

Demonstration-guided Optimal Control for Long-term Non-prehensile Planar Manipulation

Teng Xue^{1,2}, Hakan Girgin^{1,2}, Teguh Santoso Lembono^{1,2}, and Sylvain Calinon^{1,2}

Abstract—Long-term non-prehensile planar manipulation is a challenging task for robot planning and feedback control. It is characterized by underactuation, hybrid control, and contact uncertainty. One main difficulty is to determine both the continuous and discrete contact configurations, e.g., contact points and modes, which requires joint logical and geometrical reasoning. To tackle this issue, we propose a demonstration-guided hierarchical optimization framework to achieve offline task and motion planning (TAMP). Our work extends the formulation of the dynamics model of the pusher-slider system to include separation mode with face switching mechanism, and solves a warm-started TAMP problem by exploiting human demonstrations. We show that our approach can cope well with the local minima problems currently present in the state-of-the-art solvers and determine a valid solution to the task. We validate our results in simulation and demonstrate its applicability on a pusher-slider system with a real Franka Emika robot in the presence of external disturbances. Project webpage: <https://sites.google.com/view/dg-oc/>.

I. INTRODUCTION

With the trend of labor shortage and aging population, robots are required for increasingly complicated tasks, moving beyond the typical pick-and-place tasks to non-prehensile manipulation, which refers to manipulating without grasping but with mutual interaction. This requires real-time contact identification and adaptation. In this paper, we focus on long-term non-prehensile planar manipulation, which concerns achieving reliable non-prehensile planar manipulation with a long-term horizon, involving joint logistic and geometric planning and feedback control over diverse interaction modes and face switches. For example, to push an object, a prerequisite is to decide how much force should be applied, and which point to push. Moreover, in some cases such as pushing an object with small distance but large orientation, relying on a single fixed face is not feasible. Therefore, a sequence of face switching is required, as well as contact mode schedule resulting from Coulomb friction.

To achieve non-prehensile planar manipulation, four main challenges appear:

1. **Hybrid Dynamics.** The dynamics of a pusher-slider system depends on the current interaction mode and contact point. In this paper, we consider not only the motion involving various interaction modes, i.e., separation, sticking, sliding up, and sliding down, but also the switching between the contact faces, i.e., left, bottom, right, and up (unlike




Fig. 1: System setup where the robot pushes an object with changes of contact points.

- previous works [1], [2] that only work with a single face).
2. **Underactuation.** The contact force between the pusher and the slider is constrained within a motion cone, making it impossible to exert arbitrary acceleration on the object to achieve omnidirectional movement.
3. **Long-horizon TAMP.** Due to the characteristics of hybrid dynamics and underactuation, it is important to reason over long horizon using both logic and geometric descriptors, where the logic variables relate to contact modes and faces and the geometric variables relate to contact points, slider states, switching points, and control inputs.
4. **Contact Uncertainty.** Arising from the frictional contact interactions between the pusher and the slider, as well as between the slider and the table, the contact is hard to be modeled precisely, therefore a controller enabling online contact adaptation is required.

The existence of both continuous and discrete variables characterizes the hybrid nature of contact, making it challenging for gradient-based optimization algorithms. Although optimal control has demonstrated good results for contact-rich tasks through techniques such as contact smoothing [3] and the inclusion of complementary constraints [2], it is still susceptible to poor local minima. On the other hand, sampling-based methods like search tree [4] have the ability to perform long-term global planning, but can struggle with the vast number of potential contact modes. Human decision-making, conversely, excels at discrete decisions. Integrating human demonstration into the hybrid optimization framework could greatly enhance its convergence towards near-optimal solutions.

The main contribution of our paper is a hybrid framework

*This work was in part supported by the China Scholarship Council (grant No.202106230104).

¹Idiap Research Institute, Martigny, Switzerland (e-mail: firstname.lastname@idiap.ch)

²École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

that integrates optimization and learning from human demonstrations to overcome the issue of getting trapped in poor local optima, which is a common problem in many state-of-the-art solvers. Specifically, we demonstrate the effectiveness of our approach in a challenging long-term non-prehensile manipulation task that involves multiple interaction modes and face-switching cases. We collect demonstrations through an interface¹, where humans are required to guide the robot to push the block toward the target. The discrete decision variables are represented implicitly by some continuous variables. We then use the collected demonstration data as soft constraints in the cost function for optimization, enabling efficient exploration of the solution space that includes both logic and geometric variables. Additionally, we propose a real-time feedback controller that accounts for model mismatch and contact uncertainty when interacting with the physical world. It is worth noting that our approach prioritizes motion planning in the task space and then utilizes operational space control [5] to compute joint torques based on end-effector motion.

II. RELATED WORK

In this section, we discuss some related work on dynamics modeling, TAMP, and robot learning from demonstration.

Non-prehensile manipulation has been widely studied as a challenging task for model-based planning and control [6], with the pusher-slider system as one of the most prominent examples. Under the assumption of quasi-static interaction [7] and ellipsoidal limit surface [8], Lynch *et al.* put forward a set of analytical methods to compute velocities based on the motion cone [9], which was previously studied in [10] to distinguish different sticking and sliding behaviors with a friction cone. Recently, Hogan *et al.* reformulated the motion equation as a piecewise function to describe hybrid dynamics [1]. In [11], they instead expressed the limit surface in convex quadratic form to simplify the previous equation as a more general form. Nevertheless, the motion equations mentioned above are all assuming that the contact face is fixed, which is too constrained for long-term manipulation, since the pusher will lose the chance to adjust the contact point and face during pushing. Therefore, an extended dynamics model is required to provide the pusher with more flexible choices on the contact configuration.

A hybrid framework with planning and feedback control is important to achieve long-term non-prehensile planar manipulation. Optimally unifying hybrid interaction modes and geometric variables is still an open challenge. For instance, in [1], a mixed integer programming (MIP) was introduced into a model predictive controller (MPC) to incorporate the selection of interaction modes, i.e., sticking and sliding, into a common optimization framework. Due to the non-convex nature of integer variables, the computation can be time-consuming and not feasible for high-frequency feedback controller. An offline 3-layer mode sequence classifier was later proposed in [11]. A higher number of discrete modes

and contact faces would require the training dataset to increase exponentially, leading to more time-consuming data generation process. Moreover, the works mentioned above are both based on short-horizon MPC, which is capable of online feedback control but shortsighted for long-term planning, which typically does not allow for separation mode and face-switching cases. Alternatively, a mathematical program with complementarity constraints (MPCC) was proposed in [2], showing faster convergence for planning and recovery from external disturbance, but still susceptible to poor local optima due to the local smoothing and relaxation.

To cope with this issue, sampling-based methods have been used to build a search tree to optimize the symbolic nodes with geometric variables [12], [13], [14], [15]. In the field of non-prehensile manipulation, a search tree with predefined depth was used in [4], corresponding to the number of hybrid switches to optimize over a fixed contact modes sequence. A hybrid Differential Dynamic Programming (DDP) algorithm was simultaneously employed for geometric reasoning. However, the method was validated only with a simplified setup, without considering sliding and separation mode, and without considering optimal waypoints for face switching. Hereby, with additional logic variables, the nodes in the search tree would exponentially grow, yielding exponential computation. Moreover, although expensive sampling computation can be resolved by multiple threads with a powerful machine, there is no guarantee that a feasible solution can be solved in a finite time.

In contrast to such long-term TAMP formulations, humans showcase an impressive agility at exploiting (sub)optimal strategies, especially in terms of discrete variables, by relying on past experience. In this direction, Learning from Demonstration (LfD) has been studied in robotics for a long time and has proven to be an efficient way to solve motion planning challenges [16], [17]. LfD aims to extract motion features from only few demonstrations and then generalize the learned tasks to new situations. Many algorithms have been proposed to encode human demonstrations, such as Dynamic Movement Primitive (DMP) [18], Gaussian Mixture Regression (GMR) [17] and task-parameterized probabilistic models [19]. Among these techniques, K-nearest neighbor (k-NN) [20] is the simplest, yet remaining a powerful tool for discrete data classification [21].

III. DYNAMICAL SYSTEM

In this section, the hybrid dynamics of the pusher-slider system is described. We introduce the kinematics model (Sec. III-A) first, and then extend motion cone (Sec. III-B) and motion equation (Sec. III-C) by including separation mode and face-switching mechanism.

A. Kinematics


Fig. 2 shows the kinematics of the pusher-slider system. The red and blue circles represent the pusher at the initial face and the face after switching. The pose of the slider is defined as $\mathbf{q}_s = [x \ y \ \theta]^\top$ w.r.t. the global frame \mathbb{F}_g , where x and y are the Cartesian coordinates of the center of mass, and

¹<https://sites.google.com/view/dg-oc/>

TABLE I: Constraints of different interaction modes

Sticking	Sliding Up	Sliding Down	Separation
$\{\mathbf{v}_f \in \Omega\} \cap \{\mathbf{q}_{pf} \in \Psi\}$	$\{\Omega < \mathbf{v}_f < \phi\} \cap \{\mathbf{q}_{pf} \in \Psi\}$	$\{\phi < \mathbf{v}_f < \Omega\} \cap \{\mathbf{q}_{pf} \in \Psi\}$	$\{\mathbf{v}_f \notin \phi\} \cup \{\mathbf{q}_{pf} \notin \Psi\}$
$\frac{v_{tf}}{v_{nf}} \leq \gamma_{up}, \quad \frac{v_{tf}}{v_{nf}} \geq \gamma_{dn}, \quad v_{nf} > 0,$ $ p_x = r_s + r_p, \quad p_y \leq r_s.$	$\frac{v_{tf}}{v_{nf}} > \gamma_{up}, \quad v_{nf} > 0,$ $ p_x = r_s + r_p, \quad p_y \leq r_s.$	$\frac{v_{tf}}{v_{nf}} < \gamma_{dn}, \quad v_{nf} > 0,$ $ p_x = r_s + r_p, \quad p_y \leq r_s.$	$\forall v_{nf} < 0,$ $\forall p_x > r_s + r_p,$ $\forall p_y > r_s.$

θ is the rotation angle around the vertical axis. The contact position between the pusher and the slider is described as $\mathbf{q}_{pf} = [p_x \ p_y]^\top$ w.r.t. the current slider frame \mathbb{F}_f after face switching, while $\mathbf{q}_{ps} = \mathbf{R}_{\theta_f} \mathbf{q}_{pf}$ is the expression of \mathbf{q}_{pf} in the initial slider frame \mathbb{F}_s , and θ_f is the rotation angle from \mathbb{F}_s to \mathbb{F}_f . The input of this system is defined in \mathbb{F}_s as the acceleration of the pusher $\mathbf{u} = [v_n \ v_t]^\top$, while $\mathbf{v} = [v_n \ v_t]^\top$ is the pusher velocity. $\mathbf{v}_f = \mathbf{R}_{\theta_f}^\top \mathbf{v} = [v_{nf} \ v_{tf}]^\top$ denotes \mathbf{v} in frame \mathbb{F}_f .


Fig. 2: Kinematics of pusher-slider system allowing face switching.

B. Generalized Motion Cone

Based on the quasi-static approximation and ellipsoidal limit surface assumption [1], [7], [8], [10], a motion cone is introduced to determine the contact mode given by the current pusher velocity and contact position. The two boundaries of the motion cone are given as $\mathbf{v}_{up} = 1 \mathbf{f}_x + \gamma_{up} \mathbf{f}_y$ and $\mathbf{v}_{dn} = 1 \mathbf{f}_x + \gamma_{dn} \mathbf{f}_y$, resolved in the current slider frame \mathbb{F}_f , with

$$\gamma_{up} = \frac{\mu_p c^2 - p_x p_y + \mu_p p_x^2}{c^2 + p_y^2 - \mu_p p_x p_y}, \quad (1)$$

$$\gamma_{dn} = \frac{-\mu_p c^2 - p_x p_y - \mu_p p_x^2}{c^2 + p_y^2 + \mu_p p_x p_y}, \quad (2)$$

where μ_p is the friction coefficient between the pusher and the slider, and c is the parameter connecting applied force and the resulting velocity. More details are introduced in [1].

In the previous work [1], [9], the pusher is assumed to remain on one fixed contact face during the execution. We would like to enable more complex pushing strategies by involving face switching. For that, we introduce another interaction mode, which we call separation mode. Table I lists the relationship between state constraints and interaction modes, where Ω is the set within the boundaries of the motion cone, ϕ is the space where the pusher goes towards the slider, and ψ is the set where the pusher keeps touching with the slider. r_s and r_p are the half length of the slider and the radius of the pusher, respectively.

C. Generalized Motion Equation

We build the generalized motion equation on the basis of [9] and [1], by including separation mode and contact face

switching, namely

$$\dot{\mathbf{x}} = \begin{cases} \mathbf{g}_1(\mathbf{x}, \mathbf{u}), & \text{if Sticking,} \\ \mathbf{g}_2(\mathbf{x}, \mathbf{u}), & \text{if Sliding Up,} \\ \mathbf{g}_3(\mathbf{x}, \mathbf{u}), & \text{if Sliding Down,} \\ \mathbf{g}_4(\mathbf{x}, \mathbf{u}), & \text{if Separation,} \end{cases} \quad (3)$$

where $\mathbf{x} = [\mathbf{q}_s^\top \ \mathbf{q}_p^\top \ \mathbf{v}^\top]^\top$, and

$$\mathbf{g}_j(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \mathbf{R}_{\theta_f} \mathbf{R}_\theta \mathbf{Q} \mathbf{P}_j \\ \mathbf{b}_j \\ \mathbf{R}_{\theta_f} [\mathbf{d}_j \ \mathbf{c}_j]^\top \\ \mathbf{u} \end{bmatrix} \mathbf{R}_{\theta_f}^\top \mathbf{v}, \quad \mathbf{R}_\theta = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix},$$

$$\mathbf{Q} = \frac{1}{c^2 + p_x^2 + p_y^2} \begin{bmatrix} c^2 + p_x^2 & p_x p_y \\ p_x p_y & c^2 + p_y^2 \end{bmatrix},$$

$$\mathbf{b}_1 = \begin{bmatrix} -p_y & p_x \\ c^2 + p_x^2 + p_y^2 & c^2 + p_x^2 + p_y^2 \end{bmatrix}, \mathbf{b}_2 = \begin{bmatrix} -p_y + \gamma_{up} p_x & 0 \\ c^2 + p_x^2 + p_y^2 & c^2 + p_x^2 + p_y^2 \end{bmatrix},$$

$$\mathbf{b}_3 = \begin{bmatrix} -p_y + \gamma_{dn} p_x & 0 \\ c^2 + p_x^2 + p_y^2 & c^2 + p_x^2 + p_y^2 \end{bmatrix}, \quad \mathbf{b}_4 = [0 \ 0],$$

$$\mathbf{c}_1 = [0 \ 0], \quad \mathbf{c}_2 = [-\gamma_{up} \ 1], \mathbf{c}_3 = [-\gamma_{dn} \ 1], \mathbf{c}_4 = [0 \ 1],$$

$$\mathbf{d}_1 = [0 \ 0], \quad \mathbf{d}_2 = [0 \ 0], \quad \mathbf{d}_3 = [0 \ 0], \quad \mathbf{d}_4 = [1 \ 0],$$

$$\mathbf{P}_1 = \mathbf{I}_{2 \times 2}, \quad \mathbf{P}_2 = \begin{bmatrix} 1 & 0 \\ \gamma_{up} & 0 \end{bmatrix}, \mathbf{P}_3 = \begin{bmatrix} 1 & 0 \\ \gamma_{dn} & 0 \end{bmatrix}, \quad \mathbf{P}_4 = \mathbf{0}_{2 \times 2}.$$

where $j = 1, 2, 3, 4$ corresponds to sticking, sliding up, sliding down, and separation mode, respectively.

IV. METHODS

In this section, we first formulate the long-term non-prehensile planar pushing task as an optimal control problem (Sec. IV-A), and then we propose three DDP-related methods by modifying the initialization and cost function, namely Demonstration-started DDP (Sec. IV-B), Demonstration-penalized DDP (Sec. IV-C) and Warm-starting DDP (Sec. IV-D). The combination of Demonstration-penalized DDP and Warm-starting DDP leads to a demonstration-guided hierarchical optimization framework. Finally, we describe the online tracking controller (Sec. IV-E).

A. Problem formulation

An optimal control problem (OCP) can be described as

$$\min_{\mathbf{u}_t} c_T(\mathbf{x}_T) + \sum_{t=0}^{T-1} c_t(\mathbf{x}_t, \mathbf{u}_t), \quad (4)$$

$$\text{s.t. } \mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t, \mathbf{u}_t), \quad (5)$$

where (4) is the cost function and (5) is the dynamic equation. Due to the nonlinearity and high degrees of freedom in robotics, numerical optimization is commonly used to

solve such problems. This paper employs DDP (Differential Dynamic Programming) [22], which has been demonstrated to be effective in previous studies [4], [23]. Furthermore, DDP provides a local feedback mechanism that enhances the robustness of the controller, as demonstrated in [24].

After minimizing the cost-to-go function w.r.t. $\Delta \mathbf{u}_t$, a local stabilizing controller can be obtained as

$$\mathbf{u}_t = \hat{\mathbf{u}}_t + \mathbf{K}_t (\mathbf{x}_t - \hat{\mathbf{x}}_t), \quad (6)$$

where \mathbf{K}_t is the feedback gain, and $\hat{\mathbf{u}}_t$ is the feedforward term.

Given (4) is a non-convex problem, DDP solves it by optimizing around the current solution iteratively. The convergence is very sensitive to the initial guess, which means that the algorithm can converge to poor local optima if the initial guess is far away from the optimal solution.

B. Demonstration-started DDP (DS-DDP)

Demonstrations can be used for OCP initialization to address the poor local optima problem [25], [26]. Here, we also use human demonstration for initialization. The demonstrations are designed as continuous variables, which can implicitly express the discrete variables instead of explicitly specifying the mode sequence as in previous work [1], [2], [11]. In this way, the joint logic and geometric optimization problem can be converted to a continuous optimization problem that can be solved efficiently.

The collected human demonstrations are denoted as $[\tilde{\mathbf{q}}_s, \tilde{\mathbf{q}}_{p_s}, \tilde{\mathbf{v}}, \tilde{\mathbf{u}}]$, with $\tilde{\mathbf{q}}_s = [\tilde{\mathbf{q}}_{s_0}, \tilde{\mathbf{q}}_{s_1}, \dots, \tilde{\mathbf{q}}_{s_T}]$, $\tilde{\mathbf{q}}_{p_s} = [\tilde{\mathbf{q}}_{p_{s_0}}, \tilde{\mathbf{q}}_{p_{s_1}}, \dots, \tilde{\mathbf{q}}_{p_{s_T}}]$, where $\tilde{\mathbf{q}}_{s_t} \in \mathbb{R}^3$ and $\tilde{\mathbf{q}}_{p_{s_t}} \in \mathbb{R}^2$ represent the state of the slider and the pusher at timestep t , respectively. $\tilde{\mathbf{v}} = [\tilde{v}_0, \tilde{v}_1, \dots, \tilde{v}_{T-1}] \in \mathbb{R}^2$ and $\tilde{\mathbf{u}} = [\tilde{u}_0, \tilde{u}_1, \dots, \tilde{u}_{T-1}] \in \mathbb{R}^2$ denote the velocity and acceleration at each timestep.

Given a target \mathbf{q}_s^* , we use k-NN to select the index of j^* with the closest demonstration $\tilde{\mathbf{q}}_{s_T}$ to the target in the task space by evaluating

$$j^* = \arg \min_{j \in \mathbb{S}} \text{dist}(\tilde{\mathbf{q}}_{s_T}^j, \mathbf{q}_s^*), \quad (7)$$

where $\mathbb{S} = \{j : j \in \{0, 1, \dots, n_d - 1\}\}$, n_d is the number of demonstrations, and

$$\text{dist}(\mathbf{x}, \mathbf{y}) = \left(\sum_{r=1}^d |x_r - y_r|^p \right)^{1/p}, \quad (8)$$

where d is the dimension of slider state. p is set as 2.

The cost function of DS-DDP is defined as

$$c_1 = c_{re} + c_{rg} + c_{bd}, \quad (9)$$

with

$$c_{re} = (\boldsymbol{\mu}_T - \mathbf{x}_T)^\top \mathbf{Q}_T (\boldsymbol{\mu}_T - \mathbf{x}_T), \quad c_{rg} = \sum_{t=0}^{T-1} \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t, \\ c_{bd} = \sum_{t=0}^{T-1} \mathbf{f}^{\text{cut}}(\mathbf{u}_t, \mathbf{u}_t)^\top \mathbf{Q}_f \mathbf{f}^{\text{cut}}(\mathbf{u}_t, \mathbf{u}_t),$$

where c_{re} is the reaching cost, and c_{rg} , c_{bd} are the regularizer and boundary penalizer of control commands. \mathbf{u}_t is the

predefined bounding box of \mathbf{u} . \mathbf{f}^{cut} is a soft-thresholding function. The initial guess $\mathbf{u}^0 = \tilde{\mathbf{u}}$ is drawn from human demonstrations directly.

Although this method seems like an effective warm-starting method, it is restricted by the number of acquired demonstrations.

C. Demonstration-penalized DDP (DP-DDP)

To alleviate the problem mentioned above, we propose to use demonstration as soft constraints. It is achieved by designing the cost function as

$$c_2 = c_{re} + c_{rg} + c_{bd} + c_{sw} + c_{ve} + c_{ac}, \quad (10)$$

with

$$c_{sw} = \sum_{n=t_0}^{t_{N-1}} (\boldsymbol{\mu}_n - \mathbf{x}_n)^\top \mathbf{Q}_n (\boldsymbol{\mu}_n - \mathbf{x}_n), \\ c_{ve} = \sum_{t=0}^{T-1} (\tilde{\mathbf{v}}_t - \mathbf{v}_t)^\top \mathbf{R}_{dv} (\tilde{\mathbf{v}}_t - \mathbf{v}_t), \\ c_{ac} = \sum_{t=0}^{T-1} (\tilde{\mathbf{u}}_t - \mathbf{u}_t)^\top \mathbf{R}_{du} (\tilde{\mathbf{u}}_t - \mathbf{u}_t),$$

where c_{sw} , c_{ve} and c_{ac} are designed to follow the demonstrated face switching strategy, pusher velocity and pusher acceleration. $n = [t_0, \dots, t_{N-1}]$ is the timestep when the contact face switches, $\boldsymbol{\mu} = [\tilde{\mathbf{q}}_s^\top, \tilde{\mathbf{q}}_{p_s}^\top, \tilde{\mathbf{v}}^\top]^\top$ is the state of the selected demonstration, and $\tilde{\mathbf{v}}_t$, $\tilde{\mathbf{u}}_t$ are the demonstrated velocity and acceleration at timestep t .

D. Warm-starting DDP (WS-DDP)

Demonstration-penalized DDP shows good performance and can avoid poor local optima, but in order to further improve its convergence properties, we propose a hierarchical optimization framework, where the solution of DP-DDP is used to initialize another DDP problem that we call Warm-starting DDP (WS-DDP).

The cost function of WS-DDP is as same as DS-DDP, allowing it to explore much freely towards the final target. The initial guess $\mathbf{u}^0 = \mathbf{u}_{DP}^*$ is the result of previous DP-DDP.

E. Adaptation to disturbance

Typically, the feedback controller (6) is used to generate control commands at each timestep based on the current state to stabilize the motion along the nominal trajectory. However, the friction and other unmodeled nonlinearities might cause undesired behaviors, especially when the error between the current state and the planned solution $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2$ is too big. To alleviate this issue, inspired by [24], we propose to use an error filtering method based on a trust region defined as a ball of radius r as $\mathcal{B}_t(r) = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \hat{\mathbf{x}}_t\| < r\}$. By denoting the actual robot state as \mathbf{x}_t^0 , we filter the \mathbf{x}_t in (6) as follows: if $\mathbf{x}_t^0 \in \mathcal{B}_t(r)$, then we take $\mathbf{x}_t = \hat{\mathbf{x}}_t$, else we take $\mathbf{x}_t = \mathbf{x}_t^0$. This means that if the error between the actual robot state and the planned state is small, we can use the feedforward terms of the controller directly, and if it is big, then we regenerate the trajectory using the feedback controller gains. This allows us to fully exploit the feedback and feedforward gains computed offline during the online execution of the task avoiding expensive recomputations at each timestep.

V. EXPERIMENTS

We evaluate in this section the proposed offline programming method (Sec. V-A) and the online tracking controller (Sec. V-B).

A. Offline Programming

In this work, the task space is the horizontal plane on the table, and it is limited as $\mathcal{T} = \{[x, y, \theta] : x \in [-25\text{cm}, 25\text{cm}], y \in [-25\text{cm}, 25\text{cm}], \theta \in [-\pi, \pi]\}$. We collected 3 representative demonstrations, which are $[15\text{cm}, -10\text{cm}, -\pi/2], [0, -20\text{cm}, \pi/2], [15\text{cm}, -15\text{cm}, \pi/2]$, corresponding to $N_s = 0$, $N_s = 1$ and $N_s = 2$, respectively, where N_s is the number of face switches during the demonstration. The initial state is defined as $[0, 0, 0, \alpha p_x, 0, 0, 0]^\top$, where $\alpha = 1.3$, corresponding to the separation mode in Sec. III-B, allowing the pusher to select the contact point at the beginning. The cost function gains are set to $\mathbf{Q}_T = 10^6 \times \text{diag}\{1, 1, 1, 10^{6p-5}, 10^{1-6p}, 10^{-3}, 10^{-3}\}$, $\mathbf{Q}_n = 10^6 \times \text{diag}\{10^{-3}, 10^{-3}, 10^{-3}, p, (1-p), 0, 0\}$, where $p = 1$ when $\theta_f = 0$ or $\theta_f = \pi$, otherwise $p = 0$. $\mathbf{R} = \mathbf{K} = \text{diag}\{1, 1\}$, $\mathbf{R}_{du} = \mathbf{R}_{dv} = \text{diag}\{100, 100\}$.

We first randomly selected 100 targets from the task space \mathcal{T} to test our proposed Demonstration-guided Hierarchical Optimization framework, along with two benchmarks: Zero-started DDP (ZS-DDP), namely without using demonstration, and Demonstration-started DDP (DS-DDP). Successful offline programming is defined as achieving $\{x_{\text{err}} < 1\text{cm}, y_{\text{err}} < 1\text{cm}, \theta_{\text{err}} < 5^\circ\}$, where x_{err} , y_{err} , and θ_{err} are the differences between the final and target states. We list the statistical results in Table II. The success rate of ZS-DDP is about 40%, indicating many poor local minima that trap ZS-DDP in bad regions. Initializing with demonstrations increases the success rate by nearly 20%, showing the benefits of using human demonstrations to avoid poor local minima. Our proposed demonstration-guided hierarchical framework, composed of DP-DDP and WS-DDP, achieves a significantly higher success rate (85%) by using demonstrations as soft constraints and warm-starting a new DDP. We consider using demonstration as soft constraints, rather than initialization, as the key feature of our proposed method. This provides the nonlinear optimization process with a few but valuable human guidance, leading to a more efficient exploration process. WS-DDP is then warm-started by DP-DDP, yielding a much smoother and more precise optimal solution. In practice, as long as the three selected demonstrations are informative, the generalization results are not significantly impacted. It would be studied in the future about how to collect demonstrations more efficiently and actively.

To demonstrate the strengths of our method, we moreover compared it with the Mixed-Integer Programming (MIP) formulation, which was commonly used for planar pushing tasks previously [1], [2], [11]. To ensure a fair comparison, we used demonstrations as initialization for MIP as well, which we refer to as DS-MIP. We used CasADi [27] with Bonmin solver [28] and Crocodyl [29] with FDDP solver to solve DS-MIP and DS-DDP, respectively. Convergence and

TABLE II: Performance of ZS-DDP, DS-DDP, DP-DDP, and WS-DDP for offline programming

Method	x_{err} (cm)	y_{err} (cm)	θ_{err} (rad)	succ_rate
ZS-DDP	3.18 ± 8.30	4.78 ± 8.60	0.04 ± 0.39	40%
DS-DDP	1.70 ± 9.07	1.62 ± 9.66	0.21 ± 1.82	58%
DP-DDP	0.14 ± 1.54	0.57 ± 2.77	0.22 ± 0.12	77%
WS-DDP	0.07 ± 1.31	0.19 ± 2.90	0.01 ± 0.08	85%




Fig. 3: Convergence




Fig. 4: Computation time




Fig. 5: Tracking performance under disturbances drawn from uniform distributions, where x_M , y_M and θ_M are the boundaries. The dashed lines present the tolerance.

computation time of 10 randomly selected target configurations are shown in Fig. 3 and Fig. 4. The vertical axis of Fig. 3 represents the reaching error, in meters for x_e and y_e and in radians for θ_e , between the final state \mathbf{x}_T and the target $\boldsymbol{\mu}_T$. DS-MIP achieves almost the same result as DS-DDP, as shown in Fig. 3, but takes ten times longer to solve, as illustrated in Fig. 4. Moreover, our experiments demonstrate that the proposed DP-DDP and WS-DDP algorithms converge to better results, albeit over relatively longer time (still within acceptable limits), in comparison to DS-DDP. DS-DDP was found to converge to poor local optima and increasing the number of iterations did not improve the results. These findings suggest that initializing solely using demonstrations is insufficient, but constraining the search space with demonstrations as soft constraints can lead to better solutions.

B. Online Tracking

For online tracking, we investigated both numerical simulation and real robot experiments. In simulation, we introduced a disturbance on the state as $\mathbf{x} = \bar{\mathbf{x}} + \boldsymbol{\epsilon}$ from the beginning to the end, where $\boldsymbol{\epsilon}_x \sim \mathcal{U}(-x_M, x_M)$, $\boldsymbol{\epsilon}_y \sim \mathcal{U}(-y_M, y_M)$, $\boldsymbol{\epsilon}_\theta \sim \mathcal{U}(-\theta_M, \theta_M)$, are the components of $\boldsymbol{\epsilon}$, drawn from a uniform distributions. Fig. 5 shows the evolution of the errors




Fig. 6: Pushing task with face switching. The manipulator starts from (a) and selects an optimal face (orange) and contact point (b) for pushing, until reaching the planned face-switching point (c). Next, the manipulator changes to another face (d) and touches the object again at the selected point (e), followed by the next phase of pushing (f), until reaching the final target (g). This example is for $N_s = 1$. (d)~(f) should be repeated if $N_s > 1$. The colored lines are used to express the current active faces, and the black arrow in (d) represents the face-switching process.




Fig. 7: Planar pushing with face switching. Both simulation and experimental results reach final targets within tolerance.

on x , y , and θ , computed as the difference between the final point and the target, for increasing x_M , y_M and θ_M . The tolerance for online tracking is set as $\{x_{\text{err}} < 3\text{cm}, y_{\text{err}} < 3\text{cm}, \theta_{\text{err}} < 5^\circ (0.087\text{rad})\}$. The controller can successfully resist 4cm perturbation for x and y , and 0.117rad for θ .

Then, we tested the proposed method on the real robot setup (Fig. 1), using a 7-axis Franka Emika robot and a RealSense D435 camera. The prismatic slider, a $12 \times 12\text{cm}$ block made of 3D-printed PLA and weighting 110g, was positioned on a flat plywood surface, with an Aruco Marker attached to its top face. The friction coefficients between the pusher and the slider, and between the plywood surface and the slider, are 0.3 and 0.35, respectively. The robot was equipped with a wooden pusher that had a 0.5cm radius and was used to move the slider. The camera captured the motion of the slider at 30Hz, while the feedback controller ran at 100Hz. The low-level Cartesian impedance controller, responsible for actuating the robot, ran at 1000Hz.

As reported in [4], most targets can be achieved with only one to two face switches. We therefore tested our approach with one externally disturbed line tracking task and two face-switching tasks (see accompanying video). The trajectories in Fig. 7 correspond to the task requiring one face switching to push the object from $[0, 0, 0]$ to $[20\text{cm}, -20\text{cm}, \pi/2]$. The control strategy is intuitive: pushing it from the left face to slightly adjust the pose and then changing to the top face for the next phase of pushing. Fig. 7-(a) is the simulation trajectory, which is generated by using PyBullet [30], while Fig. 7-(b) shows the real robot trajectory. It is observed that these two trajectories are significantly different, because of

the different friction parameters in the two different worlds. Still, both can overcome the uncertainty to reach the final target. Despite the existence of unstructured elements such as differences in the visual system and the robot controller, several assumptions of the dynamics model, as well as immeasurable friction, the feedback controller is able to cope with these different mismatches and track the reference trajectory successfully. Fig. 6 shows the keyframes of the pusher pushing the slider toward the target, indicating that 7 steps are needed with the face-switching strategy. Another final target configuration, $[5\text{cm}, -18\text{cm}, \pi/5]$, which requires two face switches, is additionally shown in the video.

VI. DISCUSSION, LIMITATIONS AND CONCLUSION

In this paper, we proposed to add separation mode and face-switching mechanism to the problem of pushing objects on a planar surface. We showed that by introducing only few human demonstrations, the typical TAMP problem can be expressed as a classical continuous optimization problem, which is much more efficient to be solved. With the proposed demonstration-guided hierarchical optimization framework, we demonstrated significantly better results in terms of generalization and precision compared to the state-of-the-art methods. Additionally, we developed a feedback controller based on DDP feedback gains to regenerate the trajectory for online tracking. We tested the combination of these approaches in both PyBullet and real robot applications, showing good performance to resist contact uncertainty.

Currently, we are using a feedback controller to track the offline trajectory. If the system is subject to large perturbations such as rotating 180° , an online Model Predictive Control (MPC) may still be required. Nevertheless, our demonstration-guided method is also promising as an optimizer within MPC to avoid poor local optima.

As future work, we aim to apply our demonstration-guided approach to a broader range of manipulation tasks (namely, beyond pushing problems), which requires further study on how to extract and represent constraints from demonstrations. It would also be relevant to explore extensions of TAMP problems by introducing similar continuous human demonstrations. Further work could also exploit human demonstrations in model-based learning strategies to let the robot automatically refine the pushing model and its motion.

REFERENCES

- [1] F. R. Hogan and A. Rodriguez, "Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics," in *Algorithmic Foundations of Robotics XII: Proceedings of the Twelfth Workshop on the Algorithmic Foundations of Robotics*. Springer, 2020, pp. 800–815.
- [2] J. P. De Moura, T. Stouraitis, and S. Vijayakumar, "Non-prehensile planar manipulation via trajectory optimization with complementarity constraints," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2022.
- [3] T. Pang, H. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *arXiv preprint arXiv:2206.10787*, 2022.
- [4] N. Doshi, F. R. Hogan, and A. Rodriguez, "Hybrid differential dynamic programming for planar manipulation primitives," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2020, pp. 6759–6765.
- [5] B. Siciliano, O. Khatib, and T. Kröger, *Springer handbook of robotics*. Springer, 2008, vol. 200.
- [6] M. T. Mason, "Progress in nonprehensile manipulation," *Intl Journal of Robotics Research*, vol. 18, no. 11, pp. 1129–1141, 1999.
- [7] S. Goyal, A. Ruina, and J. Papadopoulos, "Limit surface and moment function descriptions of planar sliding," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 1989, pp. 794–795.
- [8] S. H. Lee and M. Cutkosky, "Fixture planning with friction," *Journal of Manufacturing Science and Engineering, Transactions of the ASME*, vol. 113, no. 3, pp. 320–327, 1991.
- [9] K. M. Lynch, H. Maekawa, and K. Tanie, "Manipulation and active sensing by pushing using tactile feedback," in *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, vol. 1, 1992, pp. 416–421.
- [10] M. T. Mason, "Mechanics and planning of manipulator pushing operations," *Intl Journal of Robotics Research*, vol. 5, no. 3, pp. 53–71, 1986.
- [11] F. R. Hogan, E. R. Grau, and A. Rodriguez, "Reactive planar manipulation with convex hybrid mpc," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2018, pp. 247–253.
- [12] A. Simeonov, Y. Du, B. Kim, F. Hogan, J. Tenenbaum, P. Agrawal, and A. Rodriguez, "A long horizon planning framework for manipulating rigid pointcloud objects," in *Conference on Robot Learning*. PMLR, 2021, pp. 1582–1601.
- [13] M. Toussaint, "Logic-geometric programming: An optimization-based approach to combined task and motion planning," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [14] T. Migimatsu and J. Bohg, "Object-centric task and motion planning in dynamic environments," *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 844–851, 2020.
- [15] R. E. Fikes and N. J. Nilsson, "Strips: A new approach to the application of theorem proving to problem solving," *Artificial intelligence*, vol. 2, no. 3-4, pp. 189–208, 1971.
- [16] A. G. Billard, S. Calinon, and R. Dillmann, "Learning from humans," in *Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Secaucus, NJ, USA: Springer, 2016, ch. 74, pp. 1995–2014, 2nd Edition.
- [17] S. Calinon and D. Lee, "Learning control," in *Humanoid Robotics: a Reference*, P. Vadakkepat and A. Goswami, Eds. Springer, 2019, pp. 1261–1312.
- [18] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: learning attractor models for motor behaviors," *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.
- [19] S. Calinon, "Robot learning with task-parameterized generative models," in *Robotics Research*. Springer, 2018, pp. 111–126.
- [20] L. E. Peterson, "K-nearest neighbor," *Scholarpedia*, vol. 4, no. 2, p. 1883, 2009.
- [21] P. Cunningham and S. J. Delany, "K-nearest neighbour classifiers-a tutorial," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–25, 2021.
- [22] D. Mayne, "A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems," *International Journal of Control*, vol. 3, no. 1, pp. 85–95, 1966.
- [23] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2014, pp. 1168–1175.
- [24] H. Girgin, J. Jankowski, and S. Calinon, "Reactive anticipatory robot skills with memory," in *International Symposium on Robotics Research (ISRR)*, 2022.
- [25] T. S. Lembono, A. Paolillo, E. Pignat, and S. Calinon, "Memory of motion for warm-starting trajectory optimization," *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 2594–2601, 2020.
- [26] N. Mansard, A. DelPrete, M. Geisert, S. Tonneau, and O. Stasse, "Using a memory of motion to efficiently warm-start a nonlinear predictive controller," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2986–2993.
- [27] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [28] P. Bonami, M. Kiliç, and J. Linderoth, "Algorithms and software for convex mixed integer nonlinear programs," in *Mixed integer nonlinear programming*. Springer, 2012, pp. 1–39.
- [29] 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 *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, 2020, pp. 2536–2542.
- [30] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning.(2016)," URL <http://pybullet.org>, 2016.