

# Contact-Implicit Model Predictive Control for Dexterous In-hand Manipulation: A Long-Horizon and Robust Approach

Yongpeng Jiang, Mingrui Yu, Xinghao Zhu, Masayoshi Tomizuka, and Xiang Li

**Abstract**—Dexterous in-hand manipulation is an essential skill of production and life. However, the highly stiff and mutable nature of contacts limits real-time contact detection and inference, degrading the performance of model-based methods. Inspired by recent advances in contact-rich locomotion and manipulation, this paper proposes a novel model-based approach to control dexterous in-hand manipulation and overcome the current limitations. The proposed approach has an attractive feature, which allows the robot to robustly perform long-horizon in-hand manipulation without predefined contact sequences or separate planning procedures. Specifically, we design a high-level contact-implicit model predictive controller to generate real-time contact plans executed by the low-level tracking controller. Compared to other model-based methods, such a long-horizon feature enables replanning and robust execution of contact-rich motions to achieve large displacements in-hand manipulation more efficiently; Compared to existing learning-based methods, the proposed approach achieves dexterity and also generalizes to different objects without any pre-training. Detailed simulations and ablation studies demonstrate the efficiency and effectiveness of our method. It runs at 20Hz on the 23-degree-of-freedom, long-horizon, in-hand object rotation task.

## I. INTRODUCTION

In-hand dexterous manipulation is an essential skill with a wide range of applications, such as turning doorknobs in the home service, operating all kinds of tools and machines on the production line, and so on. Such a skill is based on contact-rich manipulation, where the robot has to make and break contacts to interact with objects [1]. It is challenging to control the contact-rich motions because the non-smooth contact mechanism leads to intractable dynamics and stiff loss landscapes [2]. Although considerable effort has been devoted to the locomotion domain with great success [3], manipulation is quite different. For example, the robot must actively explore contacts instead of relying on gravity. Also, the number and locations of the contacts are constantly changing.

Existing works adopt both learning-based and model-based approaches to address these challenges. Despite the surprising robustness of learning, probably due to domain randomization and large-scale training, its data inefficiency slows down further deployment and generalization [4]–[8].

In contrast, model-based techniques provide a plug-and-play solution and can be divided into two categories. The first

Y. Jiang, M. Yu, and X. Li are with the Department of Automation, Beijing National Research Center for Information Science and Technology, Tsinghua University, China. X. Zhu and M. Tomizuka are with the Department of Mechanical Engineering, University of California, Berkeley, CA, USA. This work was supported in part by the National Key R&D Program of China under Grant 2022YFB4701403, in part by the National Natural Science Foundation of China under Grant U21A20517, 52075290, and 623B2059, and in part by the Institute for Guo Qiang, Tsinghua University. Corresponding author: Xiang Li (xiangli@tsinghua.edu.cn).




Fig. 1. Overview of the proposed method for long-horizon in-hand manipulation. The reference generator infers how to make and break contacts in the system’s neighborhood based on desired object motions. Then, the contact controller addresses modeling errors with tactile feedback and executes the inferred contacts to the fullest. Finally, joint-level impedance control is used to execute the contact-rich motions. The proposed method can seamlessly generalize among different objects.

category, the contact-explicit approach, controls the robot to make contacts at predefined locations [9]. This transforms the manipulation problem into the discrete search for contact sequences and the continuous optimization of control inputs [10]–[13]. Although this allows for robust execution on hardware with simple controllers, problems can arise when dealing with long-horizon tasks. This is because manipulation with specific contacts easily falls into the local optimum, and replanning new sequences is time-consuming. The second category takes a more elegant contact-implicit approach. It is common to use smooth surrogate models [2], [14] or relaxed complementary constraints [3] to efficiently explore possible contacts. Methods of this kind do not require specific contact sequences, and can adapt to potential contacts online instead. However, implicit methods always sacrifice physical fidelity for computational efficiency, thus create difficulties for long-horizon tasks [2], [3], as the system tends to evolve differently under unexpected deviation of contacts.

In this paper, we propose a novel control formulation for long-horizon in-hand manipulation, where the fingers alternately make and break contacts to achieve a large object displacement. The contact-rich motions of the multi-finger hand are regulated by a high-level reference generator and a low-level contact controller, shown in Fig. 1. Specifically, the proposed method is realized with a hierarchical structure:

- **At high-level**, we formulate a contact-implicit model predictive control (MPC) scheme to compute reference positions and contact forces given desired object motions. The controller runs a differential dynamic programming (DDP) algorithm on a smooth, differentiable quasi-

dynamic model.

- **At low-level**, we formulate a simple controller to track the high-level references. The controller incorporates a compliant contact model and utilizes tactile feedback.

The high-level module infers potential contacts in the surrounding area that result in the desired object motions. Simultaneously, the low-level module regulates the system around high-level references to ensure the inferred contacts are executed to the fullest. Compared to contact-explicit approaches, our method needs no pre-defined trajectories or separate planning procedures. Compared to contact-implicit approaches, our method has superior robustness since it addresses ubiquitous modeling errors. These characteristics enable our method to execute long-horizon manipulation online robustly. In addition, our method works with only trivial warm start solutions and can seamlessly generalize to different objects without any prior training required in learning-based approaches. Simulation results demonstrate the superior performance of the proposed method compared to existing approaches. Video and codes will be available soon on the project website <sup>1</sup>.

## II. RELATED WORKS

Dexterous in-hand manipulation refers to manipulating objects without a place-and-re-grasp process [15]. In this section, we discuss existing approaches to this problem. Several illuminating works from more general contact-rich locomotion and manipulation domains are included, while we omit innovations in purely mechanical design.

1) *With Reinforcement-Learning (RL) or Imitation-Learning (IL)*: Recently, learning-based approaches have shown excellent performance and stability in dexterous in-hand manipulation [16]. RL-based and IL-based methods are the mainstream. The empirical success of RL can be attributed to its inner stochasticity [2], domain randomization, flexible choice of sensing modalities [5], and large-scale training with memory [6]. In contrast, IL-based methods require less exploration of the action space but rely heavily on expert demonstrations [17]. Some methods use model-driven learning [16], [18] or explore the combination with non-learning approaches [4], [7], [8]. Nevertheless, learning-based methods rely mainly on data to understand the mechanism of in-hand manipulation, making it difficult to capture the underlying commonalities from a theoretical perspective. Therefore, such methods may perform poorly for unseen objects or in the presence of significant sim2real gaps.

2) *Model-Based Design Using Explicitly Represented Contacts*: Model-based techniques provide a solution for deployment and generalization without training. In dexterous in-hand manipulation, planning and control commonly involve explicit representation of contacts, including their locations, modes, reaction forces, and so on [11]. These items form the contact sequences. Existing works use different methods to obtain such discrete sequences, such as searching [10], [12], [13], sampling [11], [19], or demonstration [9]. The subsequent

procedure is simplified to solving for the continuous control inputs, such as joint torques, while following pre-defined contact sequences [10]. Thus, low-level control typically involves solving tractable optimization problems with force closure and non-sliding constraints. While these methods are easier and more precise to implement on hardware, acquiring contact sequences adds complexity and degrades real-time performance. Moreover, if contact states deviate from pre-defined ones due to disturbance, the specific contact sequence can become misleading information.

3) *Model-Based Design Using Implicitly Encoded Contacts*: Several works avoid the increasing complexity of explicit contact representation using relaxed complementary constraints [3], [20], smooth surrogate models [2], [21], or direct sampling of the control [22]. Aydinoglu et al. [23] proposed an operator splitting framework to decouple variables across time steps and accelerate the intractable Linear Complementary Problem (LCP). Smoothing is discovered as an important procedure to allow efficient exploration among contact modes and to avoid falling into the local minimum [3], [20]. Pang et al. [2] proved that randomized and analytic smoothing are equally important for dexterous manipulation planning. Another method of obtaining smooth dynamics is through compliant contact models [21]. Inspired by the Hunt and Crossley model, Kurtz et al. [14] proposed an alternative version with regularized friction and propagated the gradients through inverse dynamics. Although their method works well in the dexterous manipulation domain, the performance can be easily affected by model parameters. Although contact-implicit methods improve efficiency by eliminating the dependence on specific contact sequences, the smoothing procedure introduces the side effect of force-at-a-distance [2]. As a result, these methods often fail for even small discrepancies between actual and planned contact modes. In contrast, our method is more robust due to contact feedback.

## III. METHOD

This paper considers the quasi-dynamic manipulation through rigid frictional contacts [24]. The object being manipulated is modeled as a single rigid body, while the manipulator can be any multi-link system, e.g., a multi-finger hand. The task is defined as moving the object to follow reference motions. The method is illustrated in Fig. 2.

‘Quasi-dynamic’ refers to manipulation under negligible inertia effects, in which the momentum does not accumulate and the system can be described with first-order dynamics [2], [25]. We make such an assumption because it accords with many objects of daily life due to frictional surfaces, damped hinges, or slow motions. In this paper, we use subscripts  $o$ ,  $r$  for object and robot, and  $\|\cdot\|_W$  for weighted norm unless otherwise stated. Besides,  $[a; b]$  denotes the vertical concatenation of vectors  $a$ ,  $b$ , and  $[A; B]$  denotes the vertical stacking of matrices  $A$ ,  $B$ .

### A. Contact-Implicit Reference Generator

One key issue to achieve long-horizon in-hand manipulation is to auto-generate contact references online, i.e., hand

<sup>1</sup>[https://director-of-g.github.io/in\\_hand\\_manipulation/](https://director-of-g.github.io/in_hand_manipulation/)

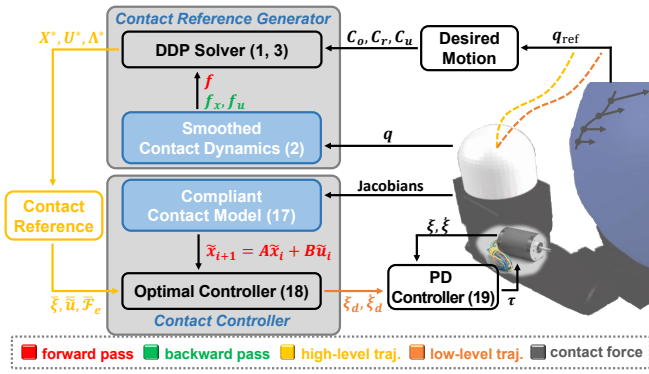


Fig. 2. **Block diagram of the proposed method.** The long-horizon contact-rich manipulation is accomplished by the contact reference generator and the contact controller working as a hierarchical structure.

configurations, contact locations, reaction forces, and so on for the low-level controller. The process accepts only desired object motion as input without pre-defined sequences or trajectories.

1) *Problem Formulation:* The generation of references is described as a finite horizon optimal control problem.

$$\begin{aligned} \min_{U=\{u_0, \dots, u_{N-1}\}} & \sum_{i=0}^{N-1} l(x_i, u_i) + l_f(x_N) \\ \text{s.t.} & \quad x_{i+1} = f(x_i, u_i) \\ & \quad u \preceq u_i \preceq \bar{u}, \quad i = 0, \dots, N-1 \end{aligned} \quad (1)$$

For quasi-dynamic manipulation, the state variables only include generalized positions, i.e.,  $x = q$ , with the dimension of  $n_q$ . Here  $q$  includes both the robot and the object, i.e.,  $q = [q_r; q_o]$ . Besides, control inputs are chosen as  $u = S_r \delta q$  with the same bounds  $\underline{u}, \bar{u}$  at different time steps, where  $S_r$  selects the actuated dimensions. We expect the dynamics  $f$  to 1) satisfy quasi-dynamic assumptions, 2) have smooth forward and backward passes, and 3) evaluate efficiently. Property 2) is indicated in [3], where smoothed gradients enable efficient transition among different contact modes. Moreover, we illustrate that unlike the exact integration as in [3], the smooth forward pass is also indispensable for contact-implicit optimization in the manipulation domain.

2) *Differentiable Dynamics Through Contact:* For the dynamics mapping  $f$ , we use the convex, quasi-dynamic, differentiable contact (CQDC) model [2], which has the properties mentioned above. The model is exactly an optimization-based dynamics that relaxes the nonlinear contact mechanism as a Cone Complementary Problem (CCP) [26]. Assuming there are  $N_c$  point contacts with known local geometry and friction, the discrete-time dynamics can be formulated as:

$$\begin{aligned} \min_v & \quad \frac{1}{2} v^\top Q v + b^\top v \\ \text{s.t.} & \quad \mu_i \|J_{i,t} v\|_2 \leq J_{i,n} v + \phi_i/h, \quad i = 1, \dots, N_c \end{aligned} \quad (2)$$

where  $h$  is the discrete time step,  $v = \delta q/h$ ,  $\phi_i, \mu_i$  are the signed distance and friction coefficient of the  $i^{\text{th}}$  contact, and  $J_i = [J_{i,n}; J_{i,t}]$  is the contact Jacobian. Besides,  $Q$  and  $b$  are known and can be computed with the object's inertia, the robot's joint stiffness and generalized forces applied on both. We refer the readers to [2] for a more detailed derivation.

The original problem (2) is a second-order cone program [27], which is convex, differentiable, but non-smooth in nature. We found that incorporating non-smooth dynamics and its gradients in the optimal control problem (1) always results in poor convergence. Thus we adopt the analytically smoothed version of (2) and explain how the control performance is influenced by the degree of smoothness in Sec IV. The dual variable  $\lambda$  of (2) stands for the concatenated contact forces.

Finally, the smoothed dynamics  $f$  and differential terms  $f_x, f_u$  are obtained with the open-sourced implementation of [2], which relies on the differentiable collision detection from [28]. For the same reason as [2], only simple geometries with differentiable boundaries are used in the experiments, i.e., spheres, capsules. The CQDC model is also integrated into a numerical optimization framework in [2] but for open-loop planning. However, close-loop control is further considered in this paper.

3) *Cost Function Design:* The running and terminal cost terms are composed of three parts

$$\begin{aligned} l(x, u) &= C_o(x) + C_r(x) + C_u(u) \\ l_f(x_N) &= C_o(x_N) + C_r(x_N) \end{aligned}, \quad (3)$$

the subscripts are omitted for brevity. The regulation costs

$$\begin{aligned} C_o(x) &= \|q_o - q_{o,\text{ref}}\|_{W_o}^2 \\ C_r(x) &= \|q_r - q_{r,\text{ref}}\|_{W_r}^2 \end{aligned} \quad (4)$$

are the keys to generate contact-rich motions, where  $q_{\cdot,\text{ref}}$  are references and  $W_{\cdot}$  are weight matrices. Among them,  $C_o$  guarantees desired object motions, and  $C_r$  encourages the periodic finger-gaiting by penalizing large finger displacements from nominal configurations. This technique is widely applied in many model-based or learning-based frameworks that include rich contacts and periodic motions [3], [29]. The negligence of  $C_r$  often causes the system to fall into local optimum, i.e., the hand fails to break contacts actively. In practice, the nominal finger configuration  $q_{o,\text{ref}}$  can be fixed as any grasping pose (i.e., decided empirically or generated by a grasping sampler) in most cases. Besides, we find it helpful to increase  $W_r$  near the end of the receding horizon.

We also find that considering control regulation terms improves numerical conditions and accelerates convergence.

$$C_u(u) = \|u\|_{W_u}^2 \quad (5)$$

Note that the framework allows other types of costs as long as they are differentiable. For example, additional costs can be introduced to encourage symmetric motions of certain fingers and to enhance contact clearance or maintenance.

4) *Generation of Contact References:* With differentiable dynamics (2) and costs (3), the optimal control problem (1) can be efficiently solved with the control-limited version of DDP [30]. As suggested by [3], we omit the second order terms  $f_{xx}, f_{xu}, f_{uu}$  since they contain derivatives with higher orders than kinematic Hessian, and have weaker effects on convergence.

Finally, the solved trajectories  $X^* = \{x_0^*, \dots, x_N^*\}, U^* = \{u_0^*, \dots, u_{N-1}^*\}, \Lambda^* = \{\lambda_0^*, \dots, \lambda_{N-1}^*\}$  are converted to




Fig. 3. **Modeling of the contact controller.** The frames  $O_r, O_c, O_d$ , frame displacements  $dp_{c,r}, dp_{r,d}$ , contact jacobians  $G_i, J_i$ , stiffness matrices  $K_c, K_r, K_o$  are illustrated. Note that  $O_d$  can be obtained with forward kinematics  $FK(\xi_d)$ . The model compliance creates a coupling effect between force and motion.

desired robot configurations and contact forces, and are interpolated before being sent to the low-level controller.

### B. Compliance-Based Contact Controller

The smoothed dynamics has side-effects that shift the object at a non-zero contact distance (i.e., force-at-a-distance). Thus directly executing the high-level references would lead to contact slippage or missing. To mitigate such modeling errors, a model-based controller is proposed to adapt the generated references locally. The controller incorporates an equivalent virtual spring system with a compliant contact model. By coupling joint positions and reaction forces with compliance, a balance is achieved between tracking both. The modeling is illustrated in Fig. 3.

1) *Compliant Contact Model:* Any compliant contact model that maps relative contact distances to reaction forces can be used. One example is [14], while the damping effect related to contact velocity is ignored here due to quasi-static assumptions. The compliant contact model is guaranteed to have continuous gradients. Thus, we can define the equivalent stiffness in the contact frame,

$${}^c K_c = \text{diag} \left( \frac{\partial f}{\partial \phi} [\alpha \quad \alpha \quad 1] \right) \quad (6)$$

where the y-axis aligns with the sliding direction, and the z-axis coincides with the contact normal,  $\alpha$  is an adjustable parameter. Note that the CQDC model is not used at low-level due to its slow, optimization-based evaluation.

2) *Contact Force-Motion Model:* We make the following assumptions: 1) quasi-dynamic manipulation, and 2) the contact stiffness and Jacobian matrices are evaluated only once at the beginning of each control step. Unless otherwise stated, we omit the subscript  $i$  since the following derivation is the same for each contact. We extend the formulation in [31] to multi-contact cases.

Under quasi-dynamic assumptions, the contact force exerted by each finger can be written as

$$F_e = K_c dp_{c,r} \quad (7)$$

where  $K_c = R^C K_c R^\top$ ,  $R$  is the rotation matrix. Besides,  $dp_{c,r}$  is the displacement between the nearest points on the

object and robot. We attach a virtual spring at each contact with the stiffness matrix (6), and the equivalent stiffness  $K_o \in \mathbb{R}^{6 \times 6}$  can be obtained as

$$K_o = \sum_{i=1}^{N_c} G_i K_{c,i} G_i^\top \quad (8)$$

where  $G_i \in \mathbb{R}^{6 \times 3}$  equals the grasping matrix of the  $i^{\text{th}}$  contact. The stacked matrix  $\mathcal{G} \in \mathbb{R}^{6 \times 3N_c}$  is defined as

$$\mathcal{G} = [G_1, \dots, G_{N_c}] \quad (9)$$

For clarity,  $\xi = q_r$  is used in this section. According to (19), with gravity compensation, it is obtained that

$$K_r(\xi_d - \xi) + D_r(\dot{\xi}_d - \dot{\xi}) = \mathcal{J}^\top \mathcal{F}_e \quad (10)$$

where  $\mathcal{J}$  is the stacking of contact Jacobians and  $\mathcal{F}_e$  is the concatenated contact forces. Typically we choose  $K_r = k_p I$ ,  $D_r = k_d I$  as diagonal matrices. Using  $J(\xi_d - \xi) \approx dp_{r,d}$  and ignoring the velocity term, we can obtain

$$k_p dp_{r,d} = J \mathcal{J}^\top \mathcal{F}_e \quad (11)$$

If we substitute (7) and (11) into

$$dp_{c,d} = dp_{c,r} + dp_{r,d} \quad (12)$$

and left-multiply  $K_c$ , we have

$$F_e + k_p^{-1} K_c J \mathcal{J}^\top \mathcal{F}_e = K_c dp_{c,d} \quad (13)$$

Note that

$$\begin{aligned} d\dot{p}_{c,d} &= \dot{p}_d - \dot{p}_c \\ &= J \dot{\xi}_d - G^\top K_o^{-1} \mathcal{G} \mathcal{F}_e \end{aligned} \quad (14)$$

Substituting (13) into (14) for all contacts, we can obtain that

$$\dot{\mathcal{F}}_e \approx \left( I + k_p^{-1} M \mathcal{J}^\top + N K_o^{-1} \mathcal{G} \right)^{-1} M \dot{\xi}_d \quad (15)$$

where  $M \in \mathbb{R}^{3N_c \times n_{qr}}$  is the stacking of  $K_c J$  and  $N \in \mathbb{R}^{3N_c \times 6}$  is the stacking of  $K_c G^\top$ . From (10) it can be obtained that

$$\dot{\xi} = \dot{\xi}_d + k_d^{-1} \left( k_p(\xi_d - \xi) - \mathcal{J}^\top \mathcal{F}_e \right) \quad (16)$$

With (15) and (16), if we choose desired joint velocity as the control input  $\tilde{u} = \dot{\xi}_d$ , the second-order contact dynamics can be written as (and can be further discretized)

$$\begin{bmatrix} \dot{\xi} \\ \dot{\xi}_d \\ \mathcal{F}_e \end{bmatrix} \triangleq \begin{bmatrix} \tilde{u} + k_d^{-1} \left( k_p(\xi_d - \xi) - \mathcal{J}^\top \mathcal{F}_e \right) \\ \tilde{u} \\ \left( I + k_p^{-1} M \mathcal{J}^\top + N K_o^{-1} \mathcal{G} \right)^{-1} M \tilde{u} \end{bmatrix} \quad (17)$$

Note that (17) describes the linear dynamics of  $\tilde{x} = [\xi; \xi_d; \mathcal{F}_e]$ . In practice, one can omit terms including  $K_o$  when the object is not free-floating. The low-level controller can be formulated as the following optimal control problem

$$\begin{aligned} \min_{\tilde{U}} & \|\xi_N - \bar{\xi}_N\|_{W_{\xi_N}} + \|\mathcal{F}_{e,N} - \bar{\mathcal{F}}_{e,N}\|_{W_{F_N}} \\ & \sum_{i=0}^{N-1} \|\xi_i - \bar{\xi}_i\|_{W_{\xi}} + \|\mathcal{F}_{e,i} - \bar{\mathcal{F}}_{e,i}\|_{W_F} + \|\tilde{u}_i - \bar{\tilde{u}}_i\|_{W_{\tilde{u}}} \\ \text{s.t.} & \tilde{x}_{i+1} = A \tilde{x}_i + B \tilde{u}_i, \quad i = 0, \dots, N-1 \end{aligned} \quad (18)$$




Fig. 4. Two in-hand manipulation systems used in the experiments. The object being manipulated is painted in blue. (a) The orange cylinder represents a damped hinge. (b) The brown body represents the supporting plane.

where the references are calculated from interpolating the generated references  $\mathbf{X}^*$ ,  $\mathbf{\Lambda}^*$ ,  $\mathbf{U}^*$ .

We use the joint-level PD controller with gravity compensation to convert the joint references  $\xi_d$ ,  $\dot{\xi}_d$  into joint torques.

$$\boldsymbol{\tau} = \boldsymbol{\tau}_g + \mathbf{K}_r(\boldsymbol{\xi}_d - \boldsymbol{\xi}) + \mathbf{D}_r(\dot{\boldsymbol{\xi}}_d - \dot{\boldsymbol{\xi}}) + \mathcal{J}^\top \mathcal{F}_{e,ff} \quad (19)$$

where  $\boldsymbol{\tau}_g$  is the generalized gravity. In addition, we add the feed-forward torque used in baseline methods and neglect inertia and Coriolis effects under quasi-dynamic assumptions.

#### IV. RESULTS

We carry out simulations and ablation studies to validate the proposed method. We use the CQDC model implementation from the quasi-static simulator [32] proposed by Pang et al. The DDP algorithm is implemented using Python bindings of the Crocoddyl library [33]. Simulations of the complete system are performed using Drake toolbox [28], and the code base of [14] with hydroelastic contacts.

Two types of in-hand manipulation systems are used to evaluate the proposed method, as shown in Fig. 4. We use a four-finger dexterous hand Allegro Hand with a fixed wrist and 16 actuated joints. The object to be manipulated could either rotate perpendicular to the palm (**RotZ**,  $n_q = 17$ ) or move freely while being supported by the table (**Free**,  $n_q = 23$ )

##### A. Long-Horizon Motion Generation with Rich Contacts

We use  $\kappa = 100$  (the same notation as [2]) to analytically smooth (2) throughout the paper, where a smaller value indicates a stronger smoothing effect. We set the input bounds as 0.05 rad for the **RotZ** system and 0.2 rad for the **Free** system, with the discretization step  $h = 0.1$  s and the horizon length of 10.

First, we demonstrate the importance of using smoothed forward dynamics for contact-rich manipulation. We command




Fig. 5. The relationship between  $\kappa$  and optimization results. The cost of DDP, the object's yaw angle, and the control input norms are plotted for 50 steps. Warmer colors indicate a stronger smoothing effect.




Fig. 6. The finger-gaiting of the in-hand manipulation on two different objects, with a duration of 200 iterations. Shaded areas indicate contacts of each finger. The object yaw angle curves are plotted in grey.

the object to rotate at 0.79 rad/s and change  $\kappa$  in the forward dynamics computation. As shown in Fig. 5, larger  $\kappa$  leads to poorer convergence and less effective manipulation of the object. The yaw angles increase rapidly when  $\kappa \geq 1000$  due to initial finger motions solved from warm-start trajectories. However, they quickly evolve into vibrating motions or even stop moving after 20 iterations without enough smoothing. Smoothed forward dynamics ( $\kappa \leq 500$ ) enables the finger-gaiting generation and achieves constant rotation speed instead of getting stuck in a local optimum. Thus, we prove it impossible to remedy high-level modeling errors through a more realistic forward pass. Hence, a low-level controller is necessary.

Second, we test the generalization ability on different objects. As shown in Fig. 1 and Fig. 6, a sphere with the radius of 6 cm and a cuboid with the side length of 7 cm are manipulated in the **RotZ** environment. The shaded areas indicate contact events ( $\phi < 3$  mm) of four fingers, where the finger-gaiting is auto-generated to accomplish the long-horizon manipulation. The sphere rotates at 0.29 rad/s and the cuboid rotates at 0.4 rad/s. Our method seamlessly generalizes to different geometries without pre-training, which shows a potential advantage over learning-based approaches. See the attached video for details.

Third, we further evaluate the robustness of our method under external disturbances or internal noises, which are assumed to obey Gaussian distributions, with the standard deviation  $\sigma$ . The sphere is commanded to rotate around the vertical axis in the **Free** system. To simulate the disturbances, we apply numerical perturbations once to the Euler angles




Fig. 7. The rotation and position error curves with the increasing amplitude of noise/disturbance. The mean values (solid lines) and half standard deviation ranges (shaded areas) are shown.

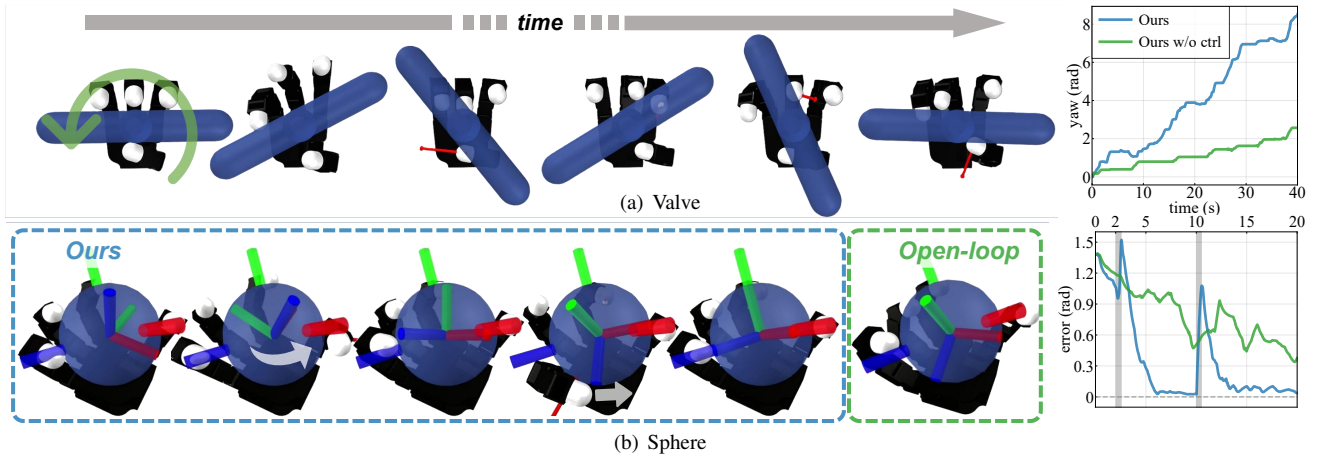


Fig. 8. Snapshots of the long-horizon (a) valve rotation task and (b) sphere rotation task carried out in the Drake simulator. Red arrows represent the contact forces. (a) The blue capsule represents the valve and rotates around the vertical axis. The fingers perform side-pushes and step over the valve in succession. (b) External disturbances are applied on the object and ring finger, as the white arrows indicate. The final state of the open-loop baseline is also shown. Object yaw angle and error curves are shown on the right.

TABLE I

CONTROLLER PERFORMANCE IN THE ROTZ SYSTEM.

Methods	Average slippage ( $\times 10^{-3}$ )	Average rotation speed (rad/s)	Average joint velocity (rad/s)
Ours	<b>3.26</b>	0.30	0.93
Ours w/o ctrl	8.04	0.05	0.83
FF+torque	9.47	0.85	1.17
FF+pos	8.80	0.85	1.18

and change  $\sigma$  (i.e., the magnitude) from 0.1 rad to 0.8 rad. To simulate the noises, at every time step, we perturb the dimensions of hand, object position, and quaternion in  $\mathbf{x}$ , with  $\sigma = 0.03$  rad, 1 mm, 0.01, respectively. We define the amplitude as the multiples of  $\sigma$ . We run 100 random experiments for each parameter and record the average tracking error within 400 steps. As shown in Fig. 7, our method could recover from various disturbances and noises. Robustness is essential for long-horizon manipulation under sensor noises, control delays, and external disturbances. The orientation error explodes under large disturbances since the system needs a longer time for recovery.

### B. Robust Execution with Contact Controller

1) *Simple tasks*: We choose three baselines for comparison. The first one executes generated trajectories in an open-loop fashion (Ours w/o ctrl), which resembles the motion planning in [2] with online re-planning. The second one utilizes feed-forward torques  $\boldsymbol{\tau}_{ff} = \boldsymbol{J}^T \boldsymbol{F}_{e,ff}$  as (19) (FF+torque), which is the quasi-dynamic variant of [14]. The third one maps feed-forward torques to desired positions, where  $\mathbf{q}_d$  is replaced by  $\mathbf{q}_d + \boldsymbol{\tau}_{ff}/\mathbf{K}_{ctrl}$  in (19) (FF+pos), as suggested by [3]. As for our feedback design, low-cost tactile sensors [34] can be used to obtain contact forces in hardware. A slippage metric is adopted from [3] to indicate the occurrence of missed or ineffective contacts.

We simulate the **RotZ** system in Drake for 10.0s with the four controllers. The reference generator outputs references at 20 Hz, and the contact controller runs at 100 Hz. We record the average slippage when  $\phi < 1$  mm, the object's average rotation speed, and the hand's average speed in joint space.

As depicted in Table I, our controller achieves the lowest average slippage by maintaining enough desired force. As a result, the average rotation speed increases significantly compared with open-loop execution, which indicates that more effective contacts are established with our method. The two feed-forward controllers perform similarly and rotate the sphere faster with higher slippage. This is because the feed-forward torques are converted into displacements due to joint impedance. Thus, feed-forward controllers often result in larger deviations from planned motions and faster joint velocities once the contact is lost, as seen in the last column of Table I. In contrast, our controller balances tracking planned motions and making contacts, thus generating more stable motions.

2) *Complex tasks*: We further evaluate the contact controller and the proposed method in more complex cases. These experiments also show potential applications in the real world.

The first case is rotating the valve, where  $n_q = 17$ . In this case, the fingers should switch between pushing aside and stepping over the valve to accomplish the long-horizon manipulation, shown in Fig. 8(a). To our knowledge, such long-horizon manipulation is almost only studied in the machine-learning literature. We find it helpful to interpolate the nominal finger configuration  $\mathbf{q}_{o,ref}$  between two grasping poses of different fingertip heights. Otherwise, a longer horizon length is needed to avoid local optimums, which degrades control frequency. As shown in Fig. 8(a), the fingers sometimes need to exert forces below the valve, which is sensitive to contact loss. Our proposed contact controller effectively avoids such problems, as shown in the curves and the attached video.

The second case is rotating the sphere in place, where  $n_q = 19$ . The sphere should be aligned with the target orientation in  $SO(3)$ . With our proposed controller, we exert spatial force disturbances on the object during 2.0s  $\sim$  2.5s and on the ring finger during 10.0s  $\sim$  10.5s. As shown in Fig. 8(b), the long-horizon manipulation fails without the contact controller due to contact loss and unintentional contact.

Hence, the importance of our contact controller is proved. We also simulate the **Free** system in Fig. 4(b) and successfully rotate the sphere for more than  $270^\circ$ . However, the rolling contacts result in highly dynamic movements of the object, which violates the quasi-dynamic assumption and degrades the performance. We temporarily increase the damping of the object. As suggested by [2], designing special high-level objectives and enforcing the system to be quasi-dynamic at a low level could help to solve this problem. We leave them for future work.

## V. CONCLUSIONS

This paper discusses the task of long-horizon in-hand manipulation. First, the reference generator computes joint trajectories and reaction forces based solely on desired motions. Then, the contact controller, which includes tactile feedback, tracks joint movements and maintains planned contacts even in the presence of modeling errors. The experiments demonstrate that the proposed method achieves similar dexterity to learning-based approaches and generalizes to different objects without pre-training. With the reference generator, our method achieves long-horizon manipulation without a separate planning procedure by recalculating trajectories in real-time. In addition, the contact controller prevents contact loss and enhances robustness to disturbances. Future work will focus on hardware implementation with tactile sensors and improving manipulation performance through hand-arm coordination.

## REFERENCES

- [1] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, 2019.
- [2] T. Pang, H. J. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on Robotics*, vol. 39, pp. 4691–4711, 2022.
- [3] G. Kim, D.-K. Kang, J. ha Kim, S. Hong, and H.-W. Park, "Contact-implicit mpc: Controlling diverse quadruped motions without pre-planned contact modes or trajectories," *ArXiv*, vol. abs/2312.08961, 2023.
- [4] W. Jin and M. Posa, "Task-driven hybrid model reduction for dexterous manipulation," *IEEE Transactions on Robotics*, pp. 1–20, 2024.
- [5] M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. W. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, pp. 20 – 3, 2018.
- [6] T. Chen, M. Tippur, S. Wu, V. Kumar, E. H. Adelson, and P. Agrawal, "Visual dexterity: In-hand reorientation of novel and complex object shapes," *Science Robotics*, vol. 8, 2022.
- [7] H. Bui and M. Posa, "Enhancing task performance of learned simplified models via reinforcement learning," *ArXiv*, vol. abs/2310.09714, 2023.
- [8] J. Xu, V. Makovychuk, Y. Narang, F. Ramos, W. Matusik, A. Garg, and M. Macklin, "Accelerated policy learning with parallel differentiable simulation," in *International Conference on Learning Representations*, 2021.
- [9] F. Khadivar and A. Billard, "Adaptive fingers coordination for robust grasp and in-hand manipulation under disturbances and unknown dynamics," *IEEE Transactions on Robotics*, vol. 39, pp. 3350–3367, 2023.
- [10] C. Chen, P. Culbertson, M. Lepert, M. Schwager, and J. Bohg, "Trajectorytree: Trajectory optimization meets tree search for planning multi-contact dexterous manipulation," *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8262–8268, 2021.
- [11] X. Cheng, S. Patil, Z. Temel, O. Kroemer, and M. T. Mason, "Enhancing dexterity in robotic manipulation via hierarchical contact exploration," *IEEE Robotics and Automation Letters*, vol. 9, pp. 390–397, 2023.
- [12] H. Zhu and L. Righetti, "Efficient object manipulation planning with monte carlo tree search," *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10628–10635, 2022.
- [13] S. Cruciani, C. Smith, D. Kragic, and K. Hang, "Dexterous manipulation graphs," *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2040–2047, 2018.
- [14] V. Kurtz, A. Castro, A. Ö. Önel, and H. Lin, "Inverse dynamics trajectory optimization for contact-implicit model predictive control," *ArXiv*, vol. abs/2309.01813, 2023.
- [15] S. Cruciani, B. Sundaralingam, K. Hang, V. Kumar, T. Hermans, and D. Kragic, "Benchmarking in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, pp. 588–595, 2020.
- [16] A. I. Weinberg, A. Shirizly, O. Azulay, and A. Sintov, "Survey of learning approaches for robotic in-hand manipulation," *ArXiv*, vol. abs/2401.07915, 2024.
- [17] Y. Qin, H. Su, and X. Wang, "From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation," *IEEE Robotics and Automation Letters*, vol. 7, pp. 10873–10881, 2022.
- [18] V. Kumar, E. Todorov, and S. Levine, "Optimal control with learned local models: Application to dexterous manipulation," *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 378–383, 2016.
- [19] N. C. Daffe, R. Holladay, and A. Rodriguez, "In-hand manipulation via motion cones," *ArXiv*, vol. abs/1810.00219, 2018.
- [20] S. L. Cleac'h, T. A. Howell, S. Yang, C.-Y. Lee, J. Zhang, A. L. Bishop, M. Schwager, and Z. Manchester, "Fast contact-implicit model predictive control," *IEEE Transactions on Robotics*, 2021.
- [21] A. Ö. Önel, P. Long, and T. Padir, "Contact-implicit trajectory optimization based on a variable smooth contact model and successive convexification," *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2447–2453, 2018.
- [22] T. A. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa, "Predictive sampling: Real-time behaviour synthesis with mujoco," *ArXiv*, vol. abs/2212.00541, 2022.
- [23] A. Aydinoglu, A. Wei, and M. Posa, "Consensus complementarity control for multi-contact mpc," *ArXiv*, vol. abs/2304.11259, 2023.
- [24] M. T. Mason, "Mechanics of robotic manipulation," 2001.
- [25] Y. Jiang, Y. Jia, and X. Li, "Contact-aware non-prehensile manipulation for object retrieval in cluttered environments," *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10604–10611, 2023.
- [26] Q. L. Lidec, W. Jallet, L.-R. Montaut, I. Laptev, C. Schmid, and J. Carpentier, "Contact models in robotics: a comparative analysis," *ArXiv*, vol. abs/2304.06372, 2023.
- [27] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge Univ. Pr., 2011.
- [28] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: <https://drake.mit.edu>
- [29] H. Qi, B. Yi, Y. Ma, S. Suresh, M. Lambeta, R. Calandra, and J. Malik, "General In-Hand Object Rotation with Vision and Touch," in *Conference on Robot Learning (CoRL)*, 2023.
- [30] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1168–1175, 2014.
- [31] T. Gold, A. Völz, and K. Graichen, "Model predictive interaction control for robotic manipulation tasks," *IEEE Transactions on Robotics*, vol. 39, pp. 76–89, 2023.
- [32] T. Pang, "A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction," *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6614–6620, 2021.
- [33] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [34] M. Lambeta, P. wei Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, and R. Calandra, "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, pp. 3838–3845, 2020.