arXiv:2212.12814v1 [cs.RO] 24 Dec 2022

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 contact points and directions, which involves joint logic and geometrical reasoning in the modes of the dynamics model. 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 cases, 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 real Franka Emika robot in the presence of external disturbances.**

I. Introduction

With the trend of labor shortage and aging population, robots are required for increasingly more 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 rotation, 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 as follows:
1. **Hybrid Dynamics**. The dynamics of a pusher-slider system depends on the current interaction mode and contact face between the pusher and the slider. In this paper, we consider not only the motion involving various interaction modes, e.g., separation, sticking, sliding up, and sliding down, but also the switching between the contact faces, i.e., left, bottom, right, and up (unlike previous works [1], [2] that only work with a single face). The transitions between different faces and the separation to the other three contacting modes pose important difficulty for gradient-based optimization methods.
2. **Underactuation**. The contact force between the pusher and the slider is constrained within a motion cone, which makes 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 commands.
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.

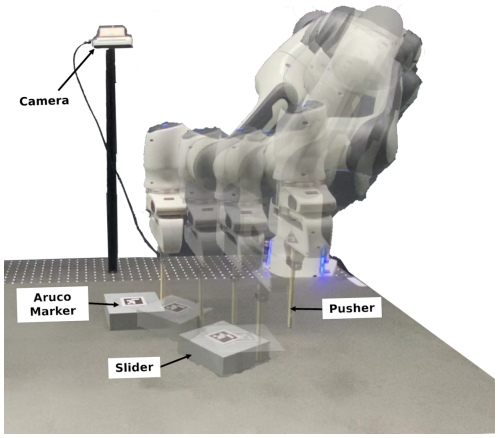

Fig. 1: System setup where the robot pushes an object with contact switching.

To address these challenges, we propose a hybrid framework based on optimization and learning from human demonstrations. An interface is first built to collect human demonstrations of the pushing task. Given a specific target configuration, we use k-nearest neighbor (k-NN) algorithm to retrieve a demonstration trajectory closest to the target configuration, and we use it to formulate the soft constraints of a hierarchical optimal control problem. The optimal control problem is solved offline to produce the optimal trajectory and the optimal control commands. A feedback controller is then used to adapt to the contact uncertainty online by following the offline trajectory.

The main contribution of our paper is an approach to solve the complete long-term non-prehensile manipulation task by covering all of the interaction modes and the face

[1]Idiap Research Institute, Martigny, Switzerland (e-mail: firstname.lastname@idiap.ch)
[2]École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

switching cases, with the ability of offline planning and online tracking. This is achieved by introducing human demonstration into the optimal control problem. By using demonstrated continuous variables to express the discrete variables implicitly, we can successfully project traditional TAMP formulation into geometric field. A demonstration-guided hierarchical optimization framework is then proposed, allowing the robot to obtain (sub)optimal solutions very quickly. A real-time feedback controller is finally proposed to replan the trajectory, by compensating for model mismatch and contact uncertainty when interacting with the real physical world.

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 [3], with the pusher-slider system as one of the most prominent example. Under the assumption of quasi-static interaction [4] and ellipsoidal limit surface [5], Lynch *et al.* put forward a set of analytical methods to compute velocities based on the motion cone [6], which was previously studied in [7] 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 [8], 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 pushing face is fixed, which is too much 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 mode schedule learning was later proposed in [8], using a 3-layer mode sequence classifier. To train the classifier, a large dataset was built by a random state generator. A drawback of this approach is that a higher number of discrete modes and contact faces would require the dataset volume 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], which was introduced into a trajectory optimization framework to replace integer variables, showing faster convergence for planning and recovery from external disturbance. However, the formulation of MPCC is based on the assumption of continuous mode transitions, which is not suitable for the separation cases appearing in long-term manipulation, because the instantaneous conversion from separation to contact is physically infeasible. Thus, a framework focusing on non-prehensile manipulation with long-term TAMP fashion is necessary.

To cope with this issue, sampling-based methods have been used to build a search tree to optimize the symbolic nodes with geometric variables [9], [10], [11], [12]. In the field of non-prehensile manipulation, a tree with predefined depth was used in [13], corresponding to the number of hybrid switches to optimize over a fixed contact modes sequence. A hybrid Differential Dynamic Programming (DDP) algorithm with a linear form was simultaneously used for geometric reasoning. However, the method was validated only with a simplified setup, without considering sliding and separation mode, and without considering optimal viapoints 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 [14], [15]. 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) [16], Gaussian Mixture Regression (GMR) [17] and task-parameterized probabilistic models [18]. Among these techniques, K-nearest neighbor (k-NN) [19] is the simplest, yet remaining a powerful tool for discrete data classification [20].

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 contact switching. The pose of the slider is defined as $\boldsymbol{q}_s = [x\ y\ \theta]^\top$ w.r.t. the global frame $\mathbb{F}_g$,

TABLE I: Constraints of different interaction modes

Sticking	**Sliding Up**	**Sliding Down**	**Separation**
$\{\boldsymbol{v}_f \in \Omega\} \cap \{\boldsymbol{q}_p \in \psi\}$	$\{\Omega < \boldsymbol{v}_f < \phi\} \cap \{\boldsymbol{q}_p \in \psi\}$	$\{\phi < \boldsymbol{v}_f < \Omega\} \cap \{\boldsymbol{q}_p \in \psi\}$	$\{\boldsymbol{v}_f \notin \phi\} \cup \{\boldsymbol{q}_p \notin \psi\}$
$\frac{v_{tf}}{v_{nf}} \leq \gamma_{up}, \quad \frac{v_{tf}}{v_{nf}} \geq \gamma_{dn}, \quad v_{nf} > 0,$ $\lvert p_x \rvert = r_s + r_p, \quad \lvert p_y \rvert \leq r_s.$	$\frac{v_{tf}}{v_{nf}} > \gamma_{up}, \quad v_{nf} > 0,$ $\lvert p_x \rvert = r_s + r_p, \quad \lvert p_y \rvert \leq r_s.$	$\frac{v_{tf}}{v_{nf}} < \gamma_{dn}, \quad v_{nf} > 0,$ $\lvert p_x \rvert = r_s + r_p, \quad \lvert p_y \rvert \leq r_s.$	$\vee\, v_{nf} < 0,$ $\vee\, \lvert p_x \rvert > r_s + r_p,$ $\vee\, \lvert p_y \rvert > r_s.$

where x and y are the Cartesian coordinates of the center of mass, and θ is the rotation angle around the vertical axis. The contact position between the pusher and the slider is described as $\boldsymbol{q}_p = [p_x \ p_y]^\top$ w.r.t. the current slider frame $\mathbb{F}_f$ after face switching, while $\boldsymbol{q}_f = \boldsymbol{R}_{\theta_f} \boldsymbol{q}_p$ is the expression of $\boldsymbol{q}_p$ 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 the acceleration of the pusher $\boldsymbol{u} = [\dot{v}_n \ \dot{v}_t]^\top$, which is resolved in $\mathbb{F}_s$, while $\boldsymbol{v} = [v_n \ v_t]^\top$ is the pusher velocity defined in frame $\mathbb{F}_s$, and $\boldsymbol{v}_f = \boldsymbol{R}_{\theta_f}^\top \boldsymbol{v} = [v_{nf} \ v_{tf}]^\top$ is the expression of $\boldsymbol{v}$ in $\mathbb{F}_f$.

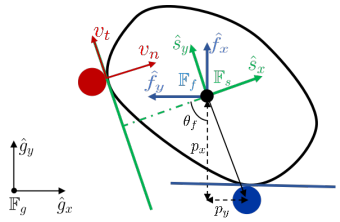

Fig. 2: Kinematics of pusher-slider system allowing face switching.

B. Generalized Motion Cone

Based on the quasi-static approximation and ellipsoidal limit surface assumption [1], [4], [5], [7], 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 $\boldsymbol{v}_{up} = 1\boldsymbol{f_x} + \gamma_{up}\boldsymbol{f_y}$ and $\boldsymbol{v}_{dn} = 1\boldsymbol{f_x} + \gamma_{dn}\boldsymbol{f_y}$, resolved in the current slider frame $\mathbb{F}_f$, with

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

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

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], [6], the pusher is assumed to remain on one face selected before and not change during the execution. We would like to enable more complex pushing that involves 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 pusher, respectively.

C. Generalized Motion Equation

We build the generalized motion equation on the basis of [6] and [1], by including separation mode and contact face switching, namely

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

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

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

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

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

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

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

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

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

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-constrained DDP (Sec. IV-C) and Warm-starting DDP (Sec. IV-D). The combination of Demonstration-constrained 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

A optimal control problem (OCP) can be described as

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

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

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

mostly used to solve this kind of problem. We use DDP [21] in this paper, which has been shown effective for this problem [13], [22]. It can also provide local feedback mechanism that was used to keep the robustness of controller in [23].

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

$$\boldsymbol{u}_t = \hat{\boldsymbol{u}}_t + \boldsymbol{K}_t (\boldsymbol{x}_t - \hat{\boldsymbol{x}}_t) + \boldsymbol{k}_t, \tag{6}$$

where $\boldsymbol{K}_t$ is a feedback gain, and $\boldsymbol{k}_t$ is a 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 it is easy to be stuck at poor local optima if the initial guess is far away from the optimal solution.

B. Demonstration-started DDP (DS-DDP)

To solve the problem of getting stuck at poor local optima, we introduce human demonstrations into the basic DDP as initialization, which is a trendy way for warm-starting OCP [24]. 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], [8]. In this way, we tactfully convert the joint logic and geometric optimization problem to a typical geometric optimization problem, which can be solved much efficiently.

The collected human demonstrations are denoted as $[\widetilde{\boldsymbol{q}}_s, \widetilde{\boldsymbol{q}}_f, \widetilde{\boldsymbol{v}}, \widetilde{\boldsymbol{u}}]$, with $\widetilde{\boldsymbol{q}}_s = [\widetilde{\boldsymbol{q}}_{s_0}, \widetilde{\boldsymbol{q}}_{s_1}, \cdots, \widetilde{\boldsymbol{q}}_{s_T}]$, $\widetilde{\boldsymbol{q}}_f = [\widetilde{\boldsymbol{q}}_{f_0}, \widetilde{\boldsymbol{q}}_{f_1}, \cdots, \widetilde{\boldsymbol{q}}_{f_T}]$, where $\widetilde{\boldsymbol{q}}_{s_t} \in \mathbb{R}^3$ and $\widetilde{\boldsymbol{q}}_{f_t} \in \mathbb{R}^2$ represent the state of the slider and the pusher at timestep t, respectively. $\widetilde{\boldsymbol{v}} = [\widetilde{\boldsymbol{v}}_0, \widetilde{\boldsymbol{v}}_1, \cdots, \widetilde{\boldsymbol{v}}_{T-1}] \in \mathbb{R}^2$ and $\widetilde{\boldsymbol{u}} = [\widetilde{\boldsymbol{u}}_0, \widetilde{\boldsymbol{u}}_1, \cdots, \widetilde{\boldsymbol{u}}_{T-1}] \in \mathbb{R}^2$ denote the velocity and acceleration at each timestep.

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

$$j^* = \min_{j \in \mathbb{S}} \quad \text{dist}(\widetilde{\boldsymbol{q}}_{s_T}^j, \boldsymbol{q}_s^*), \tag{7}$$

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

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

where d is the dimension of slider state, and $p = 2$.

The cost function of DS-DDP is defined as:

$$c_1 = c_{re} + c_{rg} + c_{bd}, \tag{9}$$

with

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

where c_{re} is the reaching cost, and c_{rg}, c_{bd} are the regularizer and boundary penalizer of control commands. $\boldsymbol{u}_l$ is the predefined bounding box of $\boldsymbol{u}$. $\boldsymbol{f}^{\text{cut}}$ is a soft-thresholding function. The initial guess $\boldsymbol{u}^0 = \widetilde{\boldsymbol{u}}$ 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-constrained DDP (DC-DDP)

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

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

with

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

where c_{sw}, c_{ve} and c_{ac} are desinged to follow the demonstrated face switching strategy, pusher velocity and pusher acceleration. $n = [t_0, \cdots, t_N]$ is the timestep when the contact face switches, $\boldsymbol{\mu} = [\widetilde{\boldsymbol{q}}_s^\top \ \widetilde{\boldsymbol{q}}_f^\top \ \widetilde{\boldsymbol{v}}^\top]^\top$ is the state of the selected demonstration, and $\widetilde{\boldsymbol{v}}_t$, $\widetilde{\boldsymbol{u}}_t$ are the demonstrated velocity and acceleration at timestep t.

D. Warm-starting DDP (WS-DDP)

Demonstration-constrained 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 DC-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 $\boldsymbol{u}^0 = \boldsymbol{u}_{DC}^*$ is the result of previous DC-DDP.

E. Adaptation to disturbance

Similarly to basic DDP, we can see that for a tracking problem, the resulting optimal control policy takes the same form as (6), characterized by a feedback gain and a feedforward term. Typically, the optimal control policy is used to generate control commands at each timestep based on the current state to stabilize the motion along the nominal trajectory. However, the frictions and other unmodeled nonlinearities might cause undesired behaviors, especially when the error between the current state and the planned solution $\|\boldsymbol{x}_t - \hat{\boldsymbol{x}}_t\|^2$ is too big. To alleviate this issue, inspired by [23], we propose to use an error filtering method based on a trust region defined as a ball of radius r as $\mathscr{B}_t(r) = \{\boldsymbol{x} \in \mathbb{R}^n | \|\boldsymbol{x} - \hat{\boldsymbol{x}}_t\| < r\}$. By denoting the actual robot state as $\boldsymbol{x}_t^0$, we filter the $\boldsymbol{x}_t$ in (6) as follows: if $\boldsymbol{x}_t^0 \in \mathscr{B}_t(r)$, then we take $\boldsymbol{x}_t = \hat{\boldsymbol{x}}_t$, else we take $\boldsymbol{x}_t = \boldsymbol{x}_t^0$. 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

replan 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 proposed 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 $\mathscr{T} = \{[x, y, \theta] : x \in [-25\text{cm}, 25\text{cm}], y \in [-25\text{cm}, 25\text{cm}], \theta \in [-\pi, \pi]\}$. 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 $\boldsymbol{Q}_T = 10^6 \times \text{diag}\{1, 1, 1, 10^{6p-5}, 10^{1-6p}, 10^{-3}, 10^{-3}\}$, $\boldsymbol{Q}_n = 10^6 \times \text{diag}\{10^{-3}, 10^{-3}, 10^{-3}, p, (1-p), 0, 0\}$, where $p = 1$ when $\theta_f = 0$ or $\theta_f = \pi$, otherwise $p = 0$. $\boldsymbol{R} = \boldsymbol{K} = \text{diag}\{1, 1\}$, $\boldsymbol{R}_{du} = \boldsymbol{R}_{dv} = \text{diag}\{100, 100\}$.

In other works, planar pushing tasks are often formulated as Mixed-Integer Programming (MIP) problems [1], [8], [2]. To compare with our method, we use demonstrations as the initialization for MIP as well, which is called as DS-MIP. CasADi [25] with the Bonmin solver [26] and Crocoddyl [27] with the FDDP solver are used separately to solve DS-MIP and DS-DDP. Fig. 3 and Fig. 4 show the convergence curve and the solving time of 10 randomly selected target configurations. The vertical axis of Fig. 3 is the norm of the reaching error at each iteration with respect to the initial error. We can see in these figures that DS-MIP can almost get the same result as DS-DDP but needs 10 times more time. This is because of the nonconvex nature of integer variables. By implicitly expressing integer variables as demonstrated continuous control commands, TAMP problems can be solved much more efficiently. Moreover, we also find that the proposed DC-DDP and WS-DDP can achieve better results in relatively longer time (but still acceptable) compared to DS-DDP. This shows that simply initializing using the demonstration as done in DS-DDP is not enough, and constraining the search space using demonstrations as the soft constraints during early iterations helps to reach a better solution.

Additionally, we tested the generalization ability of the proposed demonstration-guided offline programming method. A successful offline programming is defined as: $\{x_{\text{err}} < 1\text{cm},\ y_{\text{err}} < 1\text{cm},\ \theta_{\text{err}} < 5^\circ\}$, where x_{err}, y_{err} and θ_{err} are the difference between final pose and target pose. 100 targets are randomly selected in the task space $\mathscr{T}$ to test the generalization capability. The statistical results are listed in Table II. Clearly, with demonstrations, the success

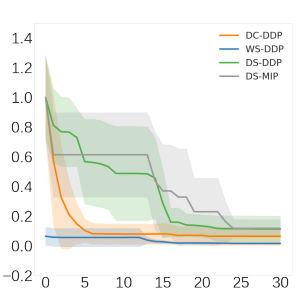

Fig. 3: Convergence curve

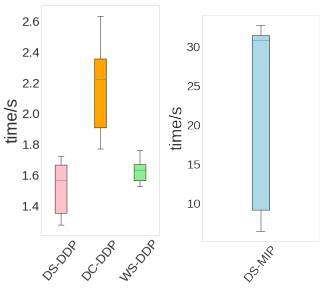

Fig. 4: Solving time

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

Method	x_{err}/cm	y_{err}/cm	θ_{err}/rad	succ_rate
DS-DDP	0.24 ± 2.22	0.96 ± 4.03	0.10 ± 0.34	74%
DC-DDP	0.04 ± 1.33	0.85 ± 1.91	0.02 ± 0.08	75%
WS-DDP	**0.11 ± 1.01**	**0.56 ± 1.63**	**0.01 ± 0.07**	**84%**

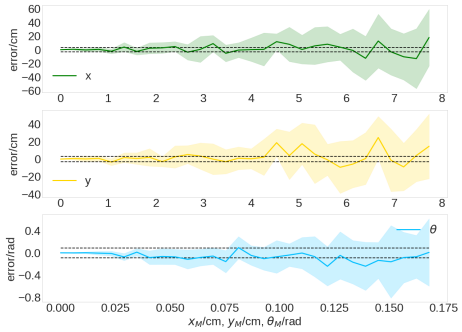

Fig. 5: Tracking performance under disturbance. The dashed lines present the tolerance.

rate significantly increases for random-selected targets. By using only 3 demonstrations, DC-DDP can accomplish 75 out of 100 random targets in task space, also with low mean and standard deviations for the errors. WS-DDP achieves a slightly higher success rate based on the DC-DDP solution, showing that this demonstration-guided hierarchical optimization framework can generalize to unknown targets very well. Practically, the generalization result doesn't change a lot as long as the selected 3 demonstrations are informative. It would be studied in the future about how to collect demonstrations more efficiently and actively.

B. Online Tracking

For online tracking, we investigated both numerical simulation and real robot experiments. In simulation, we introduce a disturbance on the state as $\boldsymbol{x} = \overline{\boldsymbol{x}} + \boldsymbol{\epsilon}$ from the beginning to the end, where $\varepsilon_x \sim \mathscr{U}(-x_M, x_M)$, $\varepsilon_y \sim \mathscr{U}(-y_M, y_M)$, $\varepsilon_\theta \sim \mathscr{U}(-\theta_M, \theta_M)$, are the components of $\boldsymbol{\epsilon}$, drawn from a uniform distributions. Fig. 5 shows the evolution of the errors 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 $\{x_{\text{err}} < 3\text{cm},\ y_{\text{err}} < 3\text{cm},\ \theta_{\text{err}} < 5^\circ/0.087\text{rad}\}$. We can find that the controller can successfully resist 4cm perturbation for x and y, and 0.117rad for θ.

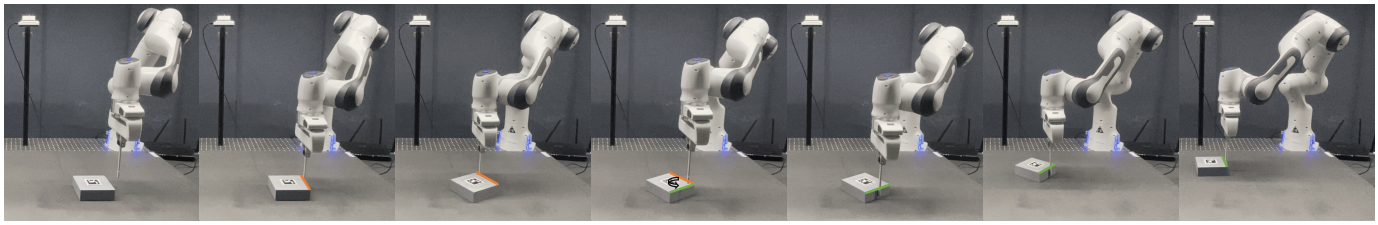

(a) Initialization **(b)** Contact **(c)** Pushing **(d)** Face switching **(e)** Contact **(f)** Pushing **(g)** Reaching

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 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 line is used to express the current active face, and the black arrow in (d) represents the face switching process.

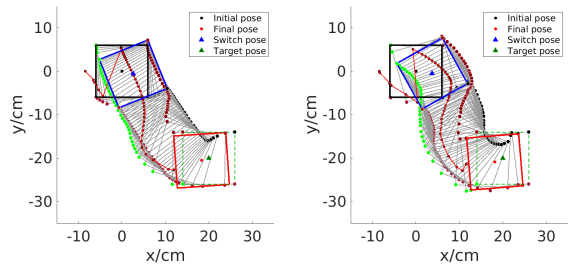

a) Simulation results b) Experimental results

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

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 slider ($r_s = 6$cm) is a 3D-printed prismatic object with PLA, lying on a flat plywood surface, with an Aruco Marker on the top face. A wooden pusher ($r_p = 0.5$cm) is attached to the robot to move the object. The motion of the object is tracked by the camera at 30 HZ, and the feedback controller runs at 100 HZ, with a low-level Cartesian impedance controller (1000 HZ) actuating the robot.

In this experiment, we tried one line tracking task with disturbance and two face switching tasks. Both simulated and experimental results achieve the targets within specified tolerance (see accompanying video). The trajectories in Fig. 7 correspond to one of 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 [28], 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 propose to add separation modes and face switching mechanisms to the problem of pushing objects on a planar surface. We showed that by introducing human demonstrations, the typical TAMP problem can be expressed as a classical geometric 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. With more demonstrations, the proposed approach has the potential to precisely reach almost any point in the task space. Additionally, we developed a feedback controller based on DDP feedback gains to replan the trajectory for online tracking. We tested the combination of these approaches in both PyBullet simulation and in real robot application, 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 perturbation 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 constraints from demonstration, and how to formulate an optimal control problem based on the extracted constraints. It would also be relevant to explore extensions of TAMP problems by introducing similar continuous human demonstration. Another future work consists in exploiting 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," *arXiv preprint arXiv:1611.08268*, 2016.

[2] J. P. De Moura, T. Stouraitis, and S. Vijayakumar, "Non-prehensile planar manipulation via trajectory optimization with complementarity constraints," in *2022 IEEE International Conference on Robotics and Automation*, 2022.

[3] M. T. Mason, "Progress in nonprehensile manipulation," *The International Journal of Robotics Research*, vol. 18, no. 11, pp. 1129–1141, 1999.

[4] S. Goyal, A. Ruina, and J. Papadopoulos, "Limit surface and moment function descriptions of planar sliding," in *1989 IEEE International Conference on Robotics and Automation*. Citeseer, 1989, pp. 794–795.

[5] 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.

[6] K. M. Lynch, H. Maekawa, and K. Tanie, "Manipulation and active sensing by pushing using tactile feedback." in *IROS*, vol. 1, 1992, pp. 416–421.

[7] M. T. Mason, "Mechanics and planning of manipulator pushing operations," *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 53–71, 1986.

[8] F. R. Hogan, E. R. Grau, and A. Rodriguez, "Reactive planar manipulation with convex hybrid mpc," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 247–253.

[9] 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.

[10] 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.

[11] T. Migimatsu and J. Bohg, "Object-centric task and motion planning in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 844–851, 2020.

[12] 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.

[13] N. Doshi, F. R. Hogan, and A. Rodriguez, "Hybrid differential dynamic programming for planar manipulation primitives," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 6759–6765.

[14] 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.

[15] S. Calinon and D. Lee, "Learning control," in *Humanoid Robotics: a Reference*, P. Vadakkepat and A. Goswami, Eds. Springer, 2019, pp. 1261–1312.

[16] 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.

[17] S. Calinon and D. Lee, "Learning control," *Humanoid Robotics: a Reference*, 2017.

[18] S. Calinon, "Robot learning with task-parameterized generative models," in *Robotics Research*. Springer, 2018, pp. 111–126.

[19] L. E. Peterson, "K-nearest neighbor," *Scholarpedia*, vol. 4, no. 2, p. 1883, 2009.

[20] P. Cunningham and S. J. Delany, "K-nearest neighbour classifiers-a tutorial," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–25, 2021.

[21] 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.

[22] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1168–1175.

[23] H. Girgin, J. Jankowski, and S. Calinon, "Reactive anticipatory robot skills with memory," in *International Symposium on Robotics Research (ISRR)*, 2022.

[24] T. S. Lembono, A. Paolillo, E. Pignat, and S. Calinon, "Memory of motion for warm-starting trajectory optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2594–2601, 2020.

[25] 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.

[26] P. Bonami, M. Kilinç, and J. Linderoth, "Algorithms and software for convex mixed integer nonlinear programs," in *Mixed integer nonlinear programming*. Springer, 2012, pp. 1–39.

[27] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocoddyl: An efficient and versatile framework for multi-contact optimal control," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.

[28] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games," *Robotics and machine learning.[Google Scholar]*, 2016.