

Dynamic Manipulation of Deformable Objects using Imitation Learning with Adaptation to Hardware Constraints

Eric Hannus¹, Tran Nguyen Le¹, David Blanco-Mulero^{1,2}, Ville Kyrki¹

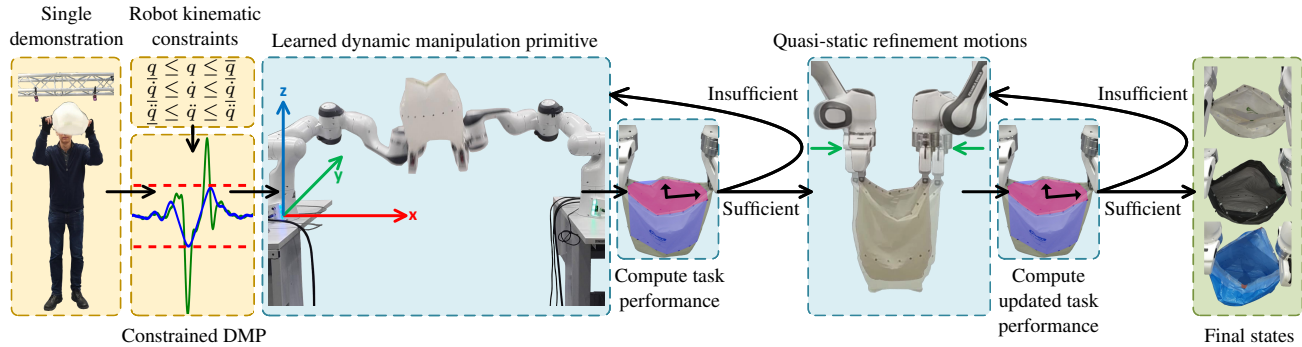


Fig. 1: Our framework combines a dynamic motion primitive that adheres to the robots’ constraints, for quick task progression, with quasi-static motions for final adjustments. We evaluate the framework by implementing *BILBO: Bimanual dynamic manipulation using Imitation Learning for Bag Opening*, a system for a practical bag-opening task.

Abstract—Imitation Learning (IL) is a promising paradigm for learning dynamic manipulation of deformable objects since it does not depend on difficult-to-create accurate simulations of such objects. However, the translation of motions demonstrated by a human to a robot is a challenge for IL, due to differences in the embodiments and the robot’s physical limits. These limits are especially relevant in dynamic manipulation where high velocities and accelerations are typical. To address this problem, we propose a framework that first maps a dynamic demonstration into a motion that respects the robot’s constraints using a constrained Dynamic Movement Primitive. Second, the resulting object state is further optimized by quasi-static refinement motions to optimize task performance metrics. This allows both efficiently altering the object state by dynamic motions and stable small-scale refinements. We evaluate the framework in the challenging task of bag opening, designing the system *BILBO: Bimanual dynamic manipulation using Imitation Learning for Bag Opening*. Our results show that *BILBO* can successfully open a wide range of crumpled bags, using a demonstration with a single bag. See supplementary material at <https://sites.google.com/view/bilbo-bag>.

I. INTRODUCTION

Dynamic manipulation has shown great potential, particularly in the manipulation of deformable objects [1]–[3]. Unlike quasi-static manipulation like pick-and-place, dynamic manipulation utilizes forces of acceleration for the success of the task [4]. Therefore dynamic actions can achieve object

configurations that are not reachable through quasi-static manipulation [2], [5]. For instance, to dynamically manipulate objects such as cloths, a typical approach is to learn the skill in simulation and then transfer it to the real world [1], [5]. However, simulation approaches can be unsuitable for more complex deformable objects such as bags, due to their 3D structure and the complex forces, for example, those related to aerodynamics [2]. One solution to bypass this challenge is to learn the manipulation skills directly from human demonstrations using Imitation Learning (IL).

IL is a learning paradigm for teaching skills to robots by providing demonstrations [6]. The demanding dynamics of dynamic manipulation including high accelerations and velocities highlight the IL challenge known as the correspondence problem [7]. Specifically, intuitive dynamic motions exhibited by human demonstrators may be impossible to transfer directly to robots, due to differences between the bodies and constraints of the robot actuation.

To overcome this problem, we propose a framework, shown in Fig. 1, that first transforms a dynamic human demonstration into a motion that adheres to the robot constraints by utilizing constrained Dynamic Movement Primitives (DMPs) [8], [9]. In the second stage of the framework, quasi-static motions are used for stable refinement of the deformable object state. Our framework decides whether to repeat an action or proceed to the next stage using the task performance metrics. These task specific metrics can be computed utilizing the same system used for recording the demonstration, e.g. a motion capture system.

We demonstrate the benefits of the proposed framework in the task of bimanual bag opening with the novel system *BILBO: Bimanual dynamic manipulation using Imitation*

This work was financially supported by Business Finland (decision 9249/31/2021) and European Union’s Horizon Europe programme project SoftEnable (grant agreement No. 101070600).

¹ Intelligent Robotics Group at the Department of Electrical Engineering and Automation, School of Electrical Engineering, Aalto University, Espoo, Finland. {firstname.lastname}@aalto.fi

² Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain. david.blanco.mulero@upc.edu

Learning for Bag Opening. BILBO employs a dynamic motion to optimize the volume and area of the bags, followed by a linear refinement motion to enhance the opening roundness (see Fig. 1). Our experiments show that a single human demonstration is sufficient for BILBO to successfully open bags of a wide range of sizes and material properties, even when they are initialized in a highly crumpled state. The experimental results additionally emphasize the importance of prioritizing high velocities and accelerations in dynamic manipulation over strictly following the demonstrated path.

In summary, our contributions include:

- A novel IL framework for learning dynamic manipulation of deformable objects with adaptation to hardware constraints using constrained DMPs and refining the object state using quasi-static motions.
- An extensive evaluation of three constrained DMP candidates in a dynamic manipulation setting.
- A new definition of a bag volume metric, and a rim area metric based on α -shapes, for the bag-opening task.
- A thorough empirical evaluation of BILBO's performance and generalization capabilities on a wide range of bags made of different materials, sizes, and shapes.

II. RELATED WORK

A. Learning Dynamic Manipulation of Deformable Objects

Dynamic manipulation of deformable objects has been studied for 1D ropes [10], 2D cloths [1], [5], [10]–[12] and 3D objects like pizza dough [13]. These methods typically utilize simulations [1], [5], [10], [11], [13], [14] thanks to the recent advancement in soft-body physics engines, making it feasible to gather large amounts of data. While prior work has successfully transferred dynamic manipulation policies learned in simulation to the real-world [1], [5], [11], there is a noticeable reality gap in dynamic tasks [15]. This gap is even more pronounced in bag manipulation, where no simulation engine is currently capable of accurately simulating plastic bags [16]. To begin with, physic engines require modeling the complex aerodynamics effects that take place in the real world when dynamically manipulating bags [2]. Furthermore, the perception of plastic bags is challenging due to their translucent and reflective material [16], [17], which are difficult to capture in simulation to ensure successful sim-to-real transfer. Although prior work has demonstrated manipulation of bags in simulation [18], [19], they provide no guarantees of a successful sim-to-real transfer.

B. Bag Manipulation

The research on bag manipulation has mainly focused on two perspectives. The first direction has focused on learning the dynamics of bags [20], [21], where the problem of learning to manipulate the bags is neglected. The second direction has used pre-defined primitives for manipulating the bag. One option for achieving this is to leverage self-supervision methods, where the bags can be marked and segmented, to learn the perception of the bag state, which is used in a pre-defined reasoning scheme to apply pre-defined manipulation primitives [3], [17], [22]. These works dealt

with the bagging problem, that is, opening a bag, inserting items, and lifting it [17]. Other works have studied learning where to direct pick-and-place actions [16], [18], [23]. While many works rely on quasi-static actions [16], [18], [23], some recent works have demonstrated the effectiveness of dynamic manipulation [2], [3]. In [2], the authors proposed a method to open bags by applying an air stream using a blower. Rather than executing pre-defined motions, or utilizing unconventional actuators, in this work we learn a dynamic motion from a single human demonstration, which can effectively open bags using standard grippers.

Another alternative to learn motions that are not pre-defined is using Reinforcement Learning (RL) to train a policy in a simulated environment [19]. However, standard model-free algorithms are unable to succeed in bag manipulation tasks [19]. Due to the aforementioned limitations of bag simulation, and the poor performance of standard RL algorithms for bag manipulation tasks, BILBO uses IL to circumvent these limitations and apply dynamic manipulation to the problem of bag opening. Specifically, we utilize the DMP framework [24] to learn directly from a single human demonstration in the real world, and bypass the perception issues discussed, by tracking robust hand poses rather than markers attached to the bag.

III. BACKGROUND

In this section, we introduce the formulation of DMPs [24]. DMPs is a framework for encoding motions in a non-linear dynamic system. For point-to-point trajectories, DMPs represent the trajectories for a single Degree of Freedom (DoF) by the system of differential equations:

$$\tau \dot{z} = \alpha_z (\beta_z (g - y) - z) + f(x), \quad (1)$$

$$\tau \dot{y} = z, \quad (2)$$

$$\tau \dot{x} = -\alpha_x x. \quad (3)$$

The transformation system, expressed by equations (1) and (2), resembles a mass-spring-damper system augmented with a forcing term $f(x)$ and a time constant τ for scaling the duration of the motion. The equilibrium for the position y in the system is given by g , which thus encodes the goal of the motion. The parameters α_z and β_z are positive constants tuned to achieve critical damping in the system when the forcing term is excluded.

The equation (3), called the canonical system, describes the decay of the phase x , which replaces the explicit time dependence via a positive constant α_x . This implicit representation of time enables synchronization between multiple DoFs or with external systems [24], and makes it possible to slow down or stop the motion in the presence of tracking errors [25]. By augmenting the forcing term $f(x)$ complex trajectories can be encoded in the system. The forcing term is calculated as a weighted sum of H exponential kernels:

$$f(x) = \frac{\sum_{i=1}^H \Psi_i(x) w_i}{\sum_{i=1}^H \Psi_i(x)} x(g - y_0), \quad (4)$$

$$\text{where } \Psi_i(x) = \exp\left(-\frac{1}{2\sigma_i^2}(x - c_i)^2\right), \quad (5)$$

where σ_i and c_i define the widths and centers of the kernels, and w_i are the weights. The kernels are spaced exponentially in phase and, therefore, evenly in time [24]. Fitting the weights is a supervised learning problem that can be solved using any function approximator. For systems with multiple DoFs, the DMPs of each DoF share one canonical system, and only the transformation systems and forcing terms are unique for each DMP [24].

IV. IL FRAMEWORK FOR DYNAMIC MANIPULATION

The proposed framework, shown in Fig. 1, consists of two main components: the dynamic primitive encoded with constrained DMPs, and the quasi-static refinement. The dynamic motions induce a drastic change in the state of a deformable object, which is useful for making fast progress in a task. However, this feature of dynamic manipulation also makes it ill-suited for fine-tuning the object state. This motivates the second stage of our framework, where repeated small-scale changes from quasi-static motions are used to gently refine the state. After each stage, the object state is estimated to compute task performance metrics that are used to decide whether to repeat the stage or proceed to the subsequent step. As the aim of the refinement stage is to improve the result of the dynamic stage, the performance metrics should have stricter requirements for deeming the state sufficient after refinement.

A. Adaptation to Constraints with Constrained DMPs

Given a set of N Cartesian poses with quaternion orientation $\mathbf{P} \in \mathbb{R}^{7 \times N}$ provided as a path in the robot workspace, and robot position constraints, we rely on the assumption that an Inverse Kinematics (IK) solver can produce feasible joint-space targets $\mathbf{Q} \in \mathbb{R}^{D \times N}$, denoted as:

$$\mathbf{Q} = \text{IK}(\mathbf{P}, \underline{\mathbf{q}}, \bar{\mathbf{q}}), \quad (6)$$

where $\underline{\mathbf{q}} \in \mathbb{R}^D$ and $\bar{\mathbf{q}} \in \mathbb{R}^D$ are the lower and upper joint position limits, respectively, and D is the number of DoFs of the robot. The corresponding velocities $\dot{\mathbf{Q}}$ and accelerations $\ddot{\mathbf{Q}}$ can be calculated using the timestamps of the provided path. Then, given \mathbf{Q} , $\dot{\mathbf{Q}}$, $\ddot{\mathbf{Q}}$, and the kinematic constraints denoted with lower and upper bars, the constrained joint-space trajectory \mathbf{Q}^* , $\dot{\mathbf{Q}}^*$, $\ddot{\mathbf{Q}}^* \in \mathbb{R}^{D \times M}$ can be computed with a constrained DMP (CDMP):

$$\mathbf{Q}^*, \dot{\mathbf{Q}}^*, \ddot{\mathbf{Q}}^* = \text{CDMP}(\mathbf{Q}, \dot{\mathbf{Q}}, \ddot{\mathbf{Q}}, \underline{\mathbf{q}}, \bar{\mathbf{q}}, \underline{\dot{\mathbf{q}}}, \bar{\dot{\mathbf{q}}}, \underline{\ddot{\mathbf{q}}}, \bar{\ddot{\mathbf{q}}}). \quad (7)$$

Note that the dimensions of the constrained trajectory matrices may differ from the unconstrained ones, due to augmented duration, velocity, and acceleration.

The original DMP formulation described in Section III does not support the encoding of constraints. Here, we present three alternatives that extend DMPs with this functionality and can be used for dynamic manipulation. To the best of our knowledge, our work is the first to use constrained DMPs for dynamic manipulation of deformable objects.

One straightforward solution is to adjust the time constant τ of the DMP, which we refer to as tau-DMP. This solution retains the path shape that a non-constrained DMP would

produce, and enforces velocity and acceleration constraints by gradually increasing the value of τ to uniformly slow down the motion until the constraints are satisfied. Consequently, the method can only be used in offline settings.

Recently, more advanced modifications have been presented that are capable of adapting the trajectory to the constraints online [8], [9] and encode additional constraints like via-points and obstacles [9]. The constrained DMP formulation in [8] is provided using temporal coupling, that is, online scaling of τ to proactively scale it before the velocity and acceleration limits are exceeded. We refer to this method as TC-DMP. One downside of this method is that it does not ensure that acceleration constraints are guaranteed in the online setting presented in [8]. In offline settings, the acceleration constraints can be guaranteed by appropriate tuning of the model parameters via trial-and-error.

The third method we consider formulates the problem of fitting the weights of a DMP with constraints as an optimization problem [9], which we refer to as Opt-DMP. In order to express position, velocity, and acceleration at any time point as affine functions in the DMP weights, Opt-DMP follows the equations from [26], thus removing the need for explicit integration. This makes it possible to formulate Opt-DMP as a solution to a Quadratic Program. As the constraints are no longer enforced by scaling τ the generated path might differ from the demonstration. Therefore, position constraints are explicitly part of the optimization problem in addition to the velocity and acceleration constraints.

B. Quasi-Static Refinement Motions

In contrast to dynamic motions which affect the entire state of an object, quasi-static motions can be designed to fine-tune a deformable object state via local, small, directed changes. We consider quasi-static motions that are reversible, while the effects of dynamic manipulation can typically not be undone. Formally, we assume that for every action a there exists a reverse action a^- such that the state dynamics $s' = f(s, a)$ follow $f(f(s, a), a^-) = s$. Furthermore, we assume the actions to cause only small changes in the object state $\|s - f(s, a)\| < \delta_1, \forall a$. Thus, the reversible action space makes quasi-static motions suitable for stable refinement of the state.

To cast the refinement as an optimization problem, we assume there exists a continuous cost function $C(s)$ that measures the quality of the state. Thus, we want to incrementally refine the state by applying actions such that the cost function is minimized:

$$a = \arg \min_a C(s'). \quad (8)$$

Assuming the cost function is convex within the region of attraction given by the initial state of the local refinement phase, continuing the local refinements allows converging to a minimum. This follows from the fact that any action will have a bounded effect on the cost $\|s - f(s, a)\| < \delta_1 \Rightarrow |C(s) - C(s')| < \delta_2$.

V. BILBO: BIMANUAL DYNAMIC MANIPULATION USING IMITATION LEARNING FOR BAG OPENING

We present BILBO, a system that uses the proposed framework for opening bags via dynamic manipulation using a bimanual system, as depicted in Fig. 1. Here, we assume that non-dynamic steps such as grasping and item insertion can be solved by existing methods [3], [17], [22]. First, we start by gathering a single human demonstration of a dynamic motion using a motion capture system. Then, we utilize the aforementioned constrained DMP approaches to augment the joint-space trajectory corresponding to the demonstration, so that it adheres to the robots' constraints. The learned constrained motion is used to dynamically improve the volume and area of the bag. As the ideal bag state should not only be expanded but also have a round opening, we design a quasi-static refinement motion for improving the rim elongation. For evaluating the bag-opening performance and deciding whether the learned dynamic manipulation primitive or the linear refinement motion needs to be applied, we estimate the bag state using a motion capture system, and extract area, volume, and elongation metrics for the bag.

A. Learning a Dynamic Motion from Human Demonstration

The human demonstration is captured using a motion capture system that tracks the hands of a human. The motion range of the demonstration should be within the robot's reach, or its amplitude must be scaled down. At each time step of the demonstration, the distance between the hands is calculated, as a percentage of the maximum observed distance, and used to define the desired distance between the robot grippers when replaying the motion for bags of different widths. Any rotation that is not around the main rotation axis of the motion is filtered out. Then, the trajectory is smoothed to account for noise in the captured trajectory. Finally, the trajectory is converted to joint space to apply the constrained DMP approaches described in Section IV-A.

B. Evaluating the Task Performance

First, to define the performance metrics for the bag-opening task, we specify how to measure the state of the bag and its regions of interest. To perceive the state, we employ a perception scheme based on reflective markers and a motion capture system. Here, we capture the bag state after either the dynamic or refinement motion have finished, which is more straightforward than tracking the bag state during its motion, especially if the motion is dynamic.

Each bag has two regions of reflective markers, one around the rim and another near the bottom. In addition, the inside of the rim has markers to facilitate detecting the rim even if it is folded. The marker points of the bag are defined as $\mathbf{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_L\} \in \mathbb{R}^3$, where L is the number of points in Cartesian space provided by the motion capture system. We then filter out the bag points \mathbf{P}^F by removing outliers based on their relative position, and identify the points $\mathbf{P}^R \in \mathbf{P}^F$ belonging to the rim¹.

¹For more details see our implementation <https://sites.google.com/view/bilbo-bag>.

The optimal state of an opened bag is characterized by a large, round rim and an expanded volume, which is required for downstream tasks such as item insertion. To quantify this state, we introduce three bag metrics that measure the bag volume, rim elongation, and rim area. These metrics utilize a convex hull function $CH(\cdot)$ of the bag points. The bag volume is defined as the 3D convex hull of the filtered points

$$V = \text{Volume}(\text{CH}_{3D}(\mathbf{P}^F)). \quad (9)$$

To measure the elongation, we first extract the Principal Component Analysis (PCA) axes from the 2D convex hull of the rim points

$$\lambda_1, \lambda_2, \mathbf{v}_1, \mathbf{v}_2 = \text{PCA}(\text{Vertices}(\text{CH}_{2D}(\mathbf{P}_{x,y}^R))). \quad (10)$$

Note that we use subscripts to denote that values for specific coordinates are extracted throughout this section. The elongation metric is then defined as the ratio of the PCA axes so that the length of the axis primarily directed along the y-axis of the robot frame is divided by the length of the axis primarily directed along the x-axis (see Fig. 1 for the axes)

$$E = \begin{cases} \sqrt{\frac{\lambda_2}{\lambda_1}} & \text{if } |\mathbf{v}_{1,x}| > |\mathbf{v}_{2,x}|, \\ \sqrt{\frac{\lambda_1}{\lambda_2}} & \text{otherwise.} \end{cases} \quad (11)$$

This definition of the elongation metric is beneficial for the refinement of the elongation using the robot system, in contrast to defining the elongation as the ratio between the major and minor PCA axes [17] which does not take into account the direction of the elongation.

Finally, the area metric is defined by fitting an alpha-shape [27] to the rim points

$$A = \text{Area}(\text{AlphaShape}(\mathbf{P}_{x,y}^R, \alpha)). \quad (12)$$

To accurately define multiple bag openings we need different α values. Thus, we define α as a linear function of the 2D convex hull of the rim

$$\alpha = k_\alpha * \text{Area}(\text{CH}_{2D}(\mathbf{P}_{x,y}^R)) + b_\alpha, \quad (13)$$

where k_α and b_α were found empirically for bags of different sizes. Our definition of the rim area provides better estimates than directly using the convex hull [17], which can lead to an overestimation of the area when the rim is slim and bent (see Fig. 2 a). However, alpha-shapes can underestimate the opening area for some bag configurations, e.g., when points of the bag are occluded (see Fig. 2 b). Nevertheless, overestimating the area is a more severe problem as it is prone to early task termination in states which are actually insufficient, whereas underestimating the area only results in superfluous motions for improving the area metric.

C. Elongation Refinement Motion

The dynamic motion from the human demonstrator targets only a high volume and rim area, without refining the roundness of the opening. Consequently, as suggested in Section IV-B we design a cost function on the elongation metric defined in equation (11). The optimal elongation




Fig. 2: Estimation of the opening area using the α -shape (orange), and convex hull (black). In a) the convex hull overestimates the rim area. In b) the α -shape underestimates the area due to the occlusion of markers.

$E = 1$ indicates a perfectly round opening and stretching in any direction is undesirable. Therefore, we define the cost function as the distance to optimal elongation: $\Delta\text{Elongation} = C(s) = |1 - E(s)|$.

As actions, we use linear, reversible motions that adjust the distance between the grippers along the x -axis by a pre-defined step. We implement the action choice in equation (8) with a state-machine based on the relationship between the gripper distance and the elongation of the bag. In case that $E < 1$, the bag is elongated primarily along the x -axis of the robot frame, and the refinement will decrease the distance between the grippers. On the other hand, if $E > 1$, the bag is elongated along the y -axis and the motion will increase the grippers' distance. Additionally, to prevent collisions and tearing apart the bag, we defined a maximum and minimum allowed distance along the x -axis.

VI. EXPERIMENTAL SET-UP

Our experiments are designed to answer the following research questions:

- What is the performance of different constrained DMPs for dynamic manipulation of deformable objects?
- Does including a quasi-static refinement stage improve the manipulation performance?
- How well does BILBO generalize to bags that are significantly larger than the one used to record the human demonstration?
- How does the difficulty of opening a bag relate to the size and stiffness of the material?

Our dual-arm setup consists of two Franka Emika's robots and an OptiTrack motion capture system for perception. Joint velocity control is used for the dynamic motion primitive, and Cartesian pose control is used for the refinement motion. Each experiment was repeated 10 times for each combination of bags, manipulation approach, and initial bag configuration. In total, 240 evaluation runs were gathered in the experiments. For a 5 s demonstration, the duration of a dynamic fling lasts 5 s (Opt-DMP), ~ 8 s (TC-DMP) and ~ 16 -20 s (tau-DMP). A refinement motion lasts ~ 1 s.

For the constrained DMP implementations, we set the constraints to 98% of the strictest position, velocity, and acceleration limits for each joint specified in [28] to leave a margin for numerical inaccuracies. Both τ of tau-DMP and γ_a of TC-DMP, a parameter which balances the trade-off between avoiding limit violations and increasing slowdown, were set by gradually increasing their values until no limits



Fig. 3: Examples of two initial bag configurations used in the experiments and an example the final state after BILBO.




Fig. 4: Plastic bags used in the experiments.

were exceeded. In Opt-DMP, a parameter λ is used to toggle between optimizing with respect to position and velocity. In our case, we optimize with respect to the position to reduce the distortion from the demonstrated trajectory, which is verified to be safe, and therefore lower the risk of collision between the robots. All kinematic constraints are enforced regardless of the optimization objective. To ensure a fair comparison between the DMP methods, we tuned the parameters so that they produced comparable trajectories when no constraints were applied.

A. Bags and Initial Configuration

In our experiments, we used five plastic bags with different material properties and sizes (see Fig. 4). To provide context on the size of the evaluated bags, prior work [17], [22] has demonstrated quasi-static manipulation of bags in the ranges 28-30 cm by 49-54 cm and 32-55 cm by 32-55 cm, respectively, while [3] dynamically manipulated bags in the range of 25-35 cm by 40-53 cm.

Regarding the material, bag A is made of biodegradable plastic, making it especially soft, while the other bags are made of polyethylene. Moreover, bags C and E contain drawstrings, which increases the rigidity of the rim. The stiffness of the bags, as perceived by a human evaluator, listed from soft to stiff, is as follows: A, C, E, D, B. Note that only bag A was used to record the human demonstration.

We define two initial states of the bags to evaluate the task performance, an easy state and a hard state (see Fig. 3). In both cases, we assume that the bags are gripped symmetrically by the robots. In the easy state, the bag is hanging extended with the rim closed. The hard state has the bottom of the bag crumpled, folded inward, and the rim pressed shut.

B. Bag Opening Evaluation

In our experiments, we evaluated the bag-opening task using the bag metrics proposed in Section V-B. The area and volume metrics are represented as the ratio between the achieved value and the observed value in a successful human demonstration with the same bag. The area and volume are considered sufficiently high if they reach the lower limits of 60% and 70%, respectively, where we consider a bag to be

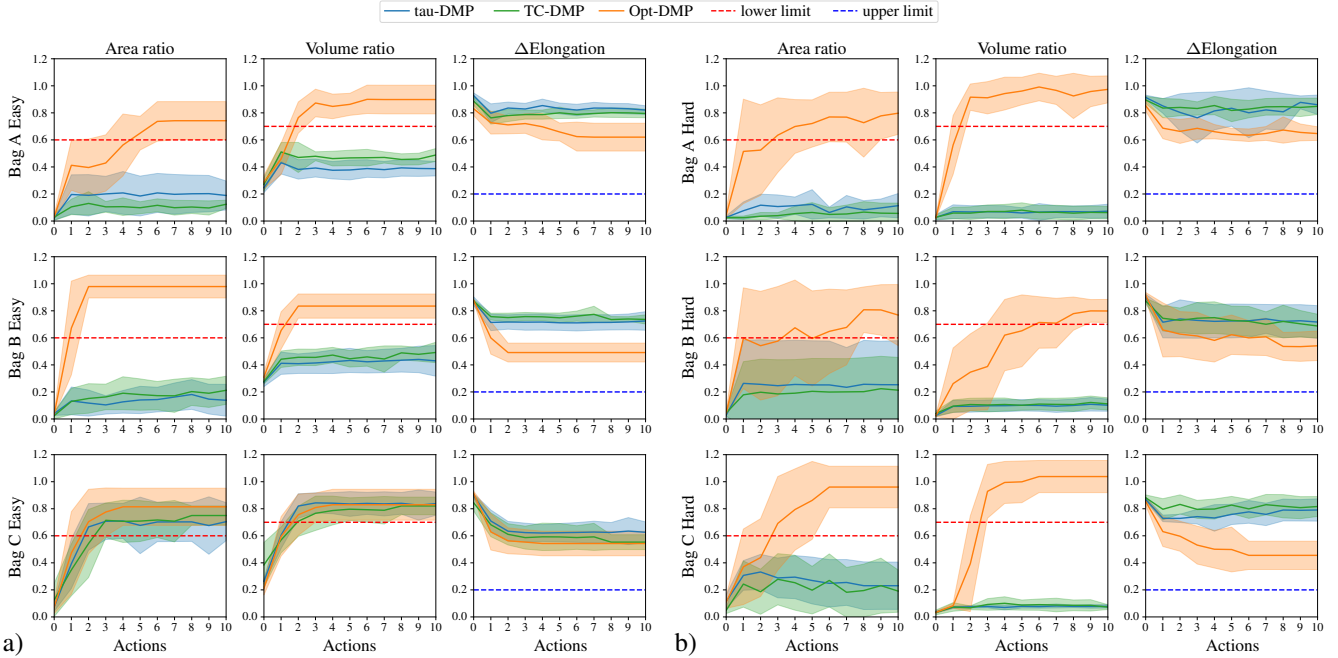


Fig. 5: Quantitative results of the learned dynamic primitives using Opt-DMP, tau-DMP, and TC-DMP for three different bags in a) easy and b) hard initial configuration. The results show the area, volume, elongation and their target limits.

open if these targets are reached. We use a more strict volume target because high volumes were expected to be easier to achieve and maintain compared to high areas of the more unstable rim.

In addition, we study the elongation during the refinement stage via the cost function defined in Section V-C. In the experiments, we use an upper limit for the cost function, which is deemed sufficiently low when $\Delta\text{Elongation} \leq 0.2$. This relatively strict upper limit is selected so that the manipulation would not terminate early at an easier target value, and the system would instead use any remaining actions to strive for a better rim elongation.

The refinement action is reversed if it causes the area or volume metrics to fall below the target levels, and in such cases, the refinement will terminate after the reversal. The reason for this rule is that the area and volume, are treated as primary objectives, which should not be compromised for better elongation, i.e., the secondary objective. The manipulation scheme terminates if the target elongation is met without reducing the area or volume metrics below their targets, or if the total number of allocated actions is reached.

VII. EXPERIMENTAL RESULTS

A. Analysis of Constrained DMP methods

In this experiment, we studied the performance of the dynamic manipulation primitive learned with each of the constrained DMP versions. We evaluated bags A, B, and C in both easy and hard initial configurations. The dynamic motion was repeated until the targets set for the area and volume metrics were achieved or a maximum of 10 actions was reached. Note that runs that needed less than 10 actions to reach the targets have the unused steps padded with the final values in the figures.

As shown in Fig. 5, Opt-DMP consistently exceeds the area and volume targets for each bag, while other methods only succeeded in reaching the targets for bag C from the easy initial state. Although the hard initial state typically requires more actions compared to the easy state, Opt-DMP was able to frequently achieve the targets with only 2-6 actions, depending on the bag and the difficulty of the initial state. This demonstrates the effectiveness of the dynamic motion. As the performance of TC-DMP is comparable to tau-DMP, the only benefit of TC-DMP in the setting we consider is faster runtime, as τ is adaptive rather than fixed.

A key reason for the superior performance of Opt-DMP compared to the τ -scaling methods is that Opt-DMP enables high peak velocities and accelerations in each joint independently, at the cost of distorting the path, as the constraints are encoded in the DMP weights. In contrast, the tau-DMP and TC-DMP tend to slow down all joints via the shared canonical system, in case any joint exceeds the limits, as they enforce the constraints by scaling τ . This result highlights that the ability of Opt-DMP to independently gain high peak velocities and accelerations in each joint is more important for dynamically manipulating the bag than an accurate reproduction of the demonstrated path.

In addition, the results show the relationship between the structure and stiffness of the bag and its opening performance. Bag A, which was used to provide the human demonstration, ended up being relatively difficult to open as the rim would easily collapse due to the softness of the material. Due to the stiffness of bag B, it needed relatively many repetitions of the dynamic motion to uncrumple the bag initialized in the hard state and reach the target volume. In contrast, the medium stiffness and drawstring in the rim




Fig. 6: Area, volume and elongation results using only the refinement motion for bag A initialized in the hard state.

of bag C made it the easiest bag to open in this experiment, despite its significant depth compared to bags A and B.

Note that the dynamic motion can only slightly enhance the elongation. However, the final elongation is far from the target value. This motivates the proposed refinement motion, since further refinement of the elongation is necessary.

B. Evaluating the Elongation Refinement Motion

In this experiment, we studied the performance of the refinement motion for improving the elongation metric. In addition, we evaluated the bag-opening performance of the refinement motion alone to rule out the possibility that it alone is sufficient without dynamic manipulation. In this experiment, bag A was initialized in the hard state, where the refinement motion was applied until a maximum of 20 actions were performed or the area and volume targets were reached. While BILBO simply prevents motion outside a range we defined to prevent collisions or the bags tearing, in this experiment we had the robots take a step in the opposite direction if they were commanded past these limits. This allowed the robots to keep adjusting the bag, and potentially uncrumple it, instead of terminating the run.

The results in Fig. 6 show that the refinement motion can slightly improve the elongation, suggesting that the motion can be used to improve the roundness of the bag opening. However, the refinement motion neither manages to uncrumple the bag nor approaches the area or volume limits, which highlights the necessity of using the dynamic primitive. Therefore, the optimal system should combine the best-performing dynamic primitive obtained from the Opt-DMP with the non-dynamic refinement scheme.

C. BILBO’s Bag-Opening and Generalization Results

Finally, we evaluated the performance of our proposed system BILBO, which combines the dynamic primitive encoded by Opt-DMP with the linear refinement. First, we evaluated the performance of the system using bags A, B, and C, which have similar widths and varying materials. A maximum of 20 total actions were allocated for both the dynamic and refinement motions. The dynamic primitive was first applied until the target area and volume were reached. Then, if the maximum number of actions had not been reached, the refinement primitive would be subsequently applied based on the measured elongation. An example of the final state of a successful run of BILBO is shown in Fig. 3 c.

The performance of both BILBO and Opt-DMP from the dynamic primitive experiment is shown in Fig. 7. The results

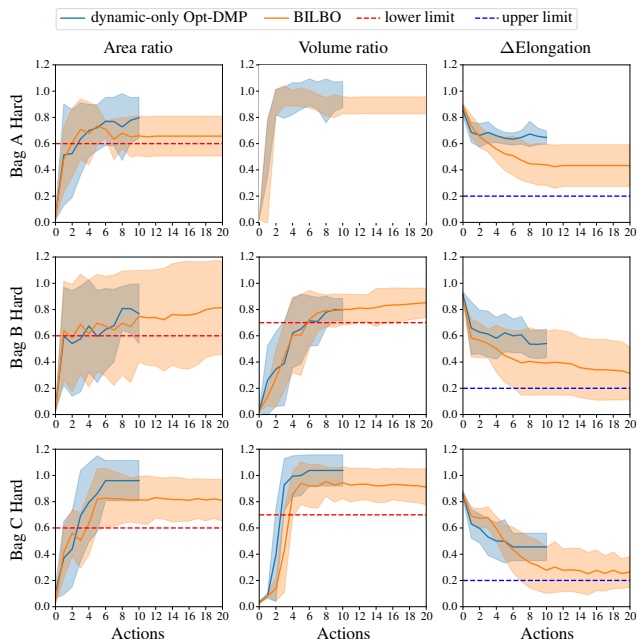


Fig. 7: Area, volume, and elongation results of BILBO for bags A-C initialized in the hard state. The results from the dynamic primitive experiment are also shown as a reference, denoted as dynamic-only Opt-DMP.

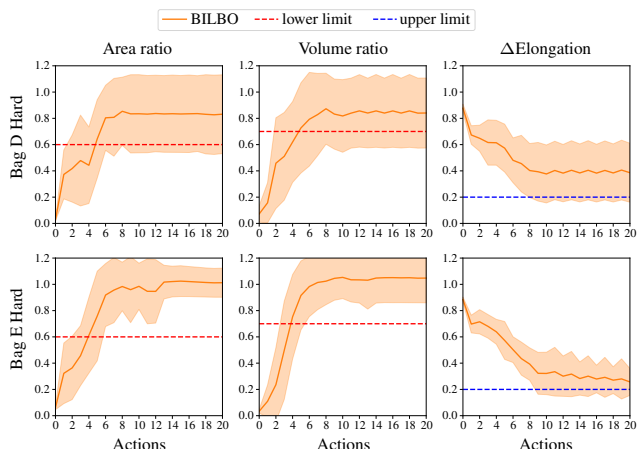


Fig. 8: Area, volume, and elongation results of BILBO for bags D & E initialized in the hard state.

show that BILBO’s refinement motion is able to keep the target state achieved by the dynamic motion, as the area and velocity targets are reached at a level similar to that of the dynamic-only experiment, even for bag A. Furthermore, the refinement successfully decreases the distance to optimal elongation. The target elongation is more challenging to reach for bag A, as the soft material causes the rim to collapse easily. In contrast, it is not uncommon for bags B and C to reach the target. One thing to note is that the mean values of Δ Elongation do not fall below the target level for any bag. The reason for this is that the target was intentionally chosen as a strict value to encourage finding the best reachable elongation as explained in Section VI-B. Therefore, Δ Elongation converges to the best reachable level for each bag in our experiments.

Additionally, we evaluated BILBO’s capability to gen-

eralize to significantly larger bags than the one used for demonstration by conducting the same experiment on Bag D and Bag E. The results in Fig. 8 show that BILBO yields similar performance on bags D and E as it does on bags A-C. Even with significantly larger bags, BILBO consistently surpasses the target values for both area and volume metrics, maintaining stability during refinement, while the Δ Elongation converges to a bag-specific lowest level.

All in all, the results show that BILBO is capable of successfully opening different bags that vary in size and material with just one human demonstration for a single bag, highlighting both the generalizability and efficiency of the proposed method. One possible extension is to extend BILBO to the full bagging task, including gripping and item insertion, thus developing the *BILBO Bagging* system.

VIII. CONCLUSIONS

We presented an IL framework that first learns a dynamic primitive for inducing significant changes in the states of deformable objects with few motions, while adhering to robot constraints with constrained DMPs. Subsequently, it employs stable quasi-static motions for small-scale adjustment. We assessed the framework in the task of bimanual bag opening, resulting in a system named BILBO.

First, we evaluated three constrained DMP versions for generating a dynamic primitive that adheres to the robots' limits. Our experiments showed that Opt-DMP achieved superior performance, which demonstrates that in the dynamic manipulation task, it is more important to achieve high velocities and accelerations than to accurately reproduce the demonstrated path. This highlights the necessity of selecting a constraint satisfaction approach that prioritizes attributes most relevant to the task. Due to the benefits of Opt-DMP, it was selected for learning the dynamic motion in the BILBO system using a human demonstration with a single bag. Our experimental results showed that BILBO was able to effectively open bags of different materials, shapes, and sizes than the one used for the demonstration. More specifically, BILBO was able to open most bags in less than 5 actions and achieve near-optimal elongation after 20 actions.

As future directions, the learned dynamic motions could be used to bootstrap learning methods that refine the motion by, for example, adapting the DMP weights, without violating the kinematic constraints. This could potentially lead to greater efficiency in dynamic manipulation tasks.

REFERENCES

- [1] J. Hietala, D. Blanco-Mulero, G. Alcan, and V. Kyrki, "Learning visual feedback control for dynamic cloth folding," in *2022 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2022, pp. 1455–1462.
- [2] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Dextairity: Deformable manipulation can be a breeze," in *Proc. of Robots: Sci. Syst. (RSS)*, 2022.
- [3] N. Gu, Z. Zhang, R. He, and L. Yu, "Shakingbot: dynamic manipulation for bagging," *Robotica*, vol. 42, no. 3, p. 775–791, 2024.
- [4] M. Mason and K. Lynch, "Dynamic manipulation," in *1993 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, vol. 1, 1993, pp. 152–159.
- [5] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Proc. of the 5th Conf. Robot Learn.*, vol. 164. PMLR, 08–11 Nov 2022, pp. 24–33.

- [6] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [7] C. L. Nehaniv and K. Dautenhahn, *The Correspondence Problem*. Cambridge, MA, USA: MIT Press, 2002, p. 41–61.
- [8] A. Dahlin and Y. Karayiannidis, "Temporal coupling of dynamical movement primitives for constrained velocities and accelerations," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 2233–2239, 2021.
- [9] A. Sidiropoulos, D. Papageorgiou, and Z. Doulgeri, "A novel framework for generalizing dynamic movement primitives under kinematic constraints," *Autonomous Robots*, vol. 47, no. 1, pp. 37–50, 2023.
- [10] C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Iterative residual policy for goal-conditioned dynamic manipulation of deformable objects," in *Proc. of Robots: Sci. Syst. (RSS)*, 2022.
- [11] D. Blanco-Mulero, G. Alcan, F. J. Abu-Dakka, and V. Kyrki, "Qdp: Learning to sequentially optimise quasi-static and dynamic manipulation primitives for robotic cloth manipulation," in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2023, pp. 984–991.
- [12] B. Balaguer and S. Carpin, "Combining imitation and reinforcement learning to fold deformable planar objects," in *2011 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2011, pp. 1405–1412.
- [13] A. C. Satici, F. Ruggiero, V. Lippiello, and B. Siciliano, "A coordinate-free framework for robotic pizza tossing and catching," in *2016 IEEE Int. Conf. Robot. Automat. (ICRA)*, 2016, pp. 3932–3939.
- [14] H. Zhang, J. Ichnowski, D. Seita, J. Wang, H. Huang, and K. Goldberg, "Robots of the lost arc: Self-supervised learning to dynamically manipulate fixed-endpoint cables," in *2021 IEEE Int. Conf. Robot. Automat. (ICRA)*, 2021, pp. 4560–4567.
- [15] D. Blanco-Mulero, O. Barbany, G. Alcan, A. Colomé, C. Torras, and V. Kyrki, "Benchmarking the sim-to-real gap in cloth manipulation," *IEEE Robot. Automat. Lett.*, pp. 1–8, 2024.
- [16] C. Gao, Z. Li, H. Gao, and F. Chen, "Iterative interactive modeling for knotting plastic bags," in *Proc. of The 6th Conf. Robot Learn.*, vol. 205. PMLR, 14–18 Dec 2023, pp. 571–582.
- [17] L. Y. Chen, B. Shi, D. Seita, R. Cheng, T. Kollar, D. Held, and K. Goldberg, "Autobag: Learning to open plastic bags and insert objects," in *2023 IEEE Int. Conf. Robot. Automat. (ICRA)*, 2023, pp. 3918–3925.
- [18] D. Seita, P. Florence, J. Tompson, E. Coumans, V. Sindhwani, K. Goldberg, and A. Zeng, "Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks," in *2021 IEEE Int. Conf. Robot. Automat. (ICRA)*, 2021, pp. 4568–4575.
- [19] R. Antonova, P. Shi, H. Yin, Z. Weng, and D. K. Jensfelt, "Dynamic environments with deformable objects," in *35th Conf. on Neural Info. Processing Systems Datasets and Benchmarks Track*, 2021.
- [20] Z. Weng, F. Paus, A. Varava, H. Yin, T. Asfour, and D. Kragic, "Graph-based task-specific prediction models for interactions between deformable and rigid objects," in *2021 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2021, pp. 5741–5748.
- [21] R. Antonova, A. Varava, P. Shi, J. F. Carvalho, and D. Kragic, "Sequential topological representations for predictive models of deformable objects," in *Proc. of the 3rd Conf. on Learning for Dynamics and Control*, vol. 144. PMLR, 07 – 08 June 2021, pp. 348–360.
- [22] L. Y. Chen, B. Shi, R. Lin, D. Seita, A. Ahmad, R. Cheng, T. Kollar, D. Held, and K. Goldberg, "Bagging by learning to singulate layers using interactive perception," in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2023, pp. 3176–3183.
- [23] A. Bahety, S. Jain, H. Ha, N. Hager, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Bag all you need: Learning a generalizable bagging strategy for heterogeneous objects," in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2023, pp. 960–967.
- [24] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical Movement Primitives: Learning Attractor Models for Motor Behaviors," *Neural Comput.*, vol. 25, no. 2, pp. 328–373, 2013.
- [25] A. J. Ijspeert, J. Nakanishi, and S. Schaal, "Movement imitation with nonlinear dynamical systems in humanoid robots," in *Proc 2002 IEEE Int. Conf. Robot. Automat. (ICRA)*, vol. 2, 2002, pp. 1398–1403.
- [26] A. Sidiropoulos and Z. Doulgeri, "A reversible dynamic movement primitive formulation," in *2021 IEEE Int. Conf. Robot. Automat. (ICRA)*, 2021, pp. 3147–3153.
- [27] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel, "On the shape of a set of points in the plane," *IEEE Transactions on Information Theory*, vol. 29, no. 4, pp. 551–559, 1983.
- [28] "Robot and interface specifications," <https://frankaemika.github.io/docs/control.parameters.html>, accessed: 2023-11-26.