension of CITO that overcomes this limitation. The innovation of STOCS is to identify salient contact points and times inside the iterative trajectory optimization process. This effectively reduces the number of variables and constraints in each MPCC invocation. The STOCS framework, instantiated with key contact identification subroutines, renders the optimization of manipulation trajectories computationally tractable even for high-fidelity geometries consisting of tens of thousands of vertices.

IEEE Transactions on Robotics (T-RO) 2025

Simultaneous Trajectory Optimization and Contact Selection for Contact-rich Manipulation with High-Fidelity Geometry

Mengchao Zhang  *Member, IEEE*, Devesh K. Jha  *Senior Member, IEEE*, Arvind U. Raghunathan ,
Kris Hauser  *Senior Member, IEEE*

Abstract—Contact-implicit trajectory optimization (CITO) is an effective method to plan complex trajectories for various contact-rich systems including manipulation and locomotion. CITO formulates a mathematical program with complementarity constraints (MPCC) that enforces that contact forces must be zero when points are not in contact. However, MPCC solve times increase steeply with the number of allowable points of contact, which limits CITO’s applicability to problems in which only a few, simple geometries are allowed to make contact. This paper introduces simultaneous trajectory optimization and contact selection (STOCS), as an extension of CITO that overcomes this limitation. The innovation of STOCS is to identify salient contact points and times inside the iterative trajectory optimization process. This effectively reduces the number of variables and constraints in each MPCC invocation. The STOCS framework, instantiated with key contact identification subroutines, renders the optimization of manipulation trajectories computationally tractable even for high-fidelity geometries consisting of tens of thousands of vertices.

Index Terms—Manipulation planning, trajectory optimization, infinite programming

I. INTRODUCTION

HUMANS and other organisms treat contact as a fact of life and utilize contact to perform dexterous manipulation of objects and agile locomotion. In contrast, the majority of current robots avoid making contact with objects as much as possible, and tend to avoid contact-rich manipulations like pushing, sliding, and rolling [1], [2]. Trajectory optimization [3] has been investigated as a tool for generating high quality manipulations, but choosing an effective mathematical representation of making and breaking contact remains a major research challenge. Two general classes of methods are available: hybrid trajectory optimization and contact-implicit trajectory optimization (CITO). Hybrid trajectory optimization divides a trajectory into segments in which the set of contacts remains constant, but it requires the contact mode sequence to be known in advance [4] or explored by an auxiliary discrete search. CITO [5], [6], [7], [8], [9] allows the optimizer to

choose the sequence of contact within the optimization loop. CITO formulates contact as a complementarity constraint to ensure that the contact forces can be non-zero if and only if a point is in contact [5]. Although the resulting mathematical programming with complementary constraint (MPCC) [10] formulation is less restrictive than hybrid trajectory optimization, it still requires a set of predefined allowable contact points on the object. Moreover, MPCC rapidly becomes more challenging to solve as the number of complementarity constraints increases, so past CITO applications were limited to a small handful of potential contact points.

This paper introduces the simultaneous trajectory optimization and contact selection (STOCS) algorithm to address the scaling problem in contact-implicit trajectory optimization. To the best of our knowledge, this is the first method capable of optimizing contact-rich manipulation trajectories with high-fidelity geometric representations in 3D. This method paves the way for manipulation planning with raw sensor input, such as point clouds derived from RGBD images, and eliminates the need for geometry simplification.

Instead of predefining a fixed set of possible contact points as in typical CITO approaches, the STOCS algorithm employs an infinite programming (IP) approach, functioning similarly to the active set method [11]. This approach dynamically generates potential contact points and contact times between the object and environment, making the resulting MPCCs significantly more tractable to solve. The primary contribution of this work is to extend prior work on IP for robot pose optimization [12] to the trajectory optimization setting. This paper describes methods for encoding quasi-dynamic constraints in trajectory optimization, and also presents a novel method for selecting salient contact points and contact times. This method, Time-active Maximum Violation Oracle (TAMVO) with spatial disturbance and temporal smoothing, encourages the IP framework to converge quickly toward a feasible solution. We demonstrate the effectiveness and efficiency of STOCS in solving 2D and 3D sliding, pivoting, and peg in hole tasks with irregular objects and environments, whose models can include up to tens of thousands of vertices. Without STOCS, CITO methods take hundreds to thousands of times longer, even for problems of moderate size (a few hundred vertices).

II. RELATED WORK

Model-based trajectory planning methods have been extensively studied for contact-rich manipulation problems. The

Mengchao Zhang is with the Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA. (e-mail: mz17@illinois.edu)

Devesh K. Jha and Arvind U. Raghunathan are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, 02139, USA. (e-mail: {jha,raghunathan}@merl.com)

Kris Hauser is with the School of Computing and Data Science, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA. (e-mail: kkhhauser@illinois.edu)

Code is available at: <https://github.com/zmccmzty/STOCS-3D>.




Fig. 1. STOCS accepts as input the high-fidelity geometry of the object (represented by a dense point cloud) and the environment (represented by a signed distance field), the robot’s contact point, and start and goal poses of the object (left). The STOCS algorithm first generates an initial trajectory by linearly interpolating between the start and goal poses, and then it iterates between selecting contact points and solving a finite-dimensional MPCC to decide a step direction until the convergence criteria are met (center). As output (right), STOCS produces the pose of the object, active object-environment contact points (green dots) and forces (green lines), and manipulation force (red lines). Nonpenetration, Coulomb friction, complementarity, and quasi-dynamic stability are enforced throughout the trajectory. [Best viewed in color.]

presence of contact presents significant challenges for optimization, primarily due to the stiff and non-smooth nature of contact dynamics [13]. Various approaches have been proposed to address this challenge.

Hybrid trajectory optimization approaches divide a trajectory into segments in which the contact mode (set of active contacts) remains constant [14]. This segmentation allows trajectory optimization to be cast as a large non-linear program that can be solved to optimize the timings and variables associated with each of the individual modes [15], [16]. However, manually determining an appropriate sequence of contact modes is intractable in all but the simplest problems.

An alternative approach assumes a predefined object trajectory, which prescribes the contact mode, and then solves for the robot’s contact interactions needed to achieve this desired trajectory [17]. However, obtaining a feasible object trajectory itself can be challenging.

Rather than predefining contact mode sequences, *mode search* methods [18], [19] automatically enumerate contact modes [20] to guide the tree expansion within a sampling-based motion planning algorithm. Although this overcomes the limitations of mode sequencing, the complexity of enumerating all 3D contact modes for one object is $\mathcal{O}(N^d)$, where N represents the number of contacts and d denotes the object’s effective degrees of freedom, which is 3 in 2D and 6 in 3D. This superlinear growth in N makes these methods impractical for high-fidelity geometric representations.

Contact-implicit trajectory optimization (CITO) offers an alternative to predefining or searching over the contact mode sequence [5], [6], [7]. CITO models all possible contacts with

complementarity constraints and casts trajectory optimization as a mathematical program with complementarity constraints (MPCC). These constraints allow forces to be nonzero if and only if the distance between a contact and the opposing object is zero. Given a set of potential allowable contacts, the optimizer loosens the complementarity constraints to let forces be applied at a distance, and then progressively tightens them. Through smoothing the contact dynamics using the above strategy, CITO can then simultaneously optimize the mode sequence as well as the contact forces. However, CITO becomes extremely challenging to solve when there are a large number of possible contacts, leading to large numbers of complementarity constraints. Consequently, the main limitation of this approach is that the set of allowable contacts must be predefined and needs to be relatively small.

The STOCS method introduced in this paper addresses the scaling problem of CITO by incorporating the identification of salient contact points and contact times inside trajectory optimization. This approach effectively reduces the number of variables and constraints in the resulting MPCC, rendering the computation of contact-rich manipulation trajectories for objects with complex, non-convex geometries computationally tractable. In prior work, we introduced the infinite programming framework used here, and applied it to optimizing stable grasping poses [12]. In contrast, this paper extends the framework to optimize entire trajectories for contact-rich manipulations. Although STOCS is contact-implicit for object-environment contact, it still requires a pre-defined robot-object contact state. To enable changes in the robot-object contact, a sampling-based planner that incorporates robot-object contact

sampling could be integrated with STOCS. It is also worth noting that STOCS is a local trajectory optimization method, so its output is sensitive to the specified temporal resolution and initial trajectory.

This article is an expanded version of a conference paper [21]. Our extensions broaden the applicability of STOCS and enables its use with more complex objects and environments. Compared to the conference version, which only worked in planar problems, the current implementation is far more efficient and we demonstrate how it can be applied to 3D problems. Moreover, we introduce a novel method for selecting contact points and contact times, called the Time-active Maximum Violation Oracle (TAMVO), that also greatly improves computational performance. This version also adds additional details, references, and experiments.

III. APPROACH

Given a start pose and a goal pose, a trajectory that includes the object's motion and the control inputs of the manipulator needs to be planned. In this section, we describe the inputs and outputs of our method in detail, and explain how we use STOCS to solve contact-rich manipulation trajectory optimization problems.

A. Problem Description

Our method requires the following information as inputs:

- 1) Object initial pose region: $Q_{init} \subset SE(2)$ or $SE(3)$.
- 2) Object goal pose region: $Q_{goal} \subset SE(2)$ or $SE(3)$.
- 3) Object properties: a rigid body \mathcal{O} whose geometry, mass distribution, and friction coefficients with both the environment μ_{env} and the manipulator μ_{mnp} are known.
- 4) Environment properties: rigid environment \mathcal{E} whose geometry is known.
- 5) Robot's contact point(s) with the object: c^{mnp} .
- 6) A time step Δt and number of time steps T in the trajectory.

Our method will output a trajectory τ that includes the following information at time t :

- 1) Object's configuration: q_t .
- 2) Object's velocity (angular and translational): v_t .
- 3) Robot's contact point(s): c_t^{mnp} .
- 4) Manipulation force: u_t .
- 5) Object's contact points with the environment: \tilde{Y}_t .
- 6) Contact force at each object-environment contact point: $z(y_t) \forall y_t \in \tilde{Y}_t$.

In this paper, we treat all objects and environments as rigid bodies and we assume the contact between the manipulator and the object is sticking contact. Furthermore, users are provided with the flexibility to opt for either quasistatic or quasidynamic assumptions based on their specific requirements. Under the quasistatic paradigm, inertial forces are considered negligible, necessitating that the object remains in a state of force and torque equilibrium at all times. Quasidynamic manipulation accounts for scenarios where tasks may involve occasional dynamic periods, during which accelerations do not integrate into substantial velocities, and both momentum and the effects of impact restitution are negligible [22].

Algorithm 1 STOCS

Require: $q_{start}, q_{goal}, c^{mnp}$

- 1: $\tilde{Y}^0 = []$ ▷ Initialize empty constraint set
- 2: $z_0 \leftarrow \emptyset$ ▷ Initialize empty force vector
- 3: $x_0 \leftarrow \text{initialize trajectory}(q_{start}, q_{goal}, c^{mnp})$
- 4: **for** $k = 1, \dots, N^{max}$ **do**
- 5: ▷ Update constraint set and guessed forces z_k
- 6: Add all points in \tilde{Y}^{k-1} to \tilde{Y}^k , and initialize their forces in z_k with the corresponding values in z_{k-1}
- 7: Call Oracle to add new points to \tilde{Y}^k , and initialize their corresponding forces in z_k
- 8: $x_k \leftarrow x_{k-1}$
- 9: ▷ Solve for step direction
- 10: Set up inner optimization $P^k = P(\tilde{Y}^k)$
- 11: Run S steps of an NLP solver on P^k , starting from x_k, z_k
- 12: Set x^*, z^* to its solution, and $\Delta x \leftarrow x^* - x_k, \Delta z \leftarrow z^* - z_k$
- 13: Do backtracking line search with at most N_{LS}^{max} steps to find optimal step size α such that $\phi(x_k + \alpha\Delta x, z_k + \alpha\Delta z; \mu) \leq \phi(x_k, z_k; \mu)$
- 14: ▷ Update state and test for convergence
- 15: $x_k \leftarrow x_k + \alpha\Delta x, z_k \leftarrow z_k + \alpha\Delta z$
- 16: **if** Convergence condition is met **then**
 return x_k, z_k
- 17: **return** NOT CONVERGED

B. STOCS Trajectory Optimizer

Overall, STOCS formulates contact-rich trajectory optimization as an *infinite program* (IP), which is a constrained optimization with a potentially infinite set of variables and constraints. It uses an *exchange method* to solve the IP, by wrapping a contact selection outer loop around a finite optimization problem. The inner problem formulates CITO with complementarity constraints and solves an MPCC. In summary, the algorithm operates as follows:

- 1) Initialize the initial guess object trajectory, e.g., a linear interpolation.
- 2) Initialize an empty *candidate set* of contact points $\tilde{Y}_1, \dots, \tilde{Y}_T$ for each time step and an empty set of contact forces.
- 3) Use an *Oracle* to identify new / removed contact points, and update $\tilde{Y}_1, \dots, \tilde{Y}_T$.
- 4) Solve an MPCC for the candidate set of contact points using a small number of iterations.
- 5) Check for convergence. If not converged, repeat from step 3.

Contact forces for existing contact points are maintained from iteration to iteration to warm-start the next MPCC. Details about the algorithm (Alg. 1) and oracle design are described below.

Infinite programming for contact-rich trajectory optimization. First, we describe the IP formulation used by STOCS. Let us start by defining *semi-infinite programming* (SIP). An SIP problem is an optimization problem in finitely many

variables $q \in \mathbb{R}^n$ on a feasible set described by infinitely many constraints:

$$\min_{q \in \mathbb{R}^n} f(q) \quad (1a)$$

$$\text{s.t. } g(q, y) \geq 0 \quad \forall y \in Y \quad (1b)$$

where $g(q, y) \in \mathbb{R}^m$ is the constraint function, y denotes the *index parameter*, and $Y \in \mathbb{R}^p$ is the *index domain* which can be continuously infinite.

As an example, to solve pose optimization with non-penetration constraints [23], q describes the pose of the object, and y is a point on the surface of the object \mathcal{O} , where $Y \equiv \partial\mathcal{O}$ denotes the surface of \mathcal{O} . The constraint $g(\cdot, \cdot)$ is the signed distance from a point to the environment. Although the inequality $g(q, y) \geq 0$ must be satisfied for all y , optimal solutions will be supported by some points, i.e., at the optimal solution q^* , the Karush-Kuhn-Tucker conditions will be established by a finite set of points y such that $g(q^*, y) = 0$.

In trajectory optimization, constraints need to be enforced in space-time [23], so we define a candidate index point $y_t \in \partial\mathcal{O}$ as being indexed by time. We denote $Y_t \equiv \partial\mathcal{O}$ as the index domain at time t . To enhance the stability of the optimization process, we assume that once an index point is instantiated, it remains active throughout the whole solving process. The removal of index points will be considered in future work.

STOCS then reasons about contact forces at each index point, and since the contact force distribution is a function over the object surface, this introduces infinitely many variables, which turns the problem into an infinite program [24]. Specifically, we define the contact force $z(\cdot) : Y_t \rightarrow \mathbb{R}^r$ as an optimization variable. Our formulation imposes friction constraints, complementarity constraints, and quasi-dynamic/quasi-static force and moment balance on these variables.

To formulate the force variables and friction cone, we encode normal and tangential forces separately in the vector z . In 2D, $z(y) = [z^N, z^+, z^-]$ is expressed in a reference frame with z^N the component of force normal to the contact surface, and z^+ , z^- tangent to the contact surface. In 3D, following [5], we use a polyhedral approximation of the friction cone [25]. We express $z(y) = [z^N, z^{D1}, \dots, z^{Dd}]$ in a reference frame with z^N normal to the contact surface, and z^{D1}, \dots, z^{Dd} tangent to the contact surface. The convex hull of the unit vectors along the directions of z^{D1}, \dots, z^{Dd} in \mathbb{R}^2 forms the polyhedral approximation. The force is required to satisfy $z \geq 0$ and $\mu z^N - \sum_{i=D_1}^{D_d} z^i \geq 0$ where μ is the friction coefficient of the given point and $\sum_{i=D_1}^{D_d} z^i$ is the sum of tangential forces.

In summary, the constraints we impose at time t are as follows:

1. State and Control Bounds:

$$q_t \in \mathcal{Q}, v_t \in \mathcal{V}, u_t \in \mathcal{U} \quad (2)$$

2. Object Dynamics:

$$q_t - q_{t+1} + v_{t+1} \Delta t = 0 \quad (3)$$

3. Distance Complementarity:

ensures that nonzero forces are only exerted at points where objects are in contact,

i.e. the normal force $z^N(y)$ is nonzero only if contact is made at y .

$$0 \leq z^N(y) \perp g(q_t, y) \geq 0 \quad \forall y \in Y_t \quad (4)$$

4. Unilateral Manipulation Velocity:

ensures the manipulator can only push the object rather than pull.

$$c(q_t, u_t) \geq 0 \quad (5)$$

Here, $c(q_t, u_t)$ represents the normal force exerted by the manipulator, and this constraint ensures that the force points inward, toward the object.

5. Force-torque Balance:

$$\underbrace{s_{q,u}(q_t, u_t) + \int_{y \in Y_t} s_z(q_t, y, z(y)) dy}_{=: s(q_t, u_t, z; Y_t)} = M \dot{v}_t \text{ or } 0, \quad (6)$$

where $s_{q,u}(q_t, u_t)$ represents the wrench applied at the object's center of mass by gravity and the manipulation force, while $s_z(q_t, y, z(y))$ denotes the wrench applied at the object's center of mass by the contact force $z(y)$.

6. Friction constraints:

the force constraints defined above are grouped into a function h

$$h(q_t, y, z(y)) \geq 0 \quad \forall y \in Y_t. \quad (7)$$

Function h is related to the friction cone constraint, and the linearized friction cone constraint used in [25] can serve as a specific instantiation of it. Details on the formulation of h for both 2D and 3D scenarios are provided separately in the Appendix. A similar constraint is applied to the manipulator contact force.

7. Friction-Velocity Complementarity:

ensures the relative tangential velocity at a contact is zero unless the contact force is at the boundary of the friction cone

$$0 \leq w(q_t, v_t, y) \perp h(q_t, y, z(y)) \geq 0 \quad \forall y \in Y_t. \quad (8)$$

Function w is related to the relative tangential velocity at a contact. Details on the formulation of w for both 2D and 3D scenarios are provided separately in the Appendix.

We elaborate a bit on (6), which integrates the force distribution over the domain Y_t . Here, $s_{q,u}(q, u)$ represents the force and torque applied by gravity and the manipulator, and $s_z(q, y, z(y))$ represents the force and torque applied on an index point y by the contact force $z(y)$. The integral gives the net force and torque experienced by the object, which is 0 if we assume quasi-static, and is $M \dot{v}_t$ if we assume quasi-dynamic, where M is the inertia matrix of \mathcal{O} , and \dot{v}_t is approximated as $v_t / \Delta t$.

Lastly, adding terminal constraints $q_0 \in \mathcal{Q}_{init}$ and $q_T \in \mathcal{Q}_{goal}$, we formulate the following infinite programming with complementarity constraints trajectory optimization (IPCC-TO) problem $P(Y)$:

$$\min_{q, v, u, z} f(q, v, u, z) \quad (9a)$$

$$\text{s.t. } q_0 \in \mathcal{Q}_{init}, q_T \in \mathcal{Q}_{goal} \quad (9b)$$

$$(2), (4), (5), (6), (8) \quad t = 0, \dots, T \quad (9c)$$

$$(3) \quad t = 0, \dots, T - 1, \quad (9d)$$

where q , v , and u concatenate the state and control variables along all time steps, and z represents the contact force distribution at all points in time. Detailed instantiations of Equation 9 for both 2D and 3D scenarios are available in the Appendix.

Exchange method. The IPCC-TO problem $P(Y)$ not only has infinitely many constraints, but also an infinity of variables in z . To solve it using numerical methods, we require that z only is non-zero at a finite number of points. Indeed, if an optimal solution q^* is supported by a finite subset of index points $\tilde{Y} \in Y$, then it suffices to solve for the values of z at these supporting points, since z should elsewhere be zero thanks to (4). This concept is used in the exchange method to solve SIP problems [26], and we extend it to solve IPCC-TO.

The exchange method progressively instantiates finite index sets $\tilde{Y} \equiv [\tilde{Y}_0, \dots, \tilde{Y}_T]$ and their corresponding finite-dimensional MPCCs whose solutions converge toward the true optimum [27]. The solving process can be viewed as a bi-level optimization. In the outer loop, index points are selected by an oracle to be added to the index set \tilde{Y} , and then in the inner loop, the optimization $P(\tilde{Y})$ is passed to the solver to solve for a small number of iterations. The outer loop will then decide how much should move toward this solution. Specifically, if we let $(\tilde{x}^* = [q^*, v^*, u^*], \tilde{z}^*)$ be the optimal solution to $P(\tilde{Y})$, then as \tilde{Y} grows denser, the iterates of $(\tilde{x}^*, \tilde{z}^*)$ will eventually approach an optimum of $P(Y)$.

Given a finite number of instantiated contact points $\tilde{Y} \subset Y$, we can solve a discretized version of the problem which only creates constraints and variables corresponding to \tilde{Y} . Also, through replacing the distribution of $z(y)$ with Dirac impulses, integrals are replaced with sums and we formulate the finite MPCC problem $P(\tilde{Y})$ in the following form:

$$\min_{q,v,u,z} \tilde{f}(q, v, u, z) \quad (10a)$$

$$\text{s.t. } q_0 \in \mathcal{Q}_{init}, q_T \in \mathcal{Q}_{goal} \quad (10b)$$

$$(2), (4), (5), (8) \quad t = 0, \dots, T \quad (10c)$$

$$(3) \quad t = 0, \dots, T-1 \quad (10d)$$

$$\tilde{s}(q_t, u_t, z; \tilde{Y}_t) = 0 \text{ or } M\dot{v} \quad t = 0, \dots, T \quad (10e)$$

where $\tilde{f}(q, v, u, z) := \sum_{t=0}^T [f_{q,v,u}(q_t, v_t, u_t) + \sum_{y \in \tilde{Y}_t} f_z(q_t, y, z(y))]$, and $\tilde{s}(q_t, u_t, z; \tilde{Y}_t) = s_{q,u}(q_t, u_t) + \sum_{y \in \tilde{Y}_t} s_z(q_t, y, z(y))$.

Oracle. Identifying an effective selection strategy for index points is a key to solving the IPs, and this strategy is denoted the Oracle. A naive Oracle would sample index points incrementally from Y_t (e.g., randomly or on a grid) at each of the discretized time steps along the trajectory, and hopefully, with a sufficiently dense set of points the iterates of solutions will eventually approach an optimum. However, this approach is inefficient, as most new index points will not yield active contact forces during the iteration. Better strategies seek to identify a small number of points that would be active at an optimal solution. This work introduces two different Oracle designs, the Maximum Violation Oracle (MVO) and the Time-active Maximum Violation Oracle (TAMVO), which will be discussed in more detail in Sec. III-C.

Merit function for the outer iteration. After solving $P(\tilde{Y})$ in an outer iteration, we get a step direction from the current iterate (\tilde{x}, \tilde{z}) toward $(\tilde{x}^*, \tilde{z}^*)$, where $x = [q, v, u]$ is a concatenation of all the optimization variables except for z . However, due to nonlinearity, the full step may lead to a worse constraint violation for the original problem $P(Y)$. To avoid this, we perform a line search over the following merit function that balances reducing the objective and reducing the constraint error on the infinite dimensional problem $P(Y)$:

$$\phi(x, z; w_c) = f(x, z) + w_c \|b(x, z)\|_1, \quad (11)$$

where b denotes the vector of constraint violations of Problem (10). Also, in SIP for collision geometries, a serious problem is that using existing instantiated index parameters, a step may go too far into areas where the minimum of the distance function $g^*(q_t) \equiv \min_{y \in Y_t} g(q_t, y)$ greatly violates the inequality, and the optimization loses reliability. So we add the max-violation $g^{*-}(q_t)$ to b , in which we denote the negative component of a term as $\cdot^- \equiv \min(\cdot, 0)$. Thus, $g^{*-} \equiv \min(\min_{y \in Y_t} g(q_t, y), 0)$.

Convergence criteria. We denote the index set \tilde{Y} instantiated at the k^{th} outer iteration as \tilde{Y}^k , the corresponding MPCC as $P_k = P(\tilde{Y}^k)$, and the solved solution as (x_k, z_k) .

The convergence condition is defined as $\|[\Delta x, \Delta z]\| \leq \epsilon_x \cdot n_{xz}$ and $|z_k|^T |g(q_k, \tilde{Y}^k)| + |v(q_k, v_k, \tilde{Y}^k)|^T |h(q_k, \tilde{Y}^k, z_k)| \leq \epsilon_{gap} \cdot n_{cc}$ and $|s(x_k, z_k, \tilde{Y}^k)| \leq \epsilon_s \cdot T$ (or $|s(x_k, z_k, \tilde{Y}^k)| - MV \leq \epsilon_s \cdot T$) and $\sum_t g_k^*(x_{k,t}) < \epsilon_p \cdot T$, where n_{xz} is the dimension of the optimization variable and n_{cc} is the number of complementarity constraints, ϵ_x is the step size tolerance, ϵ_{gap} is the complementarity gap tolerance, ϵ_s is the balance tolerance, and ϵ_p is the penetration tolerance. With a little abuse of notation, $g(x_k, \tilde{Y}^k)$ is the concatenation of the function value of all the points in \tilde{Y}^k , and similar for $v(q_k, v_k, \tilde{Y}^k)$ and $h(q_k, \tilde{Y}^k, z_k)$. \dot{V} is the concatenation of all the \dot{v}_t for $t \in \mathcal{T}$.

C. Oracle Choice

As described above, the oracle design is a key component of STOCS. We compare Maximum Violation Oracle (MVO), that adds the closest / deepest penetrating points between the object and the environment at each time step along the trajectory, along with a new method, Time Active Maximum Violation Oracle (TAMVO) with smoothing. TAMVO selects index points more judiciously and only adds the closest / deepest penetrating points at a specific time step.

MVO (Alg. 2) adds the closest or deepest penetrating points between the object and environment at each time step along the trajectory to all candidate index points across time \tilde{Y}_t^k s. Note, however, that it may still include index points that do not generate active contact forces during the iteration. For instance, as illustrated in Fig. 2(a), the closest or deepest penetrating points at time step t are typically active only around that specific period.

To address this issue, we introduce TAMVO (Alg. 3). In this refined approach, the index set is no longer the same across time steps. Lines 8–12 identify closest points at each time step.

Algorithm 2 Maximum-Violation Oracle

Input $q_{0:T}, \tilde{Y}^{k-1}$, adding threshold d_{min}^* and d_{max}^*
Output \tilde{Y}_k

- 1: $\tilde{Y}^k \leftarrow \tilde{Y}^{k-1}$
- 2: **for** $t = 0, \dots, T$ **do**
- 3: $y^* = \arg \min_{y \in Y_t} g(q_t, y)$
- 4: $d^* = g(q_t, y^*)$
- 5: **if** y^* is not in \tilde{Y}_t^k and $d^* < d_{max}^*$ **then**
- 6: **if** $\|y - y^*\| > d_{min}^* \forall y \in \tilde{Y}_t^k$ **then**
- 7: **for** $t = 0, \dots, T$ **do**
- 8: add y^* to \tilde{Y}_t^k

Algorithm 3 Time-Active Maximum-Violation Oracle

Input $q_{0:T}, \tilde{Y}^{k-1}$, max object-environment distance d_{max}^* , contact uniqueness threshold ϵ , time smoothing step n_t , spatial disturbances N_s
Output \tilde{Y}^k

- 1: $\tilde{Y}^k \leftarrow \tilde{Y}^{k-1}$
- 2: $\tilde{Y}' \leftarrow [[]_0, []_1, \dots, []_T]$
- 3: **for** $t = 0, \dots, T$ **do**
- 4: $y^* = \arg \min_{y \in \tilde{Y}_t} g(q_t, y)$
- 5: $d^* = g(q_t, y^*)$
- 6: **if** $d^* < d_{max}^*$ **then**
- 7: add y^* to $\tilde{Y}'[t]$
- 8: **for** $t = 0, \dots, T$ **do**
- 9: **for** $n_s \in N_s$ **do**
- 10: $y_s = \arg \min_{y \in Y_t} g(q_t + n_s, y)$
- 11: **if** $g(q_t + n_s, y_s) < d_{max}^*$ **then**
- 12: add y_s to $\tilde{Y}'[t]$
- 13: **for** $t = 0, \dots, T$ **do**
- 14: **for** $t' = t - n_t, \dots, t + n_t$ **do**
- 15: **if** $0 \leq t' \leq T$ **then**
- 16: **for** y' in $\tilde{Y}'[t']$ **do**
- 17: **if** $\|y' - y\| > \epsilon \forall y \in \tilde{Y}_t^k$ **then**
- 18: add y' to \tilde{Y}_t^k

Duplicate points (within threshold ϵ) are excluded in lines 13–18. Given default parameters $n_t = 0$ and $N_s = [0]$, adds only the closest or most deeply penetrating points at the current time step t to \tilde{Y}_t^k .

Choosing only the closest points at the current iterate is potentially not the most ideal choice unless the current iterate is near-optimal. A better choice would anticipate which points are active at the optimum. To address this, we introduce the following two strategies designed to mitigate this issue.

Spatial Disturbance (SD). Recognizing that the current iterate is likely to be in the neighborhood of the optimal solution, the SD approach introduces perturbations to the current solution to add new candidate contact points. Consequently, in lines 9–12 of Alg. 3, q_t is perturbed with disturbance n_s to choose closest points. We choose to perturb along each dimension of q_t in both positive and negative directions. In 3D, this strategy chooses 12 perturbations accounting for both increases and decreases in x, y, z and $roll, pitch, yaw$. An illustration of adding perturbation to rotation in 2D is shown in Fig. 2(e).

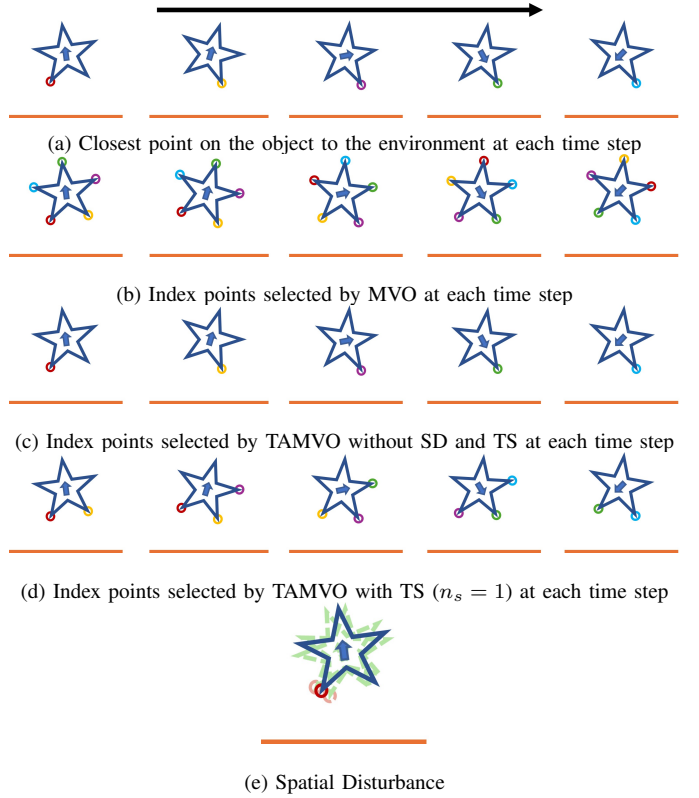


Fig. 2. Comparing various Oracles. The object trajectory is depicted as moving from left to right (as indicated by the black arrow) and undergoing clockwise rotation (as indicated by the arrow on the star). (b) In Maximum Violation Oracle (MVO), the closest point on the object is added to the candidate set at every time step. (c) The Time-Active Maximum Violation Oracle, without Spatial Disturbance and Spatial Disturbances, introduces the closest point only at the current time step. The Time Smoothing technique with $n_s = 1$, demonstrated in (d), extends constraint imposition to the closest points identified at adjacent time steps. (e) presents the Spatial Disturbance technique applied at a specific time step, with only disturbed rotation illustrated. [Best viewed in color.]

Time Smoothing (TS). Considering that the closest points may be active not just at the current time step t , but also during a surrounding interval, in line 14 of Alg. 3, the closest points detected within the adjacent time steps from $t - n_t$ to $t + n_t$, governed by a parameter n_t , are added to the index set of time step t . The effect of using TS is illustrated in Fig. 2(d).

IV. EXPERIMENTAL RESULTS

The proposed methods are implemented in Python using the optimization interface and the SNOPT solver [28] provided by Drake [29].

A. Experiments in 2D

First, we compare STOCS with vanilla MPCC to evaluate the efficacy of dynamic contact selection on an object pivoting task. We finely discretize the object geometries to better illustrate the advantages of our method. Vanilla MPCC involves adding all index points in Y to an MPCC problem without selection, resulting in a larger optimization problem than P^k in STOCS. MVO is employed in this experiment, and $T = 20$,




Fig. 3. Pivoting trajectories for 3 objects solved by STOCS. The object’s poses, the contact between the object and the robot (red triangle), and object-environment contact points (green dots) are plotted for each time step. [Best viewed in color.]

TABLE I

NUMERICAL RESULTS ON THE PIVOTING TASK. NUMBER OF POINTS IN THE OBJECT’S REPRESENTATION (# POINT), SOLVE TIME (TIME), OUTER LOOP ITERATION NUMBER (OUTER ITERS), AND AVERAGE ACTIVE INDEX POINTS FOR EACH ITERATION (INDEX POINTS) ARE REPORTED IN THE TABLE.

Object	# Points	STOCS			MPCC
		Time (s)	Outer iters.	Index points	Time (s)
Box	212	4.3	2	2.0	2908.5
Peg	214	45.2	5	4.2	5503.9
Mustard	400	33.6	5	6.8	Failed


$\Delta T = 0.1$ s, $\mu_{mnp} = 1.0$ and $\mu_{env} = 0.5$ are used for this set of the experiments.

The results are presented in Table I. We observe that STOCS can be around two to three orders of magnitude faster than MPCC, and can solve problems that MPCC cannot solve. STOCS selects only a small amount of points from the total number of points in the objects’ representation on average, which greatly decreases the dimension of the instantiated optimization problem and reduces solve time. The solved trajectories of STOCS for three different object are shown in 3.


Next, we evaluate STOCS in a 2D setting with TAMVO, focusing on testing its applicability with objects and environments whose geometries are significantly more complex than those presented in the examples of [21]. In this experiment, we set $n_t = 1$ and $N_s = [1e^{-2}]$ as the parameters for TAMVO, and we designate these as the default values for TS and SD separately. Four tasks are designed for this evaluation: pushing a dented object across uneven terrain, inserting a tilted peg into a correspondingly angled hole, and maneuvering a bean-shaped object across two distinct curvilinear terrains. The trajectories planned for these tasks are depicted in Fig. 4, and the detailed information regarding the objects’ geometries and the solve of the trajectories are presented in Table II.

$T = 10$, $\Delta T = 0.1$ s, $\mu_{mnp} = 1.0$ and $\mu_{env} = 0.5$ are used for this set of the experiments in 2D.


Same as before, STOCS selects only a small number of points from the entire set in the object’s representation to establish a solution. Typically, STOCS reaches convergence after just a few iterations. The final example, sliding the bean along curve 2, demands additional iterations and time for convergence due to the curve’s curvature shifting from




(a) Pushing a dented object on uneven terrain



(b) Pushing a tilted peg into angled hole



(c) Sliding bean on curve 1



(d) Sliding bean on curve 2

Fig. 4. Trajectories planned by STOCS on 2D examples. The progression from the start to the end of the trajectory is indicated by a dark to light gradient. For clarity, the instantiated object-environment contact points are depicted only at the first and the last time steps. [Best viewed in color.]

concave to convex. This alteration results in different contact points at different stages along the trajectory, which takes more iterations to be fully instantiated.

We visualize the solving progress of STOCS for two examples in Fig. 5. In the initial few iterations, penetrations are resolved and contacts are established. Subsequently, the solver begins to satisfy additional constraints, including object dynamics, and gradually instantiates index points to meet the object’s balance constraints. In later iterations, fine-tuning ensures that all constraints are fully satisfied.

B. Experiments in 3D

Next, we evaluate STOCS in 3D. To demonstrate its generalizability, we collect object geometries from the YCB dataset [30], the Google Scanned Objects [31], and 3D models found online [32], [33]. All the objects used in the study are shown in Fig. 6, and all the environments used in the study are shown in Fig. 7. Klampt [34] and the code in [23] are used to find the closest points between two complex shaped geometries.

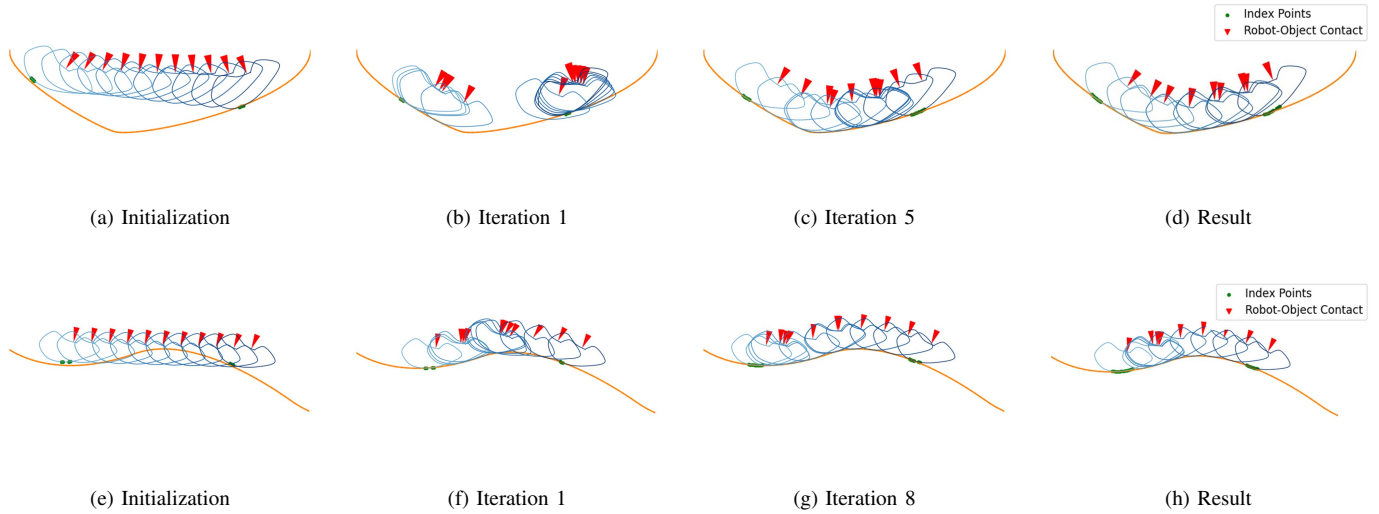


Fig. 5. Visualization of STOCS iterations for two examples (each example shown in a separate row). In the initial few iterations, penetrations are resolved and contacts are established. Subsequently, the solver begins to satisfy additional constraints, including object dynamics, and gradually instantiates index points to meet the object’s balance constraints. In later iterations, fine-tuning ensures that all constraints are fully satisfied. For clarity, the instantiated object-environment contact points are depicted only at the first and the last time steps. [Best viewed in color.]

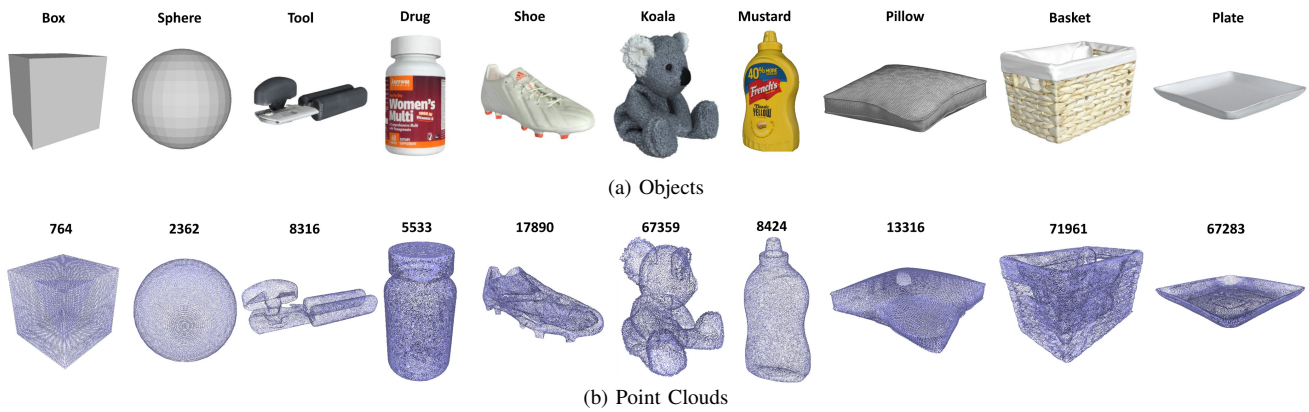


Fig. 6. 3D test objects (first row) and their point clouds (second row) used in the experiments. Not drawn to scale.




Fig. 7. 3D test environments. Not drawn to scale.





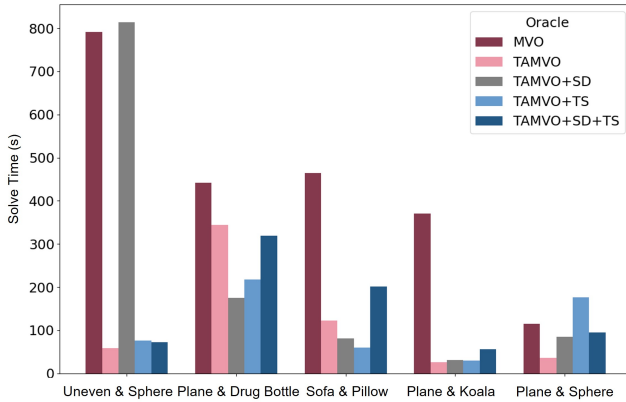
Fig. 8. Success rates of STOCS for varying choice of Oracle.

As demonstrated in the first set of experiments, employing a vanilla MPCC approach by incorporating all index points in Y into an MPCC problem without selection results in a large optimization problem. Such an approach can take a long

time to find a solution. Consequently, we have opted not to pursue experiments with the vanilla MPCC in this study, given



(a) Average Index Points



(b) Solve Time

Fig. 9. The average index points selected at each time step by the different Oracles (a) and the solve time (b) for tasks that were successfully solved by all Oracles. [Best viewed in color.]

that the 3D point clouds utilized herein contain thousands of points on average. Additionally, approximating the friction cone with a polyhedral model in a 3D context further increases the number of constraints for each index point, exacerbating the complexity beyond that encountered in 2D scenarios.

To demonstrate the effectiveness of STOCS in planning with high-fidelity geometric representations, we conduct experiments on ten different objects represented by dense point clouds sampled on the surface of the objects’ meshes, and five different environments represented by Signed Distance Field (SDF). The SDFs are calculated offline on a grid that encloses the corresponding environment given a closed polygonal mesh

TABLE II

NUMERICAL RESULTS OF STOCS IN 2D. NUMBER OF POINTS IN THE OBJECT’S REPRESENTATION (# POINT), SOLVE TIME (TIME), OUTER ITERATION COUNT (OUTER ITERS), AND AVERAGE ACTIVE INDEX POINTS FOR EACH ITERATION (INDEX POINTS) ARE REPORTED.

Environment	Object	# Point	Outer iters.	Index points	Time (s)
Uneven	Dented	543	4	9.05	30.83
Tilted Hole	Tilted Peg	214	4	12.93	40.11
Curve 1	Bean	100	7	9.43	72.74
Curve 2			14	14.29	636.89

TABLE III

NUMERICAL RESULTS OF STOCS IN 3D. NUMBER OF POINTS IN THE OBJECT’S REPRESENTATION (# POINT), SOLVE TIME (TIME), OUTER ITERATION COUNT (OUTER ITERS), AND AVERAGE ACTIVE INDEX POINTS FOR EACH ITERATION (INDEX POINTS) ARE REPORTED.

Environment	Object	Task	# Point	Outer iters.	Index points	Time (s)
Plane	Box	Push	764	3	5.75	24.84
	Shoe	Push	17890	6	4.95	144.71
	Koala	Pivot	67359	4	5.89	37.93
	Mustard	Pivot	8424	3	8.58	59.37
	Sphere	Roll	2362	6	4.39	94.42
	Tool	Rotate	8316	6	5.09	99.85
	Drug	Roll	5533	10	6.35	319.72
Curve	Koala	Pivot	67359	9	13.47	676.01
	Sphere	Roll	2362	4	6.86	72.66
Sofa	Pillow	Pivot	13316	7	9.95	201.39
Shelf	Basket	Push	71961	7	23.10	421.30
Plate	Plate	Slide	67283	3	24.67	54.02

of the environment, and values off of the grid vertices are approximated via trilinear interpolation.

Using STOCS, we plan for pushing, pivoting, rolling and rotating trajectories on these objects. The resulting planned trajectories are illustrated in Fig. 10, while detailed information regarding the objects’ geometries and the solve of the trajectories are presented in Table III. Like the experiment in 2D, we set $n_t = 1$ and $N_s = [1e^{-2}]$ as the default parameters for TAMVO. $\Delta T = 0.1$ s, $\mu_{mnp} = 1.0$ and $\mu_{env} = 1.0$ are used for all the experiments in 3D. For all experiments, $T = 10$ is used except in the tasks of pushing a basket on a shelf and sliding a plate on another plate, where $T = 5$ is used. To model the robot manipulator as having a patch contact, 3 to 5 object vertices in the neighborhood of the indicated cone are allowed to be used as contact points.

Following the initial assessments, we further evaluated the efficacy of the TAMVO alongside the SD and TS techniques through a set of comparative experiments. These experiments utilized STOCS to plan trajectories for the same set of tasks, with the primary variation being the specific oracle employed in each scenario.

Figure 8 presents the success rates of all tested Oracles. The data shows that the MVO achieves a 75% success rate. In contrast, TAMVO without SD and TS exhibits worse performance than MVO; this is particularly evident in 3D scenarios where relying solely on the nearest object-to-environment point is inadequate for fulfilling the object’s balance constraints. The SD and TS techniques, introduced to address this challenge, both demonstrated enhanced performance when combined with TAMVO, surpassing the success rate of TAMVO alone. Furthermore, the integration of SD and TS with TAMVO consistently achieved successful trajectory planning for all tasks.

Figure 9 displays the average number of index points selected at each time step by the different Oracles assessed in our study, focusing on tasks that were successfully solved by all Oracles. As depicted in Fig. 9(a), the MVO selects a larger number of points than TAMVO and all its variations.

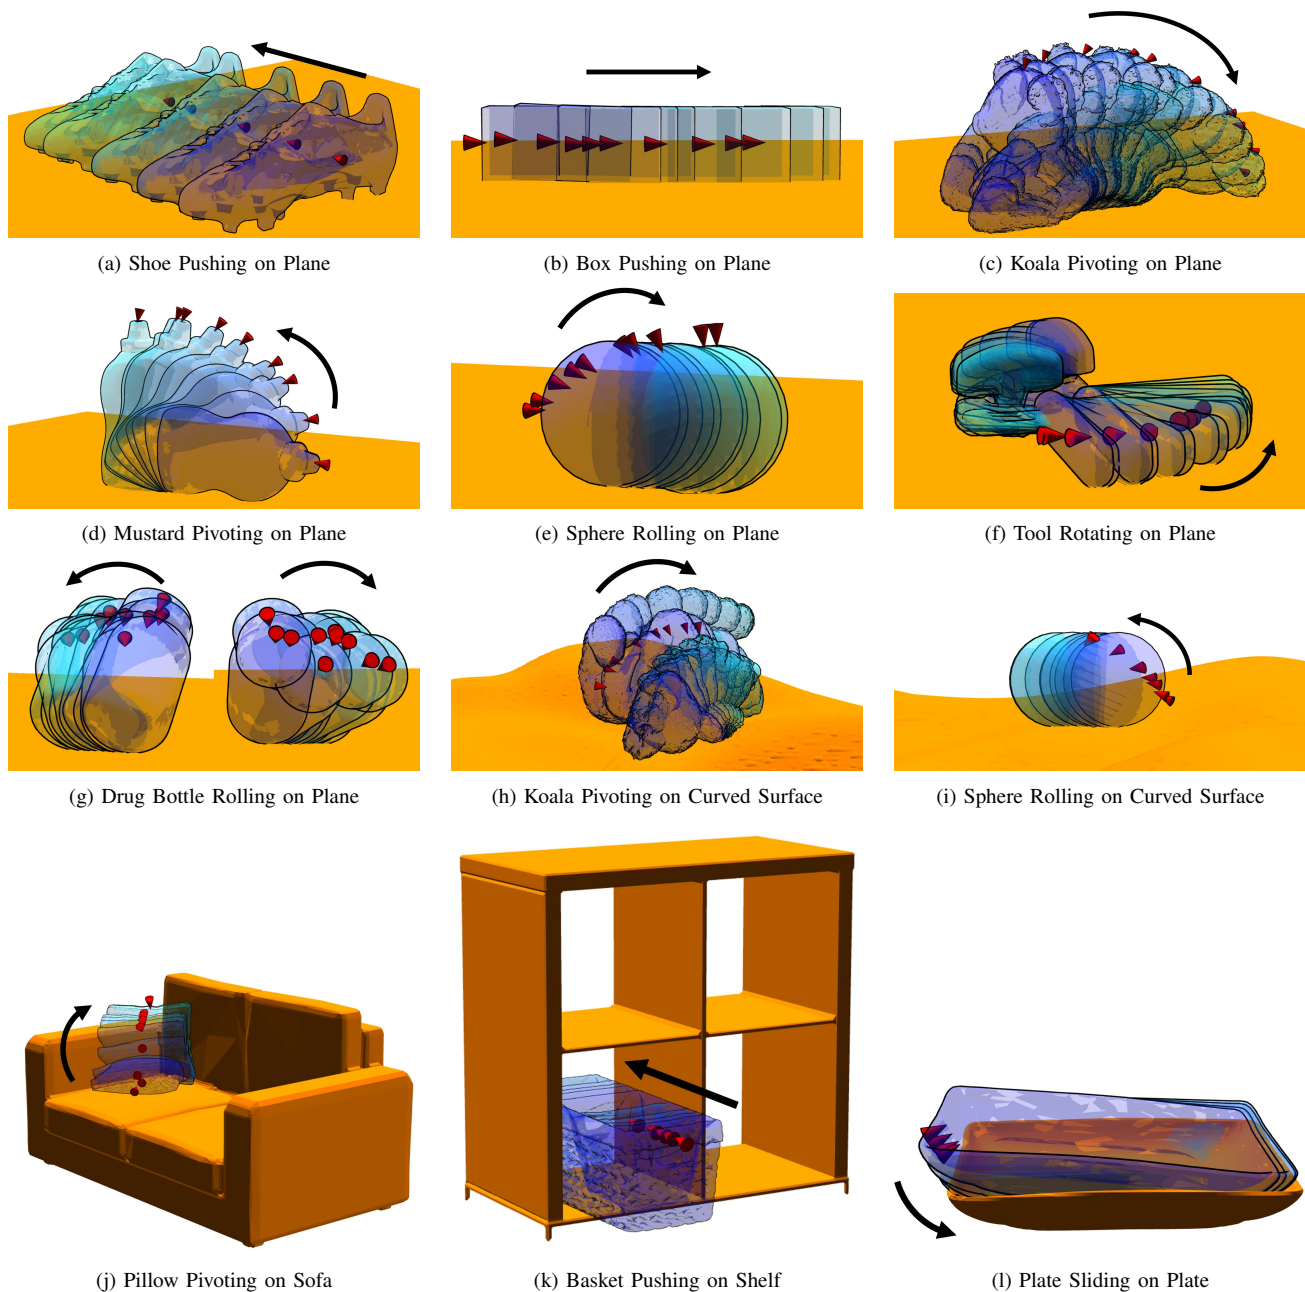


Fig. 10. Trajectories planned by STOCS for the 3D examples. Progress along the trajectory is indicated by color (dark to light). The black arrow indicates the object's movement direction. The red cone marks the contact location of the manipulator on the object. [Best viewed in color.]

Notably, TAMVO combined with SD and TS can successfully plan trajectories for all tasks while selecting fewer index points compared to MVO. These findings validate our hypothesis regarding TAMVO: an index point identified at time step t is most valuable within a temporal vicinity of t . Furthermore, these results substantiate our rationale for introducing SD and TS, affirming that a localized exploration in both temporal and spatial dimensions offers a more efficient strategy than incorporating index points identified at distant time steps along the trajectory.

Figure 9(b) presents the solve times for STOCS employing various Oracles across all successful tasks. The results indicate that the quantity of index points selected by an Oracle does not necessarily correlate with the solve time. For instance, in the

sphere rolling on plane task, TAMVO+SD+TS chooses a larger number of index points than TAMVO+TS, yet the solve time for TAMVO+SD+TS is much faster than that of TAMVO+TS. Similarly, in the case of the drug bottle rolling on plane task, it can be seen that both TAMVO+SD+TS and TAMVO+TS select a comparable number of index points, yet the solve time for TAMVO+TS is much faster than TAMVO+SD+TS. This phenomenon underscores that a larger number of index points does not invariably lead to longer solve time. Although TAMVO+SD+TS provides the best performance in terms of success rate, it is not guaranteed to give the best solve time for all different tasks.

We further assess the impact of N_s and n_t on STOCS performance by varying these parameters individually with

TAMVO+SD and TAMVO+TS for the tasks listed in Table III. The results are shown in Table IV. TAMVO+SD achieves the best performance when $N_s = [1e^{-2}]$. We observe that, on average, fewer index points are added when N_s is set to a smaller value, as index points identified under small spatial disturbances often overlap. The optimal N_s depends on factors such as object size, point cloud resolution, and signed distance field resolution. However, to demonstrate the generalizability of STOCS, we keep this parameter constant as much as possible across all experiments rather than tuning it individually for each task and geometry.

TAMVO+TS yields the best results with $n_t = 3$. As n_t increases, both the success rate and average number of index points increase. This rise in index points is inherent to the TS design, as index points identified at nearby time steps are added to the current time step. These additional points provide a more accurate local representation of the object’s geometry, which facilitates finding feasible solutions, though at the cost of longer solve times due to larger optimization problem sizes.

V. CONCLUSION AND DISCUSSION

This paper introduces Simultaneous Trajectory Optimization and Contact Selection (STOCS) to address the scaling problem in contact-implicit trajectory optimization (CITO). Through embedding CITO in an infinite programming (IP) framework to dynamically instantiate possible contact points and contact times between the object and environment inside the optimization loop, STOCS extends the application of CITO to a much higher number of contact points. The introduction of the Time-Active Maximum Violation Oracle along with Spatial Disturbance and Temporal Smoothing techniques for more efficient index point selection further enables STOCS to plan contact-rich manipulation trajectories with dense point clouds representation of complex-shaped objects. To the best of our knowledge, this is the first method capable of planning contact-rich manipulation trajectories with high-fidelity geometric representations in 3D.

Looking ahead, to further facilitate the practical application of the STOCS algorithm to raw object point clouds obtained from RGB-D sensors, it will be essential to devise effective pre-processing techniques for the point cloud data. This is crucial given that raw depth data typically contain a significant

amount of noise. Moreover, STOCS presupposes the knowledge of specific physical properties of the object, such as mass, center of mass, and friction coefficients. However, accurately measuring these parameters in practice can be challenging, and discrepancies between the nominal values used during planning and their actual real-world counterparts may lead to failures in execution. It would be of considerable interest to explore the incorporation of robust optimization techniques into the STOCS framework to develop manipulation trajectories that maintain robustness against uncertainty in physical parameters. Also, it is worth exploring methods to allow robot-object contact to change, such as incorporating a sampling-based planner to generate contact transitions between the robot and the object.

ACKNOWLEDGMENTS

This paper is partially supported by NSF Grant #IIS-1911087.

REFERENCES

- [1] M. R. Dogar, K. Hsiao, M. T. Ciocarlie, and S. S. Srinivasa, “Physics-based grasp planning through clutter,” in *Robotics: Science and Systems*, vol. 8, 2012, pp. 57–64.
- [2] C. Eppner and O. Brock, “Planning grasp strategies that exploit environmental constraints,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2015, pp. 4947–4952.
- [3] M. Kelly, “An introduction to trajectory optimization: How to do your own direct collocation,” *SIAM Review*, vol. 59, no. 4, pp. 849–904, 2017.
- [4] G. Schultz and K. Mombaur, “Modeling and optimal control of human-like running,” *IEEE/ASME Transactions on Mechatronics*, vol. 15, no. 5, pp. 783–792, 2009.
- [5] M. Posa, C. Cantu, and R. Tedrake, “A direct method for trajectory optimization of rigid bodies through contact,” *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, 2014.
- [6] I. Mordatch, Z. Popović, and E. Todorov, “Contact-invariant optimization for hand manipulation,” in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2012, pp. 137–144.
- [7] I. Mordatch, E. Todorov, and Z. Popović, “Discovery of complex behaviors through contact-invariant optimization,” *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 1–8, 2012.
- [8] A. Patel, S. L. Shield, S. Kazi, A. M. Johnson, and L. T. Biegler, “Contact-implicit trajectory optimization using orthogonal collocation,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2242–2249, 2019.
- [9] A. Ö. Önoel, R. Corcodel, P. Long, and T. Padr, “Tuning-free contact-implicit trajectory optimization,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1183–1189.
- [10] Z.-Q. Luo, J.-S. Pang, and D. Ralph, *Mathematical programs with equilibrium constraints*. Cambridge University Press, 1996.
- [11] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.
- [12] M. Zhang and K. Hauser, “Semi-infinite programming with complementarity constraints for pose optimization with pervasive contact,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6329–6335.
- [13] F. H. Clarke, *Optimization and nonsmooth analysis*. SIAM, 1990.
- [14] K. Harada, K. Hauser, T. Bretl, and J.-C. Latombe, “Natural motion generation for humanoid robots,” in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 833–839.
- [15] N. Chavan-Dafle, R. Holladay, and A. Rodriguez, “Planar in-hand manipulation via motion cones,” *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 163–182, 2020.
- [16] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Romeres, “Robust pivoting: Exploiting frictional stability using bilevel optimization,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2022, pp. 992–998.
- [17] B. Aceituno-Cabezas and A. Rodriguez, “A global quasi-dynamic model for contact-trajectory optimization in manipulation,” 2020.

TABLE IV

NUMERICAL RESULTS ANALYZING THE IMPACT OF N_s AND n_t ON THE PERFORMANCE OF STOCS. FOR TASKS SUCCESSFULLY SOLVED ACROSS ALL TESTED PARAMETERS, THE AVERAGE SOLVE TIME AND AVERAGE NUMBER OF INDEX POINTS PER ITERATION ARE CALCULATED SEPARATELY FOR EACH SET OF EXPERIMENTS.

TAMVO+SD			
N_s	$[1e^{-2}]$	$[1e^{-3}]$	$[1e^{-4}]$
Success Rate	91.7%	66.7%	66.7%
Average Solve Time (s)	291.7	151.9	160.9
Average Index Points	3.9	2.7	2.9
TAMVO+TS			
n_t	1	2	3
Success Rate	58.3%	66.7%	75.0%
Average Solve Time (s)	168.3	256.9	499.8
Average Index Points	4.6	6.2	9.1

- [18] X. Cheng, E. Huang, Y. Hou, and M. T. Mason, "Contact mode guided sampling-based planning for quasistatic dexterous manipulation in 2d," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6520–6526.
- [19] —, "Contact mode guided motion planning for quasidynamic dexterous manipulation in 3d," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2730–2736.
- [20] E. Huang, X. Cheng, and M. T. Mason, "Efficient contact mode enumeration in 3d," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2021, pp. 485–501.
- [21] M. Zhang, D. K. Jha, A. U. Raghunathan, and K. Hauser, "Simultaneous trajectory optimization and contact selection for multi-modal manipulation planning," in *Robotics: Science and Systems*, 2023.
- [22] M. T. Mason, *Mechanics of robotic manipulation*. MIT Press, 2001.
- [23] K. Hauser, "Semi-infinite programming for trajectory optimization with non-convex obstacles," *The International Journal of Robotics Research*, vol. 40, no. 10-11, pp. 1106–1122, 2021.
- [24] E. J. Anderson and A. B. Philpott, *Infinite Programming: Proceedings of an International Symposium on Infinite Dimensional Linear Programming Churchill College, Cambridge, United Kingdom, September 7–10, 1984*. Springer Science & Business Media, 2012, vol. 259.
- [25] D. E. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction," *International Journal for Numerical Methods in Engineering*, vol. 39, no. 15, pp. 2673–2691, 1996.
- [26] M. López and G. Still, "Semi-infinite programming," *European Journal of Operational Research*, vol. 180, no. 2, pp. 491–518, 2007.
- [27] R. Reemtsen and S. Görner, "Numerical methods for semi-infinite programming: a survey," in *Semi-Infinite Programming*. Springer, 1998, pp. 195–275.
- [28] P. E. Gill, W. Murray, and M. A. Saunders, "Snopt: An sqp algorithm for large-scale constrained optimization," *SIAM review*, vol. 47, no. 1, pp. 99–131, 2005.
- [29] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: <https://drake.mit.edu>
- [30] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols," *arXiv preprint arXiv:1502.03143*, 2015.
- [31] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke, "Google scanned objects: A high-quality dataset of 3d scanned household items," in *IEEE Intl. Conf. Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2553–2560.
- [32] Open Robotics. [Online]. Available: <https://app.gazebosim.org/OpenRobotics/fuel/models/Sofa>
- [33] Sketchfab. [Online]. Available: <https://sketchfab.com/3d-models/burlap-pillow-1ef567278bba4db68ac5488d6b4bc851>
- [34] KlampT – Intelligent Motion Laboratory at UIUC. [Online]. Available: <http://motion.cs.illinois.edu/klampT/>



Mengchao Zhang (Member, IEEE) received the B.S. degree in Vehicle Engineering from Tongji University, Shanghai, China, in 2017, and the M.S. degree in Mechanical Engineering from the University of Illinois at Urbana-Champaign, Urbana, IL, USA. He is currently working toward a Ph.D. degree in Mechanical Engineering with the Department of Mechanical Science & Engineering, University of Illinois at Urbana-Champaign, Urbana, USA. His research focuses on contact-rich manipulation planning and control.



Division (DSCD) in 2019.

Devesh K. Jha (Senior Member, IEEE) is a senior principal research scientist at the Mitsubishi Electric Research Laboratories (MERL). He received his PhD from Pennsylvania State University in 2016, and his PhD thesis was on the decision & control of autonomous systems. He also got a Master's degree in Mathematics from the same university. His research interests are in the areas of machine learning, time series analytics, and robotics. He was a recipient of the Rudolf Kalman Best Paper award by the ASME Dynamics, Systems and Control



systems, and security businesses.

Arvind U. Raghunathan is a senior principal research scientist at the Mitsubishi Electric Research Laboratories. He received his PhD in Chemical Engineering from Carnegie Mellon University in 2004. His research focuses on algorithms for optimization of large-scale nonlinear and mixed integer nonlinear programs with applications in power grid, transportation systems, and model-based control of processes. He previously worked at the United Technologies Research Center for 7 years developing optimization algorithms for aerospace, elevator, energy



Kris Hauser (Senior Member, IEEE) is a Professor in the School of Computing and Data Science at the University of Illinois at Urbana-Champaign, and Affiliate of the Department of Electrical and Computer Engineering and the Department of Mechanical Science and Engineering. He received his PhD in Computer Science from Stanford University in 2008, bachelor's degrees in Computer Science and Mathematics from UC Berkeley in 2003, and worked as a postdoctoral fellow at UC Berkeley. He then joined the faculty at Indiana University from 2009-2014, where he started the Intelligent Motion Lab, and then joined the faculty of Duke University from 2014-2019. Prof. Hauser is a recipient of a Stanford Graduate Fellowship, Siebel Scholar Fellowship, Best Paper Award at IEEE International Conference on Humanoid Robots 2015, the NSF CAREER award, and three Amazon Research Awards.

APPENDIX

We discuss the detailed instantiation of the MPCC problem that needs to be solved inside STOCS for both 2D and 3D scenarios separately.

A. 2D

1) *Geometry modeling*: To handle collision avoidance, we establish a semi-infinite constraint $g(q, y)$ between the object and the environment. The object is represented as a densely sampled point cloud on its surface, and the environment is represented as a signed distance field (SDF) $\psi(x) : \mathbb{R}^2 \rightarrow \mathbb{R}$ which supports $O(1)$ depth lookup. Specifically, we define $g(q, y) = \psi(T_q \cdot y_p)$ with T_q the object transform at configuration q .

2) *Friction force modeling*: The force variable of the i^{th} index point at time step t is $z_{t,i} = (z_{t,i}^N, z_{t,i}^+, z_{t,i}^-)$, which is divided into the normal component $z_{t,i}^N$ and frictional components $z_{t,i}^+$ and $z_{t,i}^-$ along the tangential direction of the contact surface [25]. Given the normal vector $n_{t,i}^{Nz}$ along the outward surface normal of the environment, and $n_{t,i}^{+z}$ and $n_{t,i}^{-z}$, the overall force applied at index point $y_{t,i}$ is $z_{t,i} = z_{t,i}^N n_{t,i}^{Nz} + z_{t,i}^+ n_{t,i}^{+z} + z_{t,i}^- n_{t,i}^{-z}$. Each of the components of $z_{t,i}$ is required to be non-negative, and given the friction coefficient μ_{env} , the friction cone constraint is given by $\mu_{env} z_{t,i}^N - (z_{t,i}^+ + z_{t,i}^-) \geq 0$.

Similarly, for the i^{th} contact between the robot and the object, we define the force variable at time step t to be $u_{t,i} = (u_{t,i}^N, u_{t,i}^+, u_{t,i}^-)$. Besides, the normal vector $n_{t,i}^{Nu}$ along the inward surface normal of the object, and the tangential vectors $n_{t,i}^{+u}$ and $n_{t,i}^{-u}$ which are perpendicular to the normal vector are defined in the object's local frame.

We adopt a heuristic to reduce the number of constraints in the complementarity condition and accelerate solve times. By only applying complementarity to the normal component of force variables, $g(q_t, y_{t,i}) \cdot z_{t,i}^N = 0$, we reduce the number of distance complementarity constraints from $\sum_0^T 3|\tilde{Y}_t|$ to $\sum_0^T |\tilde{Y}_t|$, and the problem is unchanged because $z_{t,i}^N = 0 \Rightarrow z_{t,i}^+ = 0, z_{t,i}^- = 0$ due to the friction cone constraint.

3) *Force and torque balance*: We establish an equality constraint on the object's force and torque balance. It requires the joint force and torque exerted by the robot, the environment, and the gravity on the object to be 0 if we assume quasistatic, and to be $M\dot{v}$ if we assume quasidynamic, where M is the inertia matrix of the object \mathcal{O} and v is the velocity of the object.

4) *Instantiation of MPCC problem*: With these definitions, the nonlinear programming problem P_k that needs to be solved at the k^{th} iteration of STOCS is the following:

$$\min_{q,v,u,z} \tilde{f}(q, v, u, z) \quad (12a)$$

$$s.t. \quad q_0 \in \mathcal{Q}_{init}, q_T \in \mathcal{Q}_{goal} \quad (12b)$$

$$q_t - q_{t+1} + v_{t+1} \Delta t = 0 \quad \forall t \in \mathcal{T} \setminus \{T\} \quad (12c)$$

$$q_t \in \mathcal{Q}, u_t \in \mathcal{U}, v_t \in \mathcal{V} \quad \forall t \in \mathcal{T} \quad (12d)$$

$$\mu_{mnp} u_{t,i}^N - (u_{t,i}^+ + u_{t,i}^-) \geq 0 \quad \forall i \in \mathcal{C}, \quad \forall t \in \mathcal{T} \quad (12e)$$

$$0 \leq z_{t,i}^N \perp g(q_t, y_{t,i}) \geq 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12f)$$

$$z_{t,i}^+, z_{t,i}^- \geq 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12g)$$

$$\mu_{env} z_{t,i}^N - (z_{t,i}^+ + z_{t,i}^-) \geq 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12h)$$

$$\gamma_{t,i} + \psi_i(q_t, v_t) \geq 0, \quad \gamma_{t,i} - \psi_i(q_t, v_t) \geq 0 \\ \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12i)$$

$$(\mu z_{t,i}^N - (z_{t,i}^+ + z_{t,i}^-))^T \gamma_{t,i} = 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12j)$$

$$(\gamma_{t,i} + \psi_i(q_t, v_t))^T z_{t,i}^+ = 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12k)$$

$$(\gamma_{t,i} - \psi_i(q_t, v_t))^T z_{t,i}^- = 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12l)$$

$$z_{t,i} = z_{t,i}^N \cdot n_{t,i}^{Nz} + z_{t,i}^+ \cdot n_{t,i}^{+z} + z_{t,i}^- \cdot n_{t,i}^{-z} \\ \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (12m)$$

$$u_{t,i} = u_{t,i}^N \cdot n_{t,i}^{Nu} + u_{t,i}^+ \cdot n_{t,i}^{+u} + u_{t,i}^- \cdot n_{t,i}^{-u} \\ \forall i \in \mathcal{C}, \quad \forall t \in \mathcal{T} \quad (12n)$$

$$mg + \sum_i^c u_{t,i} + \sum_i^{\mathcal{I}_t} z_{t,i} = 0 \text{ or } M_{trans} \dot{v}_{trans} \quad \forall t \in \mathcal{T} \quad (12o)$$

$$\sum_i^c (q_t \cdot (c^{mnp,i} - CoM)) \times u_{t,i} + \\ \sum_i^{\mathcal{I}_t} (q_t \cdot (y_{t,i} - CoM)) \times z_{t,i} = 0 \text{ or } M_{rot} \dot{v}_{rot} \quad \forall t \in \mathcal{T} \quad (12p)$$

where $\mathcal{T} = \{0, \dots, T\}$, $\mathcal{I}_t = \{0, \dots, |\tilde{Y}_t|\}$, $\psi_i(q_t, v_t)$ is the relative tangential velocity at contact point $y_{t,i}$. $\mathcal{C} = \{0, \dots, |c^{mnp,i}|\}$, $c^{mnp,i}$ is the i^{th} contact point between the robot and the object in the object's frame, CoM is the center of mass of the object, v_{trans} is the translational velocity of the object, M_{trans} is the translational part of the inertia matrix of the object with appropriate dimension, and v_{rot} and M_{rot} are the corresponding terms for the rotational part.

B. 3D

1) *Geometry modeling*: To handle collision avoidance, like in the 2D scenario, we establish a semi-infinite constraint $g(q, y)$ between the object and the environment. The object is represented as a densely sampled point cloud on its surface, and the environment is represented as a signed distance field (SDF) $\psi(x) : \mathbb{R}^3 \rightarrow \mathbb{R}$ which supports $O(1)$ depth lookup. We define $g(q, y) = \psi(T_q \cdot y_p)$ with T_q the object transform at configuration q .

2) *Friction force modeling*: In 3D scenario, following [5], to preserve the MPCC structure of the resulting optimization problem, we use a polyhedral approximation of the friction cone [25]. The force variable of the i^{th} index point at time step t is $z_{t,i} = (z_{t,i}^N, z_{t,i}^{D1}, \dots, z_{t,i}^{Dd})$, which is expressed in a reference frame with $z_{t,i}^N$ normal to the contact surface, and $z_{t,i}^{D1}, \dots, z_{t,i}^{Dd}$ tangent to the contact surface. The convex hull of the unit vectors along the directions of $z_{t,i}^{D1}, \dots, z_{t,i}^{Dd}$ in \mathbb{R}^2 forms the polyhedral approximation. The normal vector $n_{t,i}^{Nz}$ is along the outward surface normal of the environment, and $n_{t,i}^{D1z}, \dots, n_{t,i}^{Ddz}$ are perpendicular to $n_{t,i}^{Nz}$. Each of the components of $z_{t,i}$ is required to be non-negative, and given the friction coefficient μ_{env} , the friction cone constraint is given by $\mu_{env} z_{t,i}^N - (z_{t,i}^{D1} + \dots + z_{t,i}^{Dd}) \geq 0$.

Similarly, for the i^{th} contact between the robot and the object, we define the force variable at time step t to be $u_{t,i} = (u_{t,i}^N, u_{t,i}^{D1}, \dots, u_{t,i}^{Dd})$. The normal vector $n_{t,i}^{Nu}$ is along the inward surface normal of the object, and $n_{t,i}^{D1u}, \dots, n_{t,i}^{Ddu}$ are perpendicular to $n_{t,i}^{Nu}$.

The same heuristic is used to reduce the number of constraints in the complementarity condition and accelerate solve times. By only applying complementarity to the normal component of force variables, $g(q_t, y_{t,i}) \cdot z_{t,i}^N = 0$, we reduce the number of distance complementarity constraints from $\sum_0^T (d+1)|\tilde{Y}_t|$ to $\sum_0^T |\tilde{Y}_t|$, and the problem is unchanged because $z_{t,i}^N = 0 \Rightarrow z_{t,i}^{Dj} = 0 \forall j \in \{1, \dots, d\}$ due to the friction cone constraint.

3) *Instantiation of MPCC problem*: With these definitions the nonlinear programming problem P_k that needs to be solved at the k^{th} iteration for this pivoting task is the following:

$$\min_{q,v,u,z} \tilde{f}(q, v, u, z) \quad (13a)$$

$$s.t. q_0 \in \mathcal{Q}_{init}, q_T \in \mathcal{Q}_{goal} \quad (13b)$$

$$q_t - q_{t+1} + v_{t+1} \Delta t = 0 \quad \forall t \in \mathcal{T} \setminus \{T\} \quad (13c)$$

$$q_t \in \mathcal{Q}, u_t \in \mathcal{U}, v_t \in \mathcal{V} \quad \forall t \in \mathcal{T} \quad (13d)$$

$$\mu_{mnp} u_{t,i}^N - \sum_j u_{t,i}^{Dj} \geq 0$$

$$\forall i \in \mathcal{C}, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13e)$$

$$0 \leq z_{t,i}^N \perp g(q_t, y_{t,i}) \geq 0 \quad \forall i \in \mathcal{I}_t, \quad \forall t \in \mathcal{T} \quad (13f)$$

$$z_{t,i}^{Dj} \geq 0 \quad \forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13g)$$

$$\mu_{env} z_{t,i}^N - \sum_j z_{t,i}^{Dj} \geq 0$$

$$\forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13h)$$

$$\gamma_{t,i} + \psi(q_t, v_t, y_{t,i}) \cdot n_{t,i}^{Dj} \geq 0$$

$$\forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13i)$$

$$(\mu z_{t,i}^N - \sum_j z_{t,i}^{Dj}) \gamma_{t,i} = 0$$

$$\forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13j)$$

$$(\gamma_{t,i} + \psi(q_t, v_t, y_{t,i}) \cdot n_{t,i}^{Dj}) z_{t,i}^{Dj} = 0$$

$$\forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13k)$$

$$z_{t,i} = z_{t,i}^N n_{t,i}^{Nz} + \sum_j z_{t,i}^{Dj} n_{t,i}^{Djz}$$

$$\forall i \in \mathcal{I}_t, \quad \forall j \in \{1, \dots, d\}, \quad \forall t \in \mathcal{T} \quad (13l)$$

$$u_{t,i} = u_{t,i}^N n_{t,i}^{Nu} + \sum_j u_{t,i}^{Dj} n_{t,i}^{Dju}$$

$$\forall i \in \mathcal{C}, \quad \forall j \in \{1, \dots, d\} \quad \forall t \in \mathcal{T} \quad (13m)$$

$$mg + \sum_i^{\mathcal{C}} u_{t,i} + \sum_i^{\mathcal{I}_t} z_{t,i} = 0 \text{ or } M_{trans} \dot{v}_{trans} \quad \forall t \in \mathcal{T} \quad (13n)$$

$$\sum_i^{\mathcal{C}} (q_t \cdot (c^{mnp,i} - CoM)) \times u_{t,i} +$$

$$\sum_i^{\mathcal{I}_t} (q_t \cdot (y_{t,i} - CoM)) \times z_{t,i} = 0 \text{ or } M_{rot} \dot{v}_{rot} \quad \forall t \in \mathcal{T} \quad (13o)$$

where $\mathcal{T} = \{0, \dots, T\}$, $\mathcal{I}_t = \{0, \dots, |\tilde{Y}_t|\}$, $\psi(q_t, v_t, y_{t,i})$ is the relative tangential velocity at contact point $y_{t,i}$. $\mathcal{C} = \{0, \dots, |c^{mnp}|\}$, $c^{mnp,i}$ is the i^{th} contact point between the robot and the object in the object's frame, CoM is the center of mass of the object, v_{trans} is the translational velocity of the object, M_{trans} is the translational part of the inertia matrix of the object with appropriate dimension, and v_{rot} and M_{rot} are the corresponding terms for the rotational part.