

PegasusFlow: Parallel Rolling-Denoising Score Sampling for Robot Diffusion Planner Flow Matching

Lei Ye, Haibo Gao, Peng Xu*, Zhelin Zhang, Junqi Shan, Ao Zhang, Wei Zhang, Ruyi Zhou, Zongquan Deng and Liang Ding

Abstract—Diffusion models offer powerful generative capabilities for robot trajectory planning, yet their practical deployment on robots is hindered by a critical bottleneck: reliance on imitation learning from expert demonstrations. This paradigm is problematic as it is often impractical to produce high quality data for specialized robots, and it creates an inefficient, theoretically suboptimal training pipeline. To overcome this, we introduce PegasusFlow, a parallel rolling-denoising framework that enables direct sampling of trajectory score gradients from environmental interaction, completely bypassing the need for expert data. Our core innovation is a sampling algorithm called Weighted Basis Function Optimization (WBFO), which leverages spline basis representations to achieve superior sample efficiency and faster convergence compared to traditional methods like MPPI. The framework is embedded within a scalable, asynchronous parallel simulation architecture that supports massively parallel rollouts for efficient data collection. Extensive experiments on trajectory optimization and robotic navigation tasks demonstrate that our approach, particularly Action-Value WBFO (AVWBFO) combined with a reinforcement learning warm-start, significantly outperforms baselines. In a challenging barrier-crossing task, our method achieved a 100% success rate and was 18% faster than the next-best method, validating its effectiveness for complex terrain locomotion planning. <https://masteryip.github.io/pegasusflow.github.io/>

I. INTRODUCTION

Diffusion models have emerged as a powerful tool for trajectory generation in robotics, demonstrating a remarkable capability for refining motion plans in high-dimensional task spaces and handling complex constraints [1]–[6]. As for legged robot locomotion, this translates into a potential to solve intricate navigation tasks that challenge traditional reinforcement learning (RL) policies, such as maneuvering in confined spaces with obstacles [7]. However, the practical application of these diffusion planners is severely hampered by their reliance on extensive, high-quality expert demonstration data.

The predominant training paradigm for diffusion planners is imitation learning, where policies are trained via behavior cloning on datasets of expert trajectories, often collected from motion capture or generated by a pre-trained RL policy. This data-driven approach is problematic. First,

The authors are with the State Key Laboratory of Robotics and Systems, Harbin Institute of Technology, Harbin 150001, China.

*Corresponding author: Peng Xu (pengxu.hit@163.com).

This work was supported in part by the National Natural Science Foundation of China under Grant 52595712, Grant 52595710, Grant 52425502; the National Key R&D Program of China (Grant No. 2024YFB4710900); the New Cornerstone Science Foundation through the XPLOER PRIZE.




Fig. 1: Illustration of PegasusFlow. (a) Visualization of the rollouts of an environments within a predict horizon. (b, c) The algorithm can optimize the robot control trajectory to navigate complex environments with barriers and gaps.

collecting high-quality expert demonstrations is often effort-intensive, costly, and sometimes infeasible, particularly for specialized robots like hexapods where such data is scarce. Second, this indirect training process, generating data with one policy (e.g., RL) to train another (the diffusion model), is theoretically suboptimal, computationally wasteful, and can introduce distribution shifts that degrade performance. This naturally motivates our idea: instead of learning from data, we need methods to estimate trajectory score gradients directly from environmental interactions, enabling more direct and efficient score-matching.

Existing methods that touch upon direct score estimation, such as those inspired by Model Predictive Path Integral (MPPI) control [8], [9], are a step in the right direction but fall short in efficiency. They often require thousands of simulation rollouts to produce a single high-quality trajectory, rendering them unsuitable for parallel data generation. Furthermore, to the best of our knowledge, there is currently no known framework capable of performing this score sampling process in a massively parallel manner, which is essential for efficient data collection and diffusion planner training.

To address these gaps, we introduce a framework (Fig. 2) for direct, parallel, and efficient sampling of trajectory score gradients. Our approach formulates trajectory planning as a

discrete-time optimal control problem and leverages a new sampling algorithm, Weighted Basis Function Optimization (WBFO), which provides superior sample efficiency and faster convergence than existing methods. To enable parallel data collection, we embed this optimizer within a scalable, asynchronous parallel simulation architecture built on IsaacGym. This system supports massively parallel rollouts, warm-starting with RL trajectories, and leverages NVIDIA Warp for accelerated Signed Distance Field (SDF) and ray-casting queries.

The primary contributions of this work are:

- 1) **Parallel Score Sampling Framework:** We present a framework to directly sample robot trajectory score gradients in parallel, which makes pure score-matching training without expert demonstrations and bridging diffusion generative models with real-time robotic control more promising. This includes a scalable mini-batch environment simulation system and a rolling-denoising optimization process using IsaacGym with coordinated main and rollout-environment architecture.
- 2) **Weighted Basis Function Optimization (WBFO):** We develop an efficient trajectory optimization approach through spline basis functions with node-level importance weighting, achieving computational efficiency while maintaining trajectory quality through local optimality and temporal coherence.
- 3) **Efficient Noise Sampling Schema:** We introduce an MPC-inspired structured noise scheduling approach combining Latin Hypercube Sampling (LHS), Hierarchical Ramp Noise Scheduling (HRNS) and RL Policy Warm Start for improved space exploration and convergence in trajectory optimization tasks.

II. RELATED WORKS

A. Trajectory Optimization and Model Predictive Control

Model predictive control (MPC) algorithms search for functional gradients and perform gradient descent to solve optimal control problems [10]–[13], they implement receding horizon control with rolling updates for real-time adaptation. Direct methods like Direct Collocation and Multiple Shooting transcribe continuous control into discrete nonlinear programs. While Di Carlo et al. [10] and Grandia et al. [12] achieved impressive results, these methods suffer from local minima and may fail in complex environments. Contact planning methods in our previous work [14], [15] address foothold selection for legged robots but operate at discrete contact level rather than continuous trajectory optimization.

As for sampling-based methods, MPPI [16] and its variants [17]–[19] perform trajectory sampling based on optimal control theory, using relationships between free energy and relative entropy. However, MPPI requires extensive computation during rollout, limiting its capabilities to extend to multiple robot environments.

These concepts inspire our rolling-denoising framework, leveraging MPC’s receding horizon principle while addressing computational limitations through hierarchical noise scheduling and parallel batch rollout.

B. Score-based Models for Robot Planning

DDPM [20] establishes diffusion models through gradual noise reversal, while DDIM [21] achieves 10-50x acceleration via non-Markovian paths. SDE formulation [22], [23] provides continuous-time generalization with connections to stochastic calculus.

Diffusion Policy [2] handles multimodal action distributions for manipulation tasks. Vision-Language-Action models like π_0 [24] and $\pi_{0.5}$ [25] demonstrate broad generalization. These approaches excel in structured environments with moderate computational requirements.

For legged locomotion, DiffuseLoco [3], [4] enables real-time control from offline datasets but relies on expert trajectories. Model-Based Diffusion [9] and Dial-MPC [8] eliminate expert data dependency but lack batch processing capabilities. Several works integrate diffusion motion generator and RL tracking policies for downstream tasks [5]–[7], [26].

Diffusion policies are relatively less used in legged locomotion over challenging terrain due to: (1) real-time control requirements conflicting with low update frequency, (2) computational cost of iterative denoising, and (3) RL’s effectiveness for basic locomotion. However, diffusion offers superior potential for complex environment navigation and obstacle avoidance, motivating our direct score gradient sampling approach.

III. PRELIMINARY

A. Problem Formulation

We formulate the robot planning problem as a discrete-time optimal control problem, depicted in Fig. 2(g). The system must achieve desired objectives while satisfying constraints and maintaining operational requirements.

The discrete-time system dynamics are given by:

$$x_{k+1} = f(x_k, u_k, t_k) \quad (1)$$

where $x_k \in \mathbb{R}^n$ is the state vector at time step k containing the robot’s position, orientation, velocities, and joint configurations, $u_k \in \mathbb{R}^m$ is the control input vector, and $f(\cdot)$ represents the nonlinear system dynamics function.

The finite-horizon optimal control problem is formulated as:

$$\begin{aligned} \min_{u_{i:i+h}} J(u_{i:i+h}) &= e^{-\alpha(h+1)\Delta t} l_f(x_{i+h+1}) \\ &+ \sum_{k=i}^{i+h} e^{-\alpha\Delta t(k-i)} \cdot l(x_k, u_k, t_k; c_k, o_k) \Delta t \quad (2) \\ \text{s.t. } x_{k+1} &= f(x_k, u_k, t_k) \end{aligned}$$

where h is the planning horizon, $l_f(\cdot)$ is the terminal cost function, $\alpha > 0$ is the exponential discount factor that prioritizes immediate costs while maintaining long-term planning capability, c_k represents command inputs at time k , and o_k represents environmental observations.

The stage cost function $l(x_k, u_k, t_k; c_k, o_k)$ includes multiple components:

$$\begin{aligned} l(x_k, u_k, t_k; c_k, o_k) &= l_{\text{task}}(x_k, t_k; c_k) \\ &+ l_{\text{obstacle}}(x_k; o_k) + l_{\text{control}}(u_k) \quad (3) \end{aligned}$$

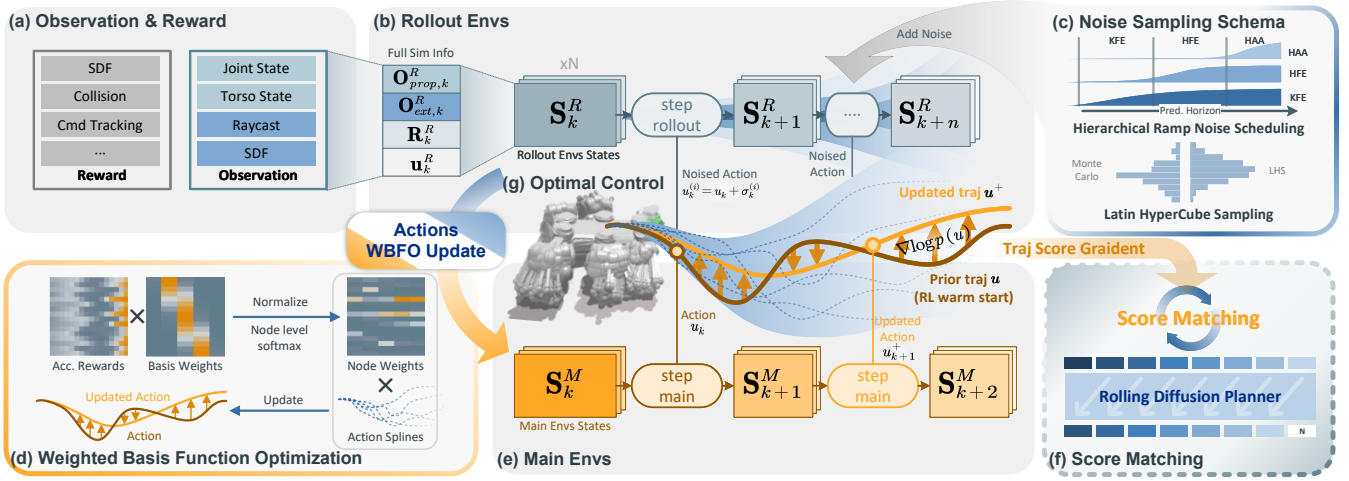


Fig. 2: Framework of PegasusFlow. (a) RL style observations and rewards setup. (b) Rollout environments for trajectory score gradient sampling. (c) Noise sampling schema (Section IV-D). (d) WBFO for trajectory optimization (Section IV-C). (e) Main environments that interact with simulation. (f) Sampled trajectory score gradient can be used for flow matching training of diffusion policy. (g) Optimal control problem and sampling-based score gradient estimation.

where l_{task} represents task-related costs such as trajectory tracking error with respect to commands c_k , l_{obstacle} penalizes constraint violations based on observations o_k , and l_{control} penalizes excessive control effort to ensure smooth operation.

B. Estimating Score Gradient using MPPI

For high-dimensional systems with complex dynamics, directly solving the optimal control problem is computationally intractable. Instead, MPPI can not only approximate the optimal control through stochastic sampling, but also estimate trajectory score gradient from sampled trajectories.

MPPI works by adding noise to the control inputs and evaluating the resulting trajectories, as is shown in Fig. 2(g). Specifically, we sample N noisy control sequences:

$$u_k^{(i)} = u_k + \sigma_k^{(i)} \quad (4)$$

where $\sigma_k^{(i)}$ represents zero-mean Gaussian noise added to the nominal control u_k for the i -th sample trajectory.

The MPPI approximation of the optimal control policy is then given by the importance-weighted average:

$$u_{i:i+h}^+ = \frac{\sum_{j=1}^N u_{i:i+h}^{(j)} e^{-\frac{1}{\lambda} J^{(j)}}}{\sum_{j=1}^N e^{-\frac{1}{\lambda} J^{(j)}}} \quad (5)$$

where $\lambda > 0$ is a temperature parameter that controls the exploration-exploitation trade-off. A smaller λ leads to more aggressive selection of low-cost trajectories, while a larger λ provides more exploration.

An important theoretical insight from [8] is that the MPPI update is mathematically equivalent to a single-stage diffusion process. The MPPI update can be viewed as a one-step gradient ascent using the score function $\nabla \log p(u)$:

$$u^+ = u + \Sigma \nabla \log p_1(u) \quad (6)$$

where $p_0(u) \propto \exp(-\frac{J(u)}{\lambda})$ represents the probability distribution induced by the cost function, and $p_1(\cdot) \propto (p_0 * \phi)(\cdot)$

is the smoothed distribution with kernel ϕ of the introduced gaussian noise σ_k .

IV. SAMPLING-BASED OPTIMIZATION METHOD

The connection between MPPI and diffusion processes provides the theoretical foundation for our approach (Fig. 2(f)). However, standard MPPI requires sampling a large number of trajectories N to achieve good approximation quality, which can be computationally expensive for batch robot simulation. This limitation motivates our development of more efficient sampling strategies based on the WBFO algorithm described in the following sections (Fig. 2(c,d)).

A. Spline Trajectory Representation

In this paper, we employ cubic spline interpolation for trajectory representation, which provides smooth, continuous trajectories through control nodes while maintaining computational efficiency. Consider a local parameter $t \in [0, 1]$ representing the normalized position within a trajectory segment.

The position at parameter t can be computed using a general cubic spline formulation:

$$\text{Pos}(t) = \mathbf{P} \cdot \mathbf{M} \cdot \mathbf{T} \quad (7)$$

where:

- $\mathbf{P} = [P_{i-1}, P_i, P_{i+1}, P_{i+2}]$ is the matrix of control points influencing the current segment
- \mathbf{M} is the cubic spline basis matrix that defines the interpolation scheme
- $\mathbf{T} = [1, t, t^2, t^3]^T$ is the power basis vector

For our implementation, we specifically choose Catmull-Rom splines due to their interpolation property, which ensures that the trajectory passes exactly through all control nodes. The spline representation allows us to construct a

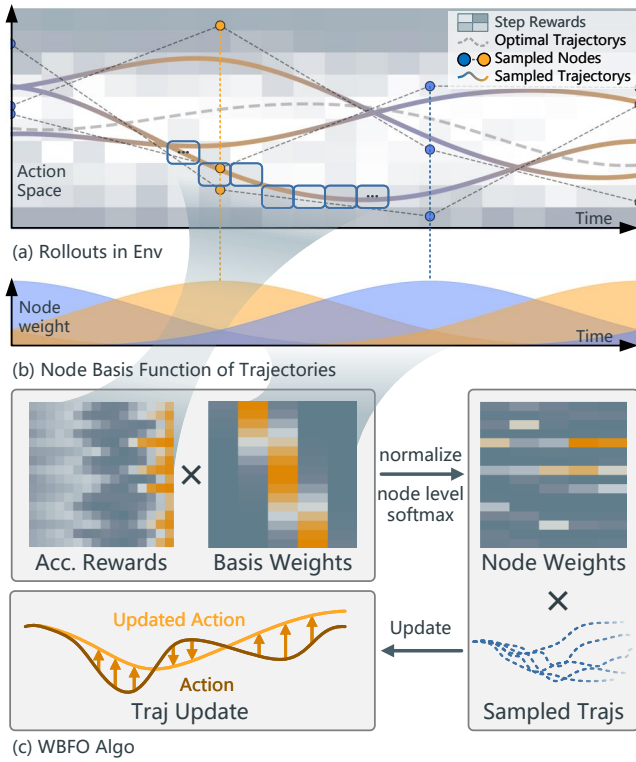


Fig. 3: Weighted Basis Function Optimization. (a) Sampled action trajectories with step-wise rewards. (b) Basis functions of the bundle of the sampled trajectories. (c) Schematic diagram of the WBFO process, illustrating the mapping from trajectory rewards to node weights via basis functions.

basis weight matrix Φ that maps directly from control nodes to trajectory samples. This matrix enables efficient conversion between the two representations. For D dimension trajectories with K control nodes and T discrete time steps, the relationship can be expressed as:

$$\mathbf{u} = \mathbf{P} \cdot \Phi \quad (8)$$

where $\mathbf{u} \in \mathbb{R}^{D \times T}$ represents the dense trajectory samples, $\Phi \in \mathbb{R}^{K \times T}$ is the precomputed basis weight matrix, and $\mathbf{P} \in \mathbb{R}^{D \times K}$ contains the control nodes.

For the reverse operation, we compute the pseudo-inverse of the basis weight matrix Φ^+ to minimize the reconstruction error:

$$\mathbf{P} = \mathbf{u} \cdot \Phi^+ \quad (9)$$

This approach ensures that the conversion between discrete samples and control nodes is mathematically optimal, minimizing the mean squared error between the original dense trajectory and the reconstructed trajectory from the control nodes.

B. Trajectory Sampling through Basis Function

Building on the spline representation, we can sample the functional space of the trajectory more efficiently by utilizing basis functions and step rewards. Instead of sampling complete trajectories directly, we sample the control points

$\{P_k\}$ in the reduced-dimensional space, which results in more efficient exploration because:

- 1) The dimensionality of the control point space (K nodes) is typically much lower than the full trajectory space (T timesteps), where $K \ll T$.
- 2) Each control point influences only a local portion of the trajectory through its associated basis functions, enabling targeted optimization.
- 3) The basis function representation inherently enforces smoothness constraints, eliminating physically implausible trajectories.

Mathematically, spline trajectory $u(t)$ can be represented as a linear combination of basis functions:

$$u(t) = \sum_{k=1}^K P_k \phi_k(t) \quad (10)$$

where $\{P_k\}_{k=1}^K$ are the control point coefficients and $\{\phi_k(t)\}_{k=1}^K$ are the basis functions. The optimization problem becomes finding optimal control points:

$$P_k^* = \arg \min_{P_k} \mathbb{E}_{u \sim p(u|\{P_k\})} \left[\sum_{t=1}^T r_t \right] \quad (11)$$

where r_t is the reward at time t and the expectation is taken over trajectories sampled based on the control points $\{P_k\}$ instead of sampling complete trajectories directly like MPPI. This functional space sampling exploits the low-dimensional structure of the trajectory manifold for efficient optimization.

C. Weighted Basis Function Optimization

We propose Weighted Basis Function Optimization (WBFO), which extends conventional MPPI by decomposing trajectories into control points and applying spatiotemporal weighting through basis function representations. This approach enables targeted optimization of trajectory segments while maintaining global coherence.

As illustrated in Fig. 3, our unified WBFO framework (Algorithm 1) handles both trajectory optimization ($\gamma = 0$) and optimal control ($\gamma > 0$) through a discount factor parameter. The algorithm computes discounted accumulated rewards from step-wise rewards collected from rollout environments in Fig. 3(a):

$$\mathbf{R}_{\text{acc}}(i, t) = \sum_{s=t}^T r_{i,s} \cdot \gamma^{s-t} \quad (12)$$

where $r_{i,s}$ is the reward of trajectory i at timestep s . When $\gamma = 0$, this reduces to step-wise rewards $\mathbf{R}_{\text{acc}}(i, t) = r_{i,t}$ for trajectory optimization. When $\gamma > 0$, it provides temporal credit assignment for systems with integrators.

The core innovation lies in computing node-specific weights through basis function mapping using basis function in Fig. 3(b):

$$\mathbf{W} = \mathbf{R}_{\text{acc}} \Phi^T \quad (13)$$

where $\Phi \in \mathbb{R}^{K \times T}$ maps timesteps to control nodes, and $\mathbf{R}_{\text{acc}} \in \mathbb{R}^{N \times T}$ contains trajectory rewards. Node-level normalization ensures balanced updates:

$$\mathbf{W}(i, j) = \frac{\mathbf{W}(i, j) - \text{mean}(\mathbf{W}(i, :))}{\text{std}(\mathbf{W}(i, :))} \quad (14)$$

Algorithm 1 Unified Weighted Basis Function Optimization

Input: nodes $\mathbf{P} \in \mathbb{R}^{D \times K}$ of prior \mathbf{u} , noise schedule $\sigma \in \mathbb{R}^{D \times K}$, discount factor γ , number of samples N

Output: Updated control points $\mathbf{P}^+ \in \mathbb{R}^{D \times K}$

- 1: Sample N trajectories $\mathbf{U} \in \mathbb{R}^{N \times T \times D}$ from prior \mathbf{u} by adding noise σ to \mathbf{P}
 - 2: Evaluate step-wise rewards $\mathbf{R} \in \mathbb{R}^{N \times T}$ for all trajectories by interacting using rollout environments in Fig. 2(b)
 - 3: Calculate basis function mask $\Phi \in \mathbb{R}^{K \times T}$
 - 4: **if** $\gamma > 0$ **then** ▷ AVWBFO-Integrated Systems
 - 5: Compute accumulated rewards \mathbf{R}_{acc} ▷ Eq. (12)
 - 6: **else** ▷ WBFO-Traj Optimization
 - 7: $\mathbf{R}_{\text{acc}} \leftarrow \mathbf{R}$ ▷ Use step-wise rewards directly
 - 8: **end if**
 - 9: Calculate node weight matrix $\mathbf{W} = \mathbf{R}_{\text{acc}} \Phi^T$ ▷ Eq. (13)
 - 10: Apply node-level normalization to \mathbf{W} ▷ Eq. (14)
 - 11: Compute node-level softmax weights: $w_{ij} = \frac{\exp(\mathbf{W}(i, j))}{\sum_k \exp(\mathbf{W}(i, k))}$
 - 12: Update control points: $P_j^+ = \sum_{i=1}^N w_{ij} \cdot P_{ij}$
 - 13: **return** $\mathbf{P}^+ = [P_1^+, P_2^+, \dots, P_K^+]$
-

This unified approach offers several advantages over conventional MPPI:

- **Local Optimality:** By applying node-level focus, we can concentrate optimization efforts on challenging segments of the trajectory.
- **Temporal Coherence:** The basis function mask enforces smoothness and temporal consistency across the optimized trajectory.
- **Efficient Exploration:** The weighted update scheme enables more effective exploration of the state space in regions of high uncertainty.
- **Unified Framework:** A single algorithm handles both trajectory optimization and optimal control problems through the discount factor parameter.

The WBFO approach forms a critical component of our rolling diffusion planner, enabling efficient trajectory refinement while maintaining the theoretical foundations of stochastic optimal control.

D. Noise Sampling Schema

We propose a dual-component approach combining Latin Hypercube Sampling (LHS) with Hierarchical Ramp Noise Scheduling (HRNS) to achieve superior sample efficiency and space exploration, shown in Fig. 2(c).

Latin Hypercube Sampling. Traditional Monte Carlo sampling can lead to clustering and poor coverage in high-dimensional trajectory optimization problems. LHS addresses this by ensuring uniform distribution across all noise

dimensions through stratified sampling, where each dimension is divided into equally probable intervals with exactly one sample per interval. This guarantees better parameter space coverage with fewer samples, translating to more diverse exploration of motion primitives.

Hierarchical Ramp Noise Scheduling. HRNS introduces a structured approach to noise level management in receding horizon planning. The scheduling follows a ramp structure where noise magnitudes are modulated according to temporal distance from the current planning step, enabling focused exploration in critical near-term decisions while maintaining broader exploration for long-term planning.

The combination creates a synergistic effect: LHS ensures comprehensive space exploration while HRNS directs this exploration toward critical regions of the planning horizon, resulting in faster convergence and improved solution quality.

V. ASYNCHRONOUS PARALLEL SIMULATION

We design an asynchronous parallel simulation framework for diffusion-based trajectory optimization that enables efficient data collection and score sampling through specialized adaptations for sampling from hundreds of sub-environments, shown in Fig. 2.

A. Batch Rollouts in Parallel

Our parallel rollout system implements a hierarchical environment structure for efficient trajectory sampling, which is illustrated in Fig. 2(b,e). The framework maintains:

$$N_{\text{total}} = N_m \times (1 + N_r) \quad (15)$$

where main environments N_m are positioned at regular intervals among rollout environments N_r : $\{0, 1 + N_r, 2(1 + N_r), \dots\}$.

The system executes two key operations: **state synchronization** resets rollout environments to match main environment states before each rollout; **rollout execution** allows parallel exploration of trajectory branches from identical initial states while main environments remain frozen. After rollout completion, main environments are restored from their cached states to continue execution, ensuring temporal consistency for receding horizon planning.

B. Rolling-Denoising Optimization

Our framework employs a rolling-denoising optimization strategy that integrates trajectory optimization within the receding horizon control loop (Fig. 2(b,e)). At each optimization interval, the system performs n_{denoise} denoising steps (i.e., the number of denoising iterations) on the current trajectory nodes using the WBFO algorithm described in Section IV-C.

The optimization process follows a Model Predictive Control (MPC) paradigm: (1) the current node trajectories are optimized through multiple denoising iterations using parallel rollout environments; (2) the optimized trajectory is executed for one timestep in the main environments; (3) trajectories are shifted forward in time, with new terminal actions appended via RL policy warm start in Section V-C; (4) the cycle repeats with updated environmental observations.




Fig. 4: Robotic tasks in PegasusFlow. (a) Hexapod timber piles navigation. (b) Franka arm collision avoidance planning. (c) Quadruped walking. (d) Hexapod confined space navigation. Videos can be found in the supplementary material or at [this page](#).

C. RL Policy Warm Start

We implement an RL policy warm start mechanism that initializes optimization with high-quality action sequences from pre-trained policies. The warm start process involves: **Trajectory initialization**. The program executes rollout procedures using the loaded policy to generate initial action sequences covering the planning horizon; **trajectory append**. For receding horizon planning, the RL policy generates the last actions for appending during trajectory shifting operations. For convenience, we use RL style observations and rewards setup both for RL warm start and trajectory optimization as shown in Fig. 2(a).

VI. EXPERIMENTS

We validate the proposed hierarchical rolling-denoising framework through three key experimental aspects: (1) validating that PegasusFlow enables effective parallel trajectory score evaluation across diverse robotic tasks, (2) demonstrating that WBFO/AVWBFO (proposed) optimization algorithms combined with Latin Hypercube Sampling (LHS) achieve superior sample efficiency over MPPI baselines, and (3) showing that RL warm-starting significantly enhances trajectory optimization performance. It should be noted that in all experiments, $\gamma = 1.0$ is used for AVWBFO (proposed). These experiments collectively demonstrate how our framework components (advanced optimization algorithms, strategic sampling, and RL integration) synergistically improve trajectory planning efficiency and success rates.

A. Various Tasks in PegasusFlow

We validate PegasusFlow’s versatility across diverse robotic platforms and challenging scenarios, as illustrated in Fig. 4. Our framework demonstrates robust performance in

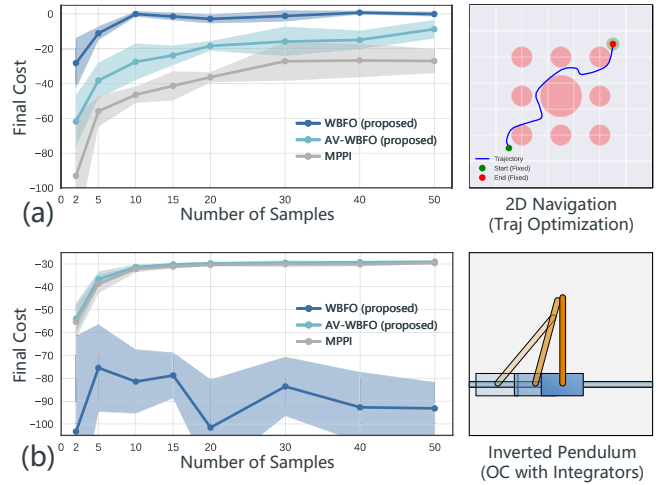


Fig. 5: Optimization performance comparison. (a) 2D Navigation (trajectory planning problem). (b) Inverted Pendulum (optimal control problem). X-axis shows the number of samples used, Y-axis shows the final cost after optimization iterations.

several distinct domains: (a) **Hexapod timber piles navigation** tests complex terrain traversal with irregular obstacles requiring precise foothold planning and dynamic balance, (b) **Franka arm collision avoidance** validates high-dimensional manipulation planning in cluttered environments with safety constraints, (c) **Quadruped walking** shows hundreds of rollout environments for trajectory score sampling, and (d) **Hexapod confined space navigation** evaluates performance in spatially constrained environments where traditional planning methods typically fail.

B. Optimization Algorithm Performance

1) *Algorithm Comparison under Varying Sample Numbers*: We evaluate WBFO/AVWBFO (proposed) performance against MPPI baselines across two representative scenarios: 2D Navigation (point-to-point trajectory optimization with 25 random circular obstacles in a 10×10 meter workspace) and Inverted Pendulum (optimal control for system with integrators). Three optimization approaches are compared: (1) AVWBFO (proposed), (2) WBFO (proposed), and (3) MPPI baseline.

For 2D Navigation, trajectories use 16 nodes interpolated to 64 dense samples over 10 optimization iterations with exponential noise decay (initial noise 3.0, decay rate 0.6). For Inverted Pendulum, trajectories use 32 nodes interpolated to 128 dense samples with identical optimization parameters. Statistical significance is ensured through 5 independent trials.

Results. As shown in Fig. 5, our proposed methods demonstrate different performance characteristics across the two experimental scenarios:

In trajectory optimization tasks (2D Navigation, Fig. 5(a)), WBFO (proposed) demonstrates superior sample efficiency compared to MPPI, particularly with limited samples (≤ 20).

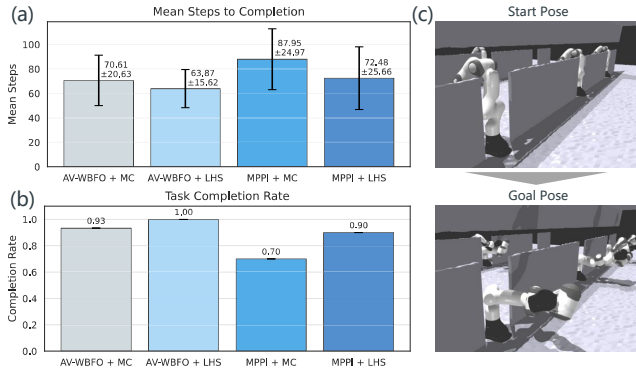


Fig. 6: Franka planning task. 30 manipulators are planned using 4 different method combinations to reach a goal behind a wall gap. (a) Mean steps to completion and standard deviation over 30 trials. (b) Completion rates of each method within 150 steps. (c) Start and final states of the planning task.

For example, with 10 samples, WBFO (proposed) achieves a final cost of -0.095 ± 1.46 , while MPPI achieves -46.53 ± 4.60 , indicating WBFO’s better performance with fewer samples. AVWBFO (proposed) performs intermediately between WBFO (proposed) and MPPI in trajectory optimization.

In optimal control tasks for system with integrators (Inverted Pendulum, Fig. 5(b)), AVWBFO (proposed) slightly outperforms MPPI, while WBFO (proposed) shows inferior performance, confirming the importance of discounted reward for dynamic systems.

2) *Franka Arm Collision Avoidance Planning*: We evaluate completion rates and efficiency of Franka arm collision avoidance planning across four method combinations: AVWBFO (proposed)/MPPI with Monte Carlo (MC) or Latin Hypercube Sampling (LHS). 30 Franka manipulators are tasked with reaching a goal behind a wall gap within 150 steps using 64 samples for all methods, as depicted in Fig. 6.

Results. AVWBFO (proposed) achieves higher completion rates than MPPI (93.3% vs 70% with MC sampling, 100% vs 90% with LHS). LHS sampling consistently improves performance, with AVWBFO (proposed)+LHS achieving optimal results (100% completion rate, 63.9 ± 15.6 steps to completion). This demonstrates that both optimization algorithm choice and sampling strategy contribute to improved planning performance.

C. Robot Navigation Performance in Rough Terrain

We validate the complete hierarchical rolling-denoising framework in realistic legged robot navigation scenarios with barrier crossing challenges. The experimental setup uses 20 Hexapod robots navigating from inside a square barrier (height 0.2m, width 0.35m) to an external goal within 300 steps (6 seconds at 0.02s step time) with 128 samples, testing agile maneuvering in constrained spaces, as shown in Fig. 7(c).

Five trajectory optimization approaches are compared: (1) **Vanilla RL** (pre-trained policy from flat terrain without

trajectory optimization), (2) **AVWBFO w RL (proposed)** (our algorithm warm-started with RL policy), (3) **AVWBFO w/o RL** (our algorithm without RL warm-start), (4) **MPPI w RL** (MPPI with RL warm-start), and (5) **MPPI w/o RL** (MPPI without RL warm-start). Performance is evaluated using navigation success rate, mean completion steps, and accumulated reward.

Results and Analysis. The experimental results, as illustrated in Fig. 7, reveal three critical insights:

- **RL Warm-starting Impact:** Both AVWBFO (proposed) w RL and MPPI w RL achieved perfect 100% success rates (Fig. 7(b)), while their non-warm-started counterparts achieved only 75% and 40% respectively. Vanilla RL completely failed (0% success), highlighting the necessity of trajectory optimization for complex navigation.
- **Algorithm Superiority:** AVWBFO (proposed) consistently outperformed MPPI in both configurations, as shown in Fig. 7(a). With RL warm-starting, AVWBFO (proposed) required 165.1 ± 47.0 steps versus MPPI’s 202.4 ± 37.1 steps. Without warm-starting, AVWBFO (proposed) needed 226.1 ± 51.9 steps versus MPPI’s 237.7 ± 48.2 steps.
- **Solution Quality:** Mean reward analysis (Fig. 7(b)) confirms trajectory optimality, with AVWBFO (proposed) w RL achieving the highest reward (0.024), followed by MPPI w RL (0.016), AVWBFO (proposed) w/o RL (0.014), MPPI w/o RL (0.011), and Vanilla RL (-0.010). Higher rewards indicate smoother, collision-free trajectories.

These results validate our hierarchical approach, demonstrating that the combination of learned behaviors with adaptive trajectory optimization significantly enhances both success rates and motion quality in challenging navigation scenarios.

VII. CONCLUSION

In this paper, we introduced PegasusFlow, a hierarchical rolling-denoising framework that fundamentally addresses the expert data dependency bottleneck in diffusion-based robot planners. By enabling direct and parallel sampling of trajectory score gradients from environmental interaction, our approach reduces reliance on imitation learning, paving the way for training diffusion policies via pure score-matching.

Our core contribution, the WBFO algorithm and its action-value variant (AVWBFO), was shown to be significantly more sample-efficient and effective than traditional sampling-based methods like MPPI. Experiments demonstrated that combining AVWBFO with an RL warm-start yields a powerful planner capable of solving complex locomotion tasks. In a challenging barrier-crossing scenario (Fig. 7), our method achieved a 100% success rate and was 18% faster than the next-best baseline, highlighting its suitability for agile navigation in constrained environments where pure RL policies fail.

By bridging the gap between the generative power of diffusion models and the practical demands of real-time

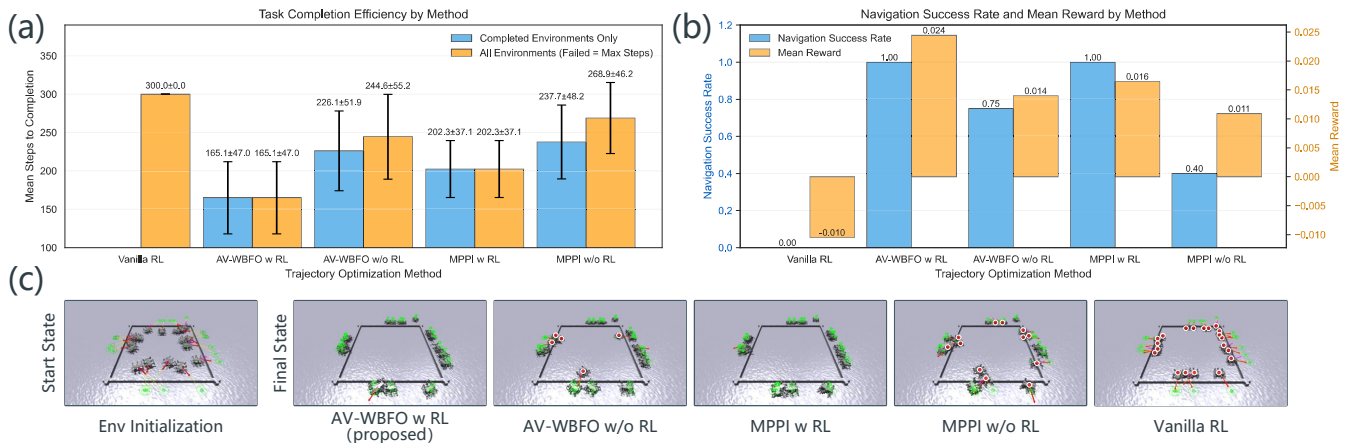


Fig. 7: Hexapod Barrier Navigation. (a) Navigation task completion efficiency comparison. Mean steps to completion and standard deviation are measured over 20 trials. (b) Navigation success rate and mean reward during trajectory optimization. (c) The navigation task setup with 20 Hexapod robots starting inside a square barrier (height 0.2m, width 0.35m) and navigating to an external goal. The snapshots at 300 steps of each method are shown, with successful robots marked in green and failed ones in red.

robotic control, this work opens new avenues for developing highly capable and autonomous legged robots. Future work will focus on deploying this framework for full diffusion policy training and extending it to multi-robot coordination and dynamic environments.

REFERENCES

- [1] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-Or, *et al.*, “Human Motion Diffusion Model,” Oct. 2022.
- [2] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [3] X. Huang, Y. Chi, R. Wang, Z. Li, X. B. Peng, S. Shao, B. Nikolic, and K. Sreenath, “DiffuseLoco: Real-Time Legged Locomotion Control with Diffusion from Offline Datasets,” in *Proceedings of The 8th Conference on Robot Learning*. PMLR, Jan. 2025, pp. 1567–1589.
- [4] Q. Liao, T. E. Truong, X. Huang, G. Tevet, K. Sreenath, *et al.*, “BeyondMimic: From Motion Tracking to Versatile Humanoid Control via Guided Diffusion,” Aug. 2025.
- [5] G. Tevet, S. Raab, S. Cohan, D. Reda, Z. Luo, *et al.*, “CLoSD: Closing the Loop between Simulation and Diffusion for multi-task character control,” Oct. 2024.
- [6] “Robot Motion Diffusion Model: Motion Generation for Robotic Characters,” 2024.
- [7] M. Xu, Y. Shi, K. Yin, and X. B. Peng, “PARC: Physics-based Augmentation with Reinforcement Learning for Character Controllers,” in *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, Aug. 2025, pp. 1–11.
- [8] H. Xue, C. Pan, Z. Yi, G. Qu, and G. Shi, “Full-Order Sampling-Based MPC for Torque-Level Locomotion Control via Diffusion-Style Annealing,” Sept. 2024.
- [9] C. Pan, Z. Yi, G. Shi, and G. Qu, “Model-based diffusion for trajectory optimization,” in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 57 914–57 943.
- [10] J. Di Carlo, P. M. Wensing, B. Katz, G. Bleedt, and S. Kim, “Dynamic Locomotion in the MIT Cheetah 3 Through Convex Model-Predictive Control,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid: IEEE, Oct. 2018, pp. 1–9.
- [11] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, “A Unified MPC Framework for Whole-Body Dynamic Locomotion and Manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4688–4695, July 2021.
- [12] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, “Perceptive Locomotion Through Nonlinear Model-Predictive Control,” *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3402–3421, Oct. 2023.
- [13] S. Xu, L. Zhu, H.-T. Zhang, and C. P. Ho, “Robust Convex Model Predictive Control for Quadruped Locomotion Under Uncertainties,” *IEEE Transactions on Robotics*, pp. 1–18, 2023.
- [14] P. Xu, L. Ding, L. Ye, T. Pang, T. Liu, H. Yang, H. Gao, Z. Deng, and J. Pajarinen, “Contact planning for multilegged robots under constraints through parallel mcts,” *IEEE Transactions on Robotics*, vol. 41, pp. 6102–6122, 2025.
- [15] L. Ye, H. Gao, H. Yang, P. Xu, H. Wang, T. Liu, J. Shan, Z. Deng, and L. Ding, “Kcfr: Kinematic collision-aware foothold reachability criteria for legged locomotion,” *IEEE Transactions on Industrial Electronics*, pp. 1–13, 2025.
- [16] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Aggressive driving with model predictive path integral control,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1433–1440.
- [17] J. Yin, Z. Zhang, E. Theodorou, and P. Tsotras, “Trajectory Distribution Control for Model Predictive Path Integral Control using Covariance Steering,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 1478–1484.
- [18] C. Pezzato, C. Salmi, E. Trevisan, M. Spahn, J. Alonso-Mora, and C. H. Corbato, “Sampling-based Model Predictive Control Leveraging Parallelizable Physics Simulations,” Jan. 2025.
- [19] Z. Yi, C. Pan, G. He, G. Qu, and G. Shi, “CoVO-MPC: Theoretical analysis of sampling-based MPC and optimal covariance design,” in *Proceedings of the 6th Annual Learning for Dynamics & Control Conference*. PMLR, June 2024, pp. 1122–1135.
- [20] J. Ho, A. Jain, and P. Abbeel, “Denosing Diffusion Probabilistic Models,” Dec. 2020.
- [21] J. Song, C. Meng, and S. Ermon, “Denosing Diffusion Implicit Models,” Oct. 2022.
- [22] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-Based Generative Modeling through Stochastic Differential Equations,” Feb. 2021.
- [23] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever, “Consistency Models,” May 2023.
- [24] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, *et al.*, “ $\$ \pi_{0.5}$: A Vision-Language-Action Flow Model for General Robot Control,” Nov. 2024.
- [25] P. Intelligence, K. Black, N. Brown, J. Darpanian, K. Dhabalia, *et al.*, “ $\$ \pi_{\{0.5\}}$: A Vision-Language-Action Model with Open-World Generalization,” Apr. 2025.
- [26] X. Huang, T. Truong, Y. Zhang, F. Yu, J. P. Sleiman, *et al.*, “Diffuse-CLoC: Guided Diffusion for Physics-based Character Look-ahead Control,” Mar. 2025.