

DA-MMP: Learning Coordinated and Accurate Throwing with Dynamics-Aware Motion Manifold Primitives

Chi Chu¹, Huazhe Xu^{2,1}

Abstract—Dynamic manipulation is a key capability for advancing robot performance, enabling skills such as tossing. While recent learning-based approaches have pushed the field forward, most methods still rely on manually designed action parameterizations, limiting their ability to produce the highly coordinated motions required in complex tasks. Motion planning can generate feasible trajectories, but the dynamics gap—stemming from control inaccuracies, contact uncertainties, and aerodynamic effects—often causes large deviations between planned and executed trajectories. In this work, we propose Dynamics-Aware Motion Manifold Primitives (DA-MMP), a motion generation framework for goal-conditioned dynamic manipulation, and instantiate it on a challenging real-world ring-tossing task. Our approach extends motion manifold primitives to variable-length trajectories through a compact parameterization and learns a high-quality manifold from a large-scale dataset of planned motions. Building on this manifold, a conditional flow matching model is trained in the latent space with a small set of real-world trials, enabling the generation of throwing trajectories that account for execution dynamics. Experiments show that our method can generate coordinated and smooth motion trajectories for the ring-tossing task. In real-world evaluations, it achieves high success rates and even surpasses the performance of trained human experts. Moreover, it generalizes to novel targets beyond the training range, indicating that it successfully learns the underlying trajectory-dynamics mapping.

I. INTRODUCTION

Dynamic manipulation [1], [2]—the ability of robots to purposefully interact with objects through high-speed motions—has recently attracted growing interest [3]–[6]. While many manipulation tasks can be accomplished under assumptions of negligible object dynamics, dynamic manipulation requires leveraging object motion and momentum to achieve the goal. This setting poses two key challenges: coping with intricate, hard-to-model dynamics and generating trajectories that effectively exploit velocity and momentum to achieve complex tasks.

Recent dynamic manipulation systems [7]–[10] primarily use learning to model or compensate for complex task dynamics, while the motion generation process is typically based on hand-crafted models or manually designed action parameterizations. This limits their applicability to more complex tasks that require complicated and highly coordinated motion sequences. In dynamic manipulation research, movement primitives (MPs) have been widely used to learn and generate complex motions from demonstrations or planned trajectories [11]–[19]. By parameterizing motions




Fig. 1. Real-world ring tossing with trajectories synthesized by DA-MMP. The first line shows the overall sequence, where superimposed snapshots (light to dark) indicate temporal order of the coordinated throw. The second line illustrates three representative frames of the hitting process.

with a small set of basis-function weights, MPs enable compact representation and efficient modulation of key attributes such as speed, amplitude, and timing, making them particularly suitable for high-speed, dynamic tasks. Building on this idea, motion manifold primitives (MMPs) [20]–[24] replace hand-crafted parameter spaces with low-dimensional manifolds learned directly from large collections of feasible trajectories. While recent works have extended MMPs to tasks such as tossing [24], they generally assume fixed-length trajectories, are trained on relatively modest-scale datasets, and have not been fully explored in challenging real-world dynamic manipulation scenarios.

In this work, we bridge these two lines of research by introducing Dynamics-Aware Motion Manifold Primitives (DA-MMP), a motion generation framework for complex dynamic manipulation that integrates motion manifold primitives with large-scale planning data and trajectory-level dynamics learning. We focus on a challenging real-world ring-tossing task, as illustrated in Fig. 1, that requires generating highly coordinated motions to throw a ring around a distant target peg. This setting is particularly difficult due to multiple sources of dynamics gaps, including control inaccuracies of the arm (control gaps), slippage and uncertainty in release timing (contact gaps), and aerodynamic effects during flight (aerodynamic gaps). Our framework extends MMPs to variable-length trajectories via a compact parameterization

¹ Shanghai Qi Zhi Institute; ² Institute for Interdisciplinary Information Sciences, Tsinghua University.

Contact: cc2018011855@gmail.com, huazhe_xu@mail.tsinghua.edu.cn

that includes execution duration, and leverages a large-scale dataset of planned trajectories generated by sampling-based motion planning to learn a high-quality motion manifold. A conditional generative model is then trained in the learned latent space to generate throwing trajectories. By training on the observed landing outcomes of executed trajectories, this model implicitly captures the relationship between trajectories and their execution dynamics. Experiments demonstrate that our approach achieves high real-world success rates, even rivaling trained human experts, and can generalize to unseen targets, which indicates that it generates valid trajectories and effectively captures task dynamics.

Our work makes the following contributions at the intersection of dynamic manipulation and motion manifold learning:

- **Challenging dynamic manipulation task.** We design and address a real-world ring-tossing task that requires generating complex, highly coordinated motions, providing a demanding benchmark for evaluating expressive motion generation.
- **Variable-length motion manifold primitives with dynamics learning.** We extend the MMP framework to represent variable-length trajectories through a compact parameterization that includes execution duration, and introduce a conditional generative model in the latent space that captures the relationship between trajectories and their execution dynamics.
- **Large-scale dataset and efficient real-world learning.** Leveraging sampling-based motion planning, we collect a large-scale dataset of feasible trajectories to learn a high-quality motion manifold, and further enable data-efficient learning from limited real-world trials.

II. RELATED WORK

Our approach builds on and connects two main lines of research. The first line of research is motion manifold primitives, which learn low-dimensional latent representations of feasible trajectories. The second line is learning-based goal-conditioned dynamic manipulation, which applies machine learning to address the challenges of high-speed, target-directed manipulation. We review both directions in the following subsections.

A. Motion Manifold Primitives

Motion Manifold Primitives [21] build upon the idea of movement primitives [11], [14], [15], which offer structured representations for generating robot motions. Instead of relying on hand-crafted parameter spaces, MMPs learn low-dimensional manifolds directly from large collections of feasible trajectories. A closely related idea was explored by Noseworthy et al. [20], who used task-conditioned variational autoencoders to learn latent spaces of movement primitives. Subsequent work has expanded MMPs in several directions, including conditional generative models for task-specific motion synthesis [23], extensions with parametric curves for temporal and via-point modulation [22], and a differentiable

formulation that enables further optimization of the learned manifold with downstream objectives [24].

While these efforts have demonstrated MMPs across diverse domains, they typically assume fixed-length trajectories, are trained on relatively modest-scale datasets, and have not been fully explored in challenging real-world dynamic manipulation. In contrast, our work extends MMPs to variable-length trajectories via a compact parameterization, learns high-quality manifolds from large-scale planned motion data without requiring downstream post-optimization, and incorporates a conditional generative model that implicitly captures the relationship between trajectories and their execution dynamics.

B. Learning-based goal-conditioned dynamic manipulation

Recent works have applied deep learning to high-speed, goal-conditioned manipulation tasks with intricate dynamics. Among the extensive literature, we highlight several representative efforts. TossingBot [7] learns residual release velocities conditioned on grasp and object’s visual features to compensate for various dynamics gaps. IRP [8] addresses dynamic manipulation of deformable objects by learning delta dynamics in an iterative residual policy framework, enabling online adaptation to new dynamics and hardware. DartBot [9] leverages tactile sensing and reinforcement learning to exploit contact interactions, learning slippage and release timing in real-world overhead dart throwing. Whole-body throwing with legged manipulators [25] combines a residual control policy to address control gaps with tube acceleration [26] to handle release-time uncertainties. The IPA policy [10] further tackles dynamic manipulation with soft-body tools by implicitly identifying the underlying physics of soft-rigid interactions through system identification. Unlike these approaches, which largely rely on hand-crafted motion generation rules with constrained expressiveness, our method synthesizes high-dimensional trajectories through planning and manifold learning. Moreover, target-level residual learning, as in TossingBot, is unsuitable in our setting since distinct trajectories for the same target can yield different landings; instead, we learn dynamics directly at the trajectory level to capture such differences.

Like our work in spirit, earlier studies also applied movement primitives to dynamic manipulation to generate smooth and consistent trajectories. Representative examples include learning motor primitives for hitting and batting [13], swinging and padding [12], and applying mixtures of motor primitives to table tennis [27], [28]. These approaches relied on adapting from a limited number of demonstrations, whereas our method learns motion manifold primitives from large-scale planned trajectories and uses limited real-world data to capture dynamics.

III. METHOD

In this section, we present DA-MMP, a framework that leverages motion manifold learning for dynamics-aware trajectory generation. The framework consists of two stages, as shown in Fig. 2: (1) data collection and parameterization,

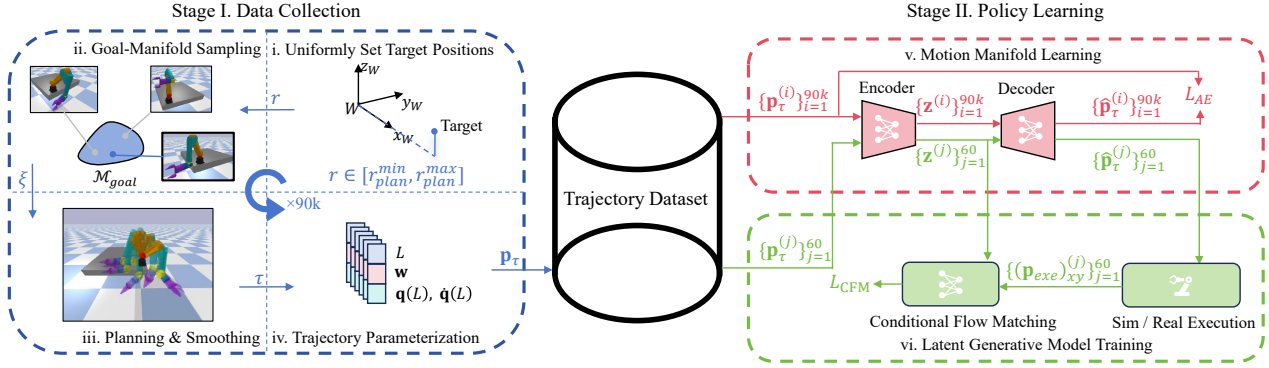


Fig. 2. Overview of DA-MMP. Stage I (blue): data collection from goal-manifold sampling, motion planning, and parameterization, yielding 90k planned trajectories. Stage II (red/green): policy learning with an autoencoder to learn motion manifolds and a conditional flow-matching model for trajectory-dynamics mapping.

and (2) policy learning using a motion manifold representation and a conditional flow-matching model. We begin by introducing the problem formulation, followed by a detailed description of each stage.

A. Problem Formulation

We study goal-conditioned dynamic manipulation tasks that require complex motions, exemplified by a ring-tossing task. In this task, the objective is to throw a ring of radius R_{ring} such that it lands around a vertical cylinder of radius R_{cyl} and height z_{cyl} on the ground. To simplify the problem, we assume a fixed hand-crafted grasp of the ring. A tossing trajectory is defined as

$$\tau = \{\mathbf{q}(t), \dot{\mathbf{q}}(t)\}_{t=0}^L, \quad (1)$$

where $\mathbf{q}(t)$ and $\dot{\mathbf{q}}(t)$ denote the joint positions and velocities at time t , and L is the execution horizon, which can vary across different trajectories. At time $t = L$, the gripper starts opening. During a subsequent settling horizon h_{settle} , the end-effector maintains its linear and angular velocity, allowing sufficient time for the ring to be released from the gripper. Inspired by how humans perform ring tossing, we design the robot to start and end at the same joint configuration with zero initial velocity:

$$\mathbf{q}(0) = \mathbf{q}(L), \quad \dot{\mathbf{q}}(0) = \mathbf{0}. \quad (2)$$

Given the target position, we exploit the rotational symmetry of the manipulator's first joint and represent the target in polar coordinates. Without loss of generality, we fix the polar angle $\alpha = 0$, so that the problem reduces to generating a trajectory $\tau(r)$ for a given radial distance r . In practice, however, the target placement and the ring landing position are not perfectly aligned with $\alpha = 0$, so we represent both using their (x, y) coordinates, where y is nearly zero. We consider a trial as successful if, at height $z = z_{\text{cyl}}$,

$$\|(x_{\text{ring}}(z_{\text{cyl}}), y_{\text{ring}}(z_{\text{cyl}})) - (x_T, y_T)\|_2 \leq R_{\text{ring}} - R_{\text{cyl}}, \quad (3)$$

where $(x_{\text{ring}}(z_{\text{cyl}}), y_{\text{ring}}(z_{\text{cyl}}))$ are the horizontal coordinates of the ring center at $z = z_{\text{cyl}}$, and (x_T, y_T) denotes the horizontal coordinates of the target top center. This serves as

an approximation, since the ring may not remain perfectly horizontal during flight.

B. Data Collection

1) *Goal-Manifold Sampling*: We adopt the idea of goal-manifold sampling from prior work on robotic throwing [29], [30], where projectile equations are used to generate candidate throw states ξ , forming a manifold of feasible goals $\mathcal{M}_{\text{goal}}$ used for motion planning. Each candidate specifies the desired end-effector (EE) pose and velocity at release. If the planner fails to generate a feasible trajectory, a new candidate is drawn until success.

Formally, we parameterize each throw state in the ring space, since the EE state is uniquely determined via a constant grasp transform ${}^E T_R$. The grasp is designed such that the ring center lies along the EE z -axis and the release direction aligns with the EE x -axis, which reduces the likelihood of post-release collisions with the gripper. We denote a ring throw state by ξ_R , consisting of its pose and twist in the world frame:

$$\xi_R = ({}^W T_R, {}^W \mathbf{v}_R, {}^W \boldsymbol{\omega}_R) \in SE(3) \times \mathbb{R}^6,$$

which comprises twelve degrees of freedom. In the following we describe the constraints that define this feasible manifold of throws.

The twelve degrees of freedom of ξ_R are constrained as follows, with reference to the coordinate frames in Fig. 3. Let ${}^W \mathbf{p}_R$ and ${}^W \mathbf{p}_T$ denote the positions of the ring center and the target top center in the world frame W . The orientation is fixed by geometry, where (x_R, y_R, z_R) denote the unit axes of the ring frame R :

$$x_R \parallel ({}^W \mathbf{p}_T - {}^W \mathbf{p}_R)_{xy}, \quad z_R \parallel z_W, \quad y_R = z_R \times x_R.$$

The angular velocity is restricted to spin about the ring normal:

$${}^W \boldsymbol{\omega}_R = (0, 0, \omega_z)^\top,$$

where ω_z is sampled from the range $[\omega_{\text{min}}, \omega_{\text{max}}]$. The ring center position ${}^W \mathbf{p}_R = (x, y, z)$ is sampled within the workspace. In practice we use a set of surrogate variables

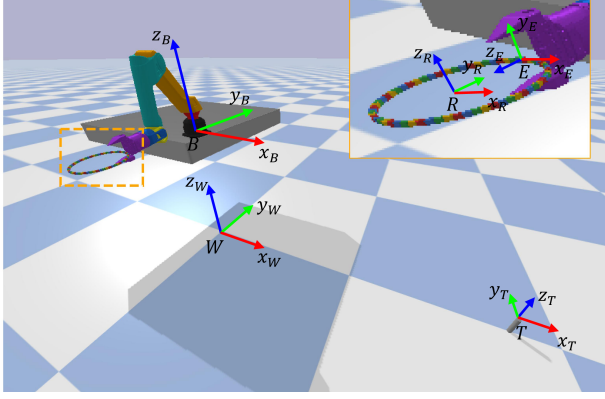


Fig. 3. Coordinate frames used for goal-manifold sampling. World W is fixed to the ground; B attaches to the robot base; the end-effector E rigidly grasps the ring frame R via ${}^E T_R$; the target frame T sits at the cylinder top center. Inset: enlarged view of E and R for clarity.

to generate these samples, but omit the details here for simplicity. For linear velocity, we impose

$$v_z = 0, \quad {}^W \mathbf{v}_R \parallel x_R,$$

where v_z denotes the vertical component of ${}^W \mathbf{v}_R$. Its magnitude is determined by projectile motion to reach the target height $z = z_{\text{cyl}}$:

$$\|{}^W \mathbf{v}_R\| = d_{xy} \sqrt{\frac{g}{2(z_R - z_{\text{cyl}})}}, \quad (4)$$

where $d_{xy} = \|({}^W \mathbf{p}_T - {}^W \mathbf{p}_R)_{xy}\|$ is the horizontal distance to the target, and g is the gravitational acceleration constant. Together these constraints define a low-dimensional manifold of throws. In practice, candidates are drawn from this manifold and checked for feasibility using inverse kinematics and collision detection; if a candidate fails, sampling is repeated until a feasible one is found.

2) *Sampling-Based Motion Planning*: Sampling-based motion planning is renowned for its computational efficiency in high-dimensional spaces [31], which enables us to generate a large-scale dataset of feasible trajectories. Given a sampled throw state, we construct the motion planning problem by setting the goal state to this specified throw state and the initial state to the same joint configuration with zero velocity. We employ DIMIT-RRT as our kinodynamic planning algorithm [32], where each planning attempt is allowed up to N_{planning} samples. Successful plans are further refined with $N_{\text{smoothing}}$ iterations of a trajectory smoothing algorithm based on [33]. We also adopt the quick rejection strategy proposed in [29] to efficiently discard infeasible candidates before planning. During planning, we perform discrete collision checking with a time resolution of Δt_{col} .

After a trajectory is found, we execute it twice in simulation and compare the resulting landing positions; the trajectory is retained only if the deviation between the two trials is below a threshold, thus filtering out unstable throwing trajectories. Finally, the accepted trajectory is interpolated to match the control frequency f_{ctrl} .




Fig. 4. Illustration of basis functions used for trajectory parameterization. Left: normalized Gaussian bases over the phase. Right: gated bases obtained by multiplying the normalized bases with the polynomial gate $(s(1-s))^2$. For illustration, we show $K = 6$.

C. Policy Learning

1) *Variable-Length Parameterization of Motion Trajectories*: A key challenge in trajectory parameterization is that planned trajectories have variable durations and cannot be directly resampled or aligned, since naive interpolation effectively accelerates or decelerates the motion and alters the release velocity that defines the throw state. To handle variable durations without distorting the motion, we include the trajectory length L as an explicit component of the trajectory parameterization. Another challenge is ensuring that the generated trajectories are sufficiently smooth for high-speed execution, which we address by drawing inspiration from movement primitives and representing trajectories with radial basis functions. Formally, each trajectory dimension is approximated as a weighted combination of normalized Gaussian radial basis functions,

$$\begin{aligned} q(s; \mathbf{w}) &= \mathbf{w}^\top \boldsymbol{\phi}(s), \\ \phi_i(s) &= \frac{\kappa_i(s)}{\sum_{j=1}^K \kappa_j(s)}, \quad i = 1, \dots, K, \\ \kappa_i(s) &= \exp\left(-\frac{(s - c_i)^2}{h}\right), \end{aligned} \quad (5)$$

where $s \in [0, 1]$ is the normalized phase, c_i are the basis centers, h is the bandwidth, and \mathbf{w} are the basis weights for this trajectory dimension. Here $\boldsymbol{\phi}(s) = [\phi_1(s), \dots, \phi_K(s)]^\top$ denotes the vector of normalized basis functions.

We find that directly using radial basis functions often leads to oscillations in the end-point velocity. To address this, we adopt via-point movement primitives [15], which explicitly constrain the start- and end-point positions and velocities. Formally, each trajectory is expressed as

$$q(s; \mathbf{w}) = \psi(s) + \mathbf{w}^\top \boldsymbol{\phi}(s), \quad (6)$$

where $\psi(s)$ is a cubic Hermite spline interpolating the boundary conditions $\{q(0), q(1), \dot{q}(0), \dot{q}(1)\}$, and $\boldsymbol{\phi}(s)$ denotes the gated basis functions, obtained by multiplying the normalized Gaussian bases with the polynomial gate function $(s(1-s))^2$. For illustration, we show examples of the normalized Gaussian and gated bases in Fig. 4.

Given a reference trajectory defined on L phase points $\{(q(s_l), \dot{q}(s_l))\}_{l=1}^L$ with $s_l = \frac{l}{L}$, the weights \mathbf{w} are obtained




Fig. 5. Experimental setup of real-world ring tossing.

by least-squares fitting with equal weighting of position and velocity terms,

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_{l=1}^L (q(s_l) - \psi(s_l) - \phi(s_l)^\top \mathbf{w})^2 + \sum_{l=1}^L \left(\frac{1}{\alpha} \dot{q}(s_l) - \psi'(s_l) - \phi'(s_l)^\top \mathbf{w} \right)^2, \quad (7)$$

where $\psi(s_l)$ and $\psi'(s_l) = \frac{d\psi}{ds}(s_l)$ denote the spline and its phase derivative, $\phi(s_l)$ and $\phi'(s_l) = \frac{d\phi}{ds}(s_l)$ are the gated basis functions and their phase derivatives evaluated at s_l , and $\alpha = \frac{ds}{dt} = \frac{1}{L}$ is the time-to-phase scaling factor (so $(1/\alpha)\dot{q} = dq/ds$). Each trajectory is therefore represented by weights \mathbf{w} , the end-point conditions $(q(1), \dot{q}(1))$, and the trajectory length L , which we collectively denote as the parameter vector \mathbf{p}_τ .

2) *Learning Motion Manifold Primitives*: Given the trajectory parameterization \mathbf{p}_τ , we learn motion manifold primitives by first embedding trajectories into a low-dimensional latent space \mathbf{z}_τ using an autoencoder, which captures the underlying manifold structure and reduces dimensionality. The autoencoder is trained to minimize the reconstruction loss

$$\mathcal{L}_{\text{AE}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{p}_\tau^{(i)} - \hat{\mathbf{p}}_\tau^{(i)}\|_2^2, \quad (8)$$

where $\hat{\mathbf{p}}_\tau^{(i)}$ is the decoded reconstruction of $\mathbf{p}_\tau^{(i)}$.

On top of this latent representation, we train a conditional flow-matching model [34]–[37], where the conditioning signal is given by the actual landing point $(x_{\text{exe}}, y_{\text{exe}})$ of the executed trajectory measured at $z = z_{\text{cyl}}$, implemented via classifier-free guidance [38]. The model is optimized with

TABLE I

SUCCESS RATE (SR) COMPARISON OF DIFFERENT BASELINES AND OUR METHOD IN SIMULATION AND REAL-WORLD EXPERIMENTS.

Method	Simulation SR (%)	Real SR (%)
Motion planning (1 attempt)	0.0	13.3
Motion planning (2 attempts)	0.0	23.3
Residual-style correction	93.3	6.7
DA-MMP (Ours)	73.3	60.0*
Human novice	–	13.3
Human expert	–	56.7

*Evaluated on an early implementation version where the trajectory parameterization inadvertently omitted a coefficient (magnitude ~ 1) related to the time-to-phase scaling factor in two terms, causing the executed trajectories to deviate from the planned ones. However, since our method trains on actual executed landing positions rather than planned targets, the model is robust to this systematic bias as long as it does not severely distort the motion or violate hardware safety limits. Correcting this omission in simulation improved the success rate from 53.3% to 73.3%, while the real-world result retains the early evaluation due to hardware setup depreciation.

the standard conditional flow-matching loss

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{u, \mathbf{z}_\tau, \mathbf{z}_{\text{noise}}} [\|v_\theta(\mathbf{z}(u), u, \mathbf{c}) - v^*(u)\|_2^2], \quad (9)$$

where $u \in [0, 1]$ denotes the flow time parameter, $\mathbf{z}(u)$ is the linear interpolation between a Gaussian prior sample $\mathbf{z}_{\text{noise}}$ and the ground-truth latent variable of the trajectory \mathbf{z}_τ , $v^*(u)$ denotes the corresponding target velocity field, and $\mathbf{c} = (x_{\text{exe}}, y_{\text{exe}})$ is the conditioning signal.

At inference time, given the conditioning signal $\mathbf{c} = (x_T, y_T)$, a noisy latent sample $\mathbf{z}_{\text{noise}}$ is progressively transported to a trajectory latent \mathbf{z}_τ by integrating the learned vector field $v_\theta(\mathbf{z}(u), u, \mathbf{c})$ over $u \in [0, 1]$. The resulting latent \mathbf{z}_τ is then decoded by the autoencoder to obtain the parameterized trajectory, which is used for the ring-tossing task.

This design implicitly captures the dynamics gap between planned and executed motions. Notably, the autoencoder is trained on a large set of planned trajectories to learn a high-quality motion manifold, while the flow-matching model is learned in the latent space and thus requires only a modest amount of real-world data to capture trajectory structure and dynamics information.

IV. EXPERIMENTS

A. Setup

1) *Simulation and Real-World Setup*: We evaluate our approach in both simulation and real-world experiments using a 6-DoF Galaxea A1 fixed-base robot arm, with simulation implemented in PyBullet [39]. The arm is mounted on a fixed table of height 0.75 m, providing sufficient flight time for the ring. Its joint acceleration limits are set to $[12.5, 12.5, 12.5, 15.0, 15.0, 15.0]$ m/s² to balance safety and motion capability. The ring has a radius of $R_{\text{ring}} = 0.075$ m, and the target is a vertical cylinder of radius $R_{\text{cyl}} = 0.005$ m and height $z_{\text{cyl}} = 0.1$ m placed on the ground. In the real setup (Fig. 5), a RealSense D435i depth camera is used to detect the x, y coordinates of the target and the ring position at $z = 0.1$ m. For the ring, contours are extracted by OpenCV Canny edge detection and fitted with ellipses for robust localization [40].




Fig. 6. Comparison of landing distributions between simulation and the real world using motion planning (Baseline). Ten trajectories are planned targeting $r = 1.8$ m, and the actual landing positions are plotted. Note that the significant deviation and variance in the real world (orange dots) illustrate the dynamics gap that our proposed DA-MMP aims to close. This corresponds to the low success rate of the baseline reported in Table I.

2) *Data Collection Details*: We collect a dataset of 90k feasible trajectories for training the autoencoder, with target radial distances uniformly sampled within $[r_{\text{plan}}^{\min}, r_{\text{plan}}^{\max}] = [1.0, 2.5]$ m. For goal-manifold sampling, the ring angular velocity ω_z is drawn from $[\omega_{\min}, \omega_{\max}] = [1.5\pi, 3\pi]$ rad/s. We set $N_{\text{planning}} = 80$, $N_{\text{smoothing}} = 100$, $\Delta t_{\text{col}} = 1/30$ s, and $f_{\text{ctrl}} = 240$ Hz. For training the flow-matching model, we collect 60 executed trajectories by uniformly setting target distances from $[r_{\text{exec}}^{\min}, r_{\text{exec}}^{\max}] = [1.5, 2.0]$ m in both simulation and real-world settings. Although the dataset is small, we will demonstrate in the experiments that the model effectively captures dynamics information and exhibits generalization capability. The difference between $[r_{\text{plan}}^{\min}, r_{\text{plan}}^{\max}]$ and $[r_{\text{exec}}^{\min}, r_{\text{exec}}^{\max}]$ reflects both safety considerations in real execution and the need to reserve part of the range for evaluating generalization.

3) *Policy Learning Details*: We use $K = 30$ basis functions, which provides a good balance between representation accuracy and model complexity. The autoencoder maps each trajectory parameterization \mathbf{p}_τ into a d_z -dimensional latent space, with $d_z = 64$ in our implementation. Both the encoder and decoder are 3-layer MLPs with hidden dimensions [256, 512, 256], and Leaky ReLU activations. The model is trained for up to 30,000 epochs with early stopping, using batch size 256 and Adam optimizer (learning rate 1×10^{-4} , weight decay 1×10^{-5}). The learned latent space then serves as input for the conditional flow-matching model. For this model, the network is a 6-layer MLP with hidden dimensions [256, 512, 1024, 1024, 512, 256] and Swish activations. Training uses batch size 450 with Adam optimizer (learning rate 3×10^{-4} , weight decay 1×10^{-6}) for up to 20,000 epochs. During trajectory generation, we integrate the learned vector field with the midpoint method using a step size of 0.001. All neural network inputs are normalized by subtracting the mean and dividing by the standard deviation.

B. Dynamics-Aware Trajectory Generation

To evaluate the effectiveness of our method in generating accurate throwing trajectories under dynamics gap, we compare it against multiple baselines in both simulation and real-world experiments, where in simulation we additionally




Fig. 7. **Effect of the autoencoder on trajectory generation.** Shown are positions (q_1, q_3) and velocities (\dot{q}_1, \dot{q}_3) of two representative joints. Our method (orange) produces smoother and more coherent trajectories, whereas the no-AE variant (blue) exhibits irregular and oscillatory motions.

introduce air drag by setting linear and angular damping. We consider three baselines in simulation: motion planning with a single attempt, motion planning with two attempts, and a residual-style correction method that replans the second throw by compensating for the error observed in the first attempt, which is applied only when the first throw is unsuccessful. In the real world, we further include two human baselines: novice participants throwing without any prior practice, and expert participants who have practiced for 60 trials before evaluation. We report success rate (SR) as the evaluation metric, defined as the fraction of trials in which the ring successfully lands around the target peg. All results are averaged over three random seeds, each tested on 10 randomly sampled target positions.

The success rates of all methods are summarized in Table I. Residual-style correction performs best in simulation but fails in the real world. Our method achieves the highest success rate in reality. Interestingly, it slightly outperforms trained human experts given the same number of practice trials, highlighting the efficiency of our framework which learns dynamics on top of structured motion primitives.

To analyze the result reported in Table I, we investigate three key questions:

- 1) Why does residual-style correction excel in simulation but fail in the real world?
- 2) Why does our method achieve a high success rate in both the simulation and the real world?
- 3) Can our method generalize to target distances not covered in the training data?

To answer these questions, we visualize the landing distributions of motion planned throws in both simulation and the real world, as shown in Fig. 6. In this evaluation, we plan ten trajectories targeting $r = 1.8$ m and plot the actual landing positions of the executed throws. In simulation, the landing points exhibit a consistent bias with minimal variance, as the discrepancy is dominated by the added air drag rather than control or contact uncertainties. In contrast, the real-world distribution exhibits both bias and variance due to additional control and contact dynamics gaps as well as visual perception errors. This explains why residual-style correction is highly effective in simulation, as

TABLE II

EFFECT OF DATASET SCALE ON AUTOENCODER RECONSTRUCTION. RMSE IS MEASURED IN THE NORMALIZED TRAJECTORY PARAMETER SPACE, AND LRE IS DEFINED IN (10).

Dataset size	Parameter-Space RMSE	LRE (%)
0.09k	0.201	12.4
0.9k	0.007	1.9
9k	0.007	1.1
90k	0.001	0.9

a systematic bias can be compensated for by adjusting the target. However, it fails in the real world because stochastic variance cannot be corrected in the same way (Question 1). In fact, its performance (6.7%) is even less than the marginal improvement gained by simply replanning a second time (23.3% - 13.3% = 10.0%), as the compensation may even exacerbate the error when the variance is high. Our method, however, encodes trajectories rather than target positions, enabling it to capture trajectory-level dynamics and remain effective under real-world variability (Question 2). Finally, as for Question 3, our method demonstrates promising generalization capability. When attempting throws to a distance of 1.2 m, for which no real-world trajectories are collected—the policy generates feasible motions that successfully hit the target (as shown in the supplementary video). This indicates that the model learns the underlying trajectory-dynamics mapping, rather than merely memorizing training samples.

C. Ablations

1) *Effect of the Autoencoder*: To evaluate the role of the autoencoder in facilitating efficient training of the downstream generative model, we compare our approach against training the flow-matching model directly on raw trajectory parameterizations. This variant lacks the manifold regularization provided by the embedding stage and therefore must model high-dimensional inputs directly. On this test, we generate trajectories in an unconditional manner. As shown in Fig. 7, the non-embedded variant yields irregular joint motion profiles that are rarely executable, whereas our method generates smooth and feasible trajectories. This demonstrates that the autoencoder is essential for capturing the intrinsic structure of feasible trajectories and for mapping trajectories with their underlying dynamics.

2) *Effect of Dataset Scale*: We further study the effect of dataset scale on the quality of the learned motion manifold. The autoencoder is trained with four different dataset sizes: 0.09k, 0.9k, 9k, and the full 90k planned trajectories, with the network capacity reduced proportionally for smaller datasets to mitigate overfitting. For evaluation, we sample 100 trajectories with a different random seed as a held-out test set. We report two metrics: the normalized parameter-space reconstruction error (RMSE), and the relative length reconstruction error

$$\text{LRE} = \frac{1}{N} \sum_{j=1}^N \frac{|L^{(j)} - \hat{L}^{(j)}|}{L^{(j)}} \times 100\%. \quad (10)$$

TABLE III

EFFECT OF PARAMETERIZATION AND DATASET SCALE ON GEOMETRIC SMOOTHNESS OVER 100 HELD-OUT TRAJECTORIES.

Parameterization	MSSD vs. Dataset Size		
	0.9k	9k	90k
Waypoints	596.4	558.9	555.9
DA-MMP (Ours)	280.7	282.4	296.3

where $L^{(j)}$ and $\hat{L}^{(j)}$ denote the ground-truth and the reconstructed trajectory lengths. As summarized in Table II, larger training sets consistently produce lower RMSE and LRE, demonstrating that a sufficiently large corpus of planned trajectories is necessary to capture the diversity of feasible motions and establish a high-quality manifold that generalizes across targets.

3) *Effect of Radial Basis Functions*: To evaluate the utility of radial basis functions, we compare our parameterization against a linear interpolation baseline using $N_{\text{wp}} = 32$ uniform waypoints and length L , matching our method’s parameter count. We train autoencoders for both representations on three dataset scales (0.9k, 9k and 90k) and evaluate the reconstructed geometric smoothness on 100 held-out trajectories.

To assess geometric smoothness independent of execution duration, we sample $N_{\text{eval}} = 100$ uniform points along the normalized phase $s \in [0, 1]$ and approximate the second derivative at each interior point s_i :

$$\mathbf{q}''(s_i) = \frac{\mathbf{q}(s_{i+1}) - 2\mathbf{q}(s_i) + \mathbf{q}(s_{i-1}))}{\Delta s^2} \quad (11)$$

where $s_i = \frac{i-1}{N_{\text{eval}}-1}$ and $\Delta s = \frac{1}{N_{\text{eval}}-1}$. The Mean Squared Second Derivative (MSSD) is then computed as:

$$\text{MSSD} = \frac{1}{D(N_{\text{eval}} - 2)} \sum_{i=2}^{N_{\text{eval}}-1} \|\mathbf{q}''(s_i)\|_2^2 \quad (12)$$

where $D = 6$ is the robot’s degrees of freedom. A lower MSSD indicates a smoother curve.

As shown in Table III, two key conclusions emerge. First, our formulation yields significantly lower MSSD across all dataset scales, demonstrating that its inherent C^2 continuity acts as a strong structural prior to guarantee geometric smoothness. Second, for the waypoint-based baseline, trajectory smoothness noticeably improves as the dataset size increases (MSSD drops from 596.4 to 555.9). This indicates that training on a larger dataset helps the autoencoder learn a better latent representation, thereby mitigating high-frequency spatial reconstruction noise. Nevertheless, even with the largest dataset, the waypoint approach falls short of the smoothness naturally enforced by our radial basis parameterization, making it inherently suited for high-speed dynamic tasks.

V. CONCLUSION

In this work, we present DA-MMP, a framework that builds upon motion manifold primitives for complex goal-conditioned dynamic manipulation tasks. Our approach extends motion manifold primitives to variable-length trajectories through a compact parameterization, and leverages a

large-scale dataset of planned trajectories to learn a high-quality manifold. On top of this manifold, we introduce a conditional flow matching model in the latent space that captures the relationship between trajectories and their execution dynamics, while relying on only limited real-world trials. Our experimental evaluation confirms that DA-MMP consistently produces smooth and coordinated throwing motions with high real-world success rates. Future directions include extending the framework to generalize across objects with diverse shapes and physical properties, exploiting its expressive motion generation to control object orientation and spin at release, and applying it to a broader range of goal-conditioned dynamic manipulation tasks.

ACKNOWLEDGMENT

The authors would like to thank Zhecheng Yuan, Kun Lei, Qianwei Han, Xueyi Liu, and Lingxiao Guo for hardware assistance, Kaizhe Hu, Yu Qi for idea discussions, and Prof. Kris Hauser and Dr. Tobias Kunz for helpful discussions on kinodynamic planning algorithms. This work is supported by Tsinghua University Dushi Program.

REFERENCES

- [1] M. T. Mason and K. M. Lynch, "Dynamic manipulation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 1993, pp. 152–159.
- [2] K. M. Lynch and M. T. Mason, "Dynamic nonprehensile manipulation: Controllability, planning, and experiments," *Int. J. Robot. Res.*, vol. 18, no. 1, pp. 64–92, 1999.
- [3] Y. Ma, A. Cramariuc, F. Farshidian, and M. Hutter, "Learning coordinated badminton skills for legged manipulators," *Sci. Robot.*, vol. 10, no. 102, 2025.
- [4] Z. Su, B. Zhang, N. Rahmadian, Y. Gao, Q. Liao, C. Regan, K. Sreenath, and S. S. Sastry, "Hitter: A humanoid table tennis robot via hierarchical planning and learning," 2025.
- [5] D. B. D'Ambrosio, S. Abeyruwan, L. Graesser, A. Iscen, H. B. Amor, A. Bewley, B. J. Reed, K. Reymann, L. Takayama, Y. Tassa, K. Choromanski, E. Coumans, D. Jain, N. Jaitly, N. Jaques, S. Kataoka, Y. Kuang, N. Lazic, R. Mahjourian, S. Moore, K. Oslund, A. Shankar, V. Sindhwani, V. Vanhoucke, G. Vesom, P. Xu, and P. R. Sanketi, "Achieving human level competitive robot table tennis," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2025, pp. 74–82.
- [6] Y. Zhang, T. Liang, Z. Chen, Y. Ze, and H. Xu, "Catch it! learning to catch in flight with mobile dexterous hands," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2025, pp. 14 385–14 391.
- [7] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. A. Funkhouser, "Tossingbot: Learning to throw arbitrary objects with residual physics," *IEEE Trans. Robot.*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [8] C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Iterative residual policy: For goal-conditioned dynamic manipulation of deformable objects," *Int. J. Robot. Res.*, vol. 43, no. 4, pp. 389–404, 2024.
- [9] S. Aslam, K. Kumar, P. Zhou, H. Yu, M. Y. Wang, and Y. She, "Dartbot: Overhand throwing of deformable objects with tactile sensing and reinforcement learning," *IEEE Trans. Autom. Sci. Eng.*, vol. 22, pp. 13 644–13 661, 2025.
- [10] Z. Wang and A. H. Qureshi, "Implicit physics-aware policy for dynamic manipulation of rigid objects via soft body tools," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2025, pp. 13 174–13 180.
- [11] 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.
- [12] J. Kober and J. Peters, "Learning motor primitives for robotics," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2009, pp. 2112–2118.
- [13] J. Kober, K. Mülling, O. Kroemer, C. H. Lampert, B. Schölkopf, and J. Peters, "Movement templates for learning of hitting and batting," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2010, pp. 853–858.
- [14] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Probabilistic movement primitives," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2013, pp. 2616–2624.
- [15] Y. Zhou, J. Gao, and T. Asfour, "Learning via-point movement primitives with inter- and extrapolation capabilities," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2019, pp. 4301–4308.
- [16] A. Pervez, Y. Mao, and D. Lee, "Learning deep movement primitives using convolutional neural networks," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)*, 2017, pp. 191–197.
- [17] S. Bahl, M. Mukadam, A. Gupta, and D. Pathak, "Neural dynamic policies for end-to-end sensorimotor learning," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020.
- [18] G. Li, Z. Jin, M. Volpp, F. Otto, R. Lioutikov, and G. Neumann, "Prodmpp: A unified perspective on dynamic and probabilistic movement primitives," *IEEE Robot. Autom. Lett.*, vol. 8, no. 4, pp. 2325–2332, 2023.
- [19] M. Saveriano, F. J. Abu-Dakka, A. Kramberger, and L. Peternel, "Dynamic movement primitives in robotics: A tutorial survey," *Int. J. Robot. Res.*, vol. 42, no. 13, pp. 1133–1184, 2023.
- [20] M. Noseworthy, R. Paul, S. Roy, D. Park, and N. Roy, "Task-conditioned variational autoencoders for learning movement primitives," in *Proc. Conf. Robot Learning (CoRL)*, 2019, pp. 933–944.
- [21] B. Lee, Y. Lee, S. Kim, M. Son, and F. C. Park, "Equivariant motion manifold primitives," in *Proc. Conf. Robot Learning (CoRL)*, 2023, pp. 1199–1221.
- [22] Y. Lee, "MMP++: motion manifold primitives with parametric curve models," *IEEE Trans. Robot.*, vol. 40, pp. 3950–3963, 2024.
- [23] Y. Lee, B. Lee, S. Kim, and F. C. Park, "Motion manifold flow primitives for task-conditioned trajectory generation under complex task-motion dependencies," *IEEE Robot. Autom. Lett.*, vol. 10, no. 7, pp. 7412–7419, 2025.
- [24] Y. Lee, "Differentiable motion manifold primitives for reactive motion generation under kinodynamic constraints," 2025.
- [25] Y. Ma, Y. Liu, K. Qu, and M. Hutter, "Learning accurate whole-body throwing with high-frequency residual policy and pullback tube acceleration," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2025, pp. 1771–1778.
- [26] Y. Liu and A. Billard, "Tube acceleration: Robust dexterous throwing against release uncertainty," *IEEE Trans. Robot.*, vol. 40, pp. 2831–2849, 2024.
- [27] K. Mülling, J. Kober, and J. Peters, "Learning table tennis with a mixture of motor primitives," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)*, 2010, pp. 411–416.
- [28] K. Mülling, J. Kober, O. Kroemer, and J. Peters, "Learning to select and generalize striking movements in robot table tennis," *Int. J. Robot. Res.*, vol. 32, no. 3, pp. 263–279, 2013.
- [29] Y. Zhang, J. Luo, and K. K. Hauser, "Sampling-based motion planning with dynamic intermediate state objectives: Application to throwing," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2012, pp. 2551–2556.
- [30] A. Pekarovskiy and M. Buss, "Optimal control goal manifolds for planar nonprehensile throwing," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2013, pp. 4518–4524.
- [31] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006, available at <http://planning.cs.uiuc.edu/>.
- [32] T. Kunz and M. Stilman, "Probabilistically complete kinodynamic planning for robot manipulators with acceleration limits," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2014, pp. 3713–3719.
- [33] K. K. Hauser and V. Ng-Thow-Hing, "Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2010, pp. 2493–2498.
- [34] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," 2023.
- [35] X. Liu, C. Gong, and Q. Liu, "Flow straight and fast: Learning to generate and transfer data with rectified flow," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023.
- [36] M. S. Albergo, M. Goldstein, N. M. Boffi, R. Ranganath, and E. Vanden-Eijnden, "Stochastic interpolants with data-dependent couplings," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2024, pp. 921–937.
- [37] Y. Lipman, M. Havasi, P. Holderrieth, N. Shaul, M. Le, B. Karrer, R. Chen, D. Lopez-Paz, H. Ben-Hamu, and I. Gat, "Flow matching guide and code," 2024.
- [38] J. Ho and T. Salimans, "Classifier-free diffusion guidance," 2022.
- [39] E. Coumans and Y. Bai, "PyBullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [40] G. Bradski, "The opencv library," *Dr. Dobbs Journal of Software Tools*, 2000.