

000 001 002 003 004 005 SAD-FLOWER: FLOW MATCHING FOR SAFE, ADMIS- 006 SIBLE, AND DYNAMICALLY CONSISTENT PLANNING 007 008 009

010 **Anonymous authors**
011

012 Paper under double-blind review
013
014
015
016
017
018
019
020
021
022
023
024
025

ABSTRACT

026 Flow matching (FM) has shown promising results in data-driven planning. How-
027 ever, it inherently lacks formal guarantees for ensuring state and action constraints,
028 whose satisfaction is a fundamental and crucial requirement for the safety and ad-
029 missibility of planned trajectories on various systems. Moreover, existing FM
030 planners do not ensure the dynamical consistency, which potentially renders tra-
031 jectories inexecutable. We address these shortcomings by proposing SAD-Flower,
032 a novel framework for generating Safe, Admissible, and **Dynamically** consistent
033 trajectories. Our approach relies on an augmentation of the flow with a virtual
034 control input. Thereby, principled guidance can be derived using techniques from
035 nonlinear control theory, providing formal guarantees for state constraints, action
036 constraints, and dynamic consistency. Crucially, SAD-Flower operates without re-
037 training, enabling test-time satisfaction of unseen constraints. Through extensive
038 experiments across several tasks, we demonstrate that SAD-Flower outperforms
039 various generative-model-based baselines in ensuring constraint satisfaction.
040

1 INTRODUCTION

041 Generative models have recently emerged as powerful tools for trajectory planning, with diffusion
042 models (Ho et al., 2020) and flow matching (FM) (Lipman et al., 2023) enabling the generation
043 of complex, long-horizon behaviors by directly learning from data. Unlike traditional data-driven
044 planners that combine learned dynamics with optimization routines (Posa et al., 2014; Kalakrishnan
045 et al., 2011), generative approaches avoid model exploitation—where optimizers produce trajec-
046 tories that perform well under the model but fail in reality due to approximation errors (Talvitie, 2014;
047 Ke et al., 2019). By training models to generate full trajectories that implicitly encode system dy-
048 namics, generative planners naturally capture multimodal (Huang et al., 2025), high-dimensional
049 behaviors while mitigating compounding errors and supporting task compositionality. These advan-
050 tages make generative approaches increasingly attractive for real-world planning and control tasks.
051

052 Despite these advantages, a critical limitation of generative model planners lies in their inability to
053 guarantee constraint satisfaction – specifically, state and action constraints. State constraints ensure
054 safety (Wabersich et al., 2023) (e.g., avoiding collisions), while action constraints guarantee admissi-
055 bility (Shen et al., 2018) (e.g., respecting torque or power limits), which makes these constraints
056 essential in domains such as robotics (Craig, 2009). However, constraint satisfaction at individual
057 time steps is insufficient: since trajectories are sequences of states (and actions), subsequent states in
058 trajectories cannot be chosen independently (Kelly, 2017). For a planned trajectory to be physically
059 realizable and its guarantees to transfer from plan to execution, the trajectory must be dynamically
060 consistent, i.e., states must evolve according to the dynamics of the system for which the trajec-
061 tory is planned. However, existing generative planners offer no inherent mechanism to enforce such
062 properties, and the non-static nature of real-world environments often leads constraints to be un-
063 derrepresented, or entirely absent, in the training data. These factors make constraint satisfaction
064 particularly challenging at test time and motivate the need for formally grounded methods that can
065 reliably ensure safety, admissibility, and consistency in generated trajectories.
066

067 A number of recent works have explored constrained planning for generative models by incorpo-
068 rating constraints into the training phase (Ho & Salimans, 2022; Ajay et al., 2023; Zheng et al., 2023),
069 injecting guidance signals during sampling (Dhariwal & Nichol, 2021; Yuan et al., 2023; Kondo
070 et al., 2024; Xiao et al., 2025), or applying post-processing (Mazé & Ahmed, 2023) corrections af-
071

054 ter generation. While these approaches improve constraint adherence, they remain fundamentally
 055 limited. Training-time methods struggle to handle unseen constraints, guidance-based techniques
 056 offer only soft biases without guarantees, and post-processing can distort the distribution learned by
 057 the generative model. These issues highlight three key challenges that need to be jointly addressed:
 058 (1) guaranteeing state and action constraint satisfaction, (2) ensuring dynamic consistency so that
 059 trajectories are physically valid, and (3) handling novel test-time constraints without retraining.

060 To address these challenges, we propose **SAD-Flower** – a novel control-augmented flow matching
 061 framework designed to generate **Safe**, **Admissible**, and **Dynamically** consistent trajectories. Inspired
 062 by guidance-based approaches, SAD-Flower introduces a virtual control input into the generation
 063 process. This control-theoretic interpretation of the sampling process enables a principled design of
 064 test-time guidance signals to ensure strong guarantees. The foundation of the design lies in a novel
 065 reformulation of state and action constraints into Control Barrier Function (CBF) conditions (Ames
 066 et al., 2017) for the generation process, while dynamic consistency is transformed into a Control
 067 Lyapunov Function (CLF) condition (Sontag, 1983). We exploit these conditions to design prin-
 068 cipled guidance signals via constrained optimal control problems, which can be solved efficiently
 069 using quadratic program solvers, such that SAD-Flower can handle novel constraints at test time
 070 without retraining. Unlike general constraint-projection approaches (Römer et al., 2025; Bouvier
 071 et al., 2025), which can significantly deviate from the original distribution learned by the genera-
 072 tive model, we leverage prescribed-time control concepts (Song et al., 2017) to flexibly schedule
 073 constraint enforcement. This allows us to formally prove the safety, admissibility, and dynamic
 074 consistency of trajectories generated by SAD-Flower. These theoretical guarantees translate into
 075 empirical performance: across a range of domains, SAD-Flower consistently ensures constraint sat-
 076 isfaction while achieving competitive or superior task performance compared to existing generative
 077 planners. Notably, SAD-Flower remains robust even under increasingly stricter test-time constraints,
 078 validating its reliability in challenging deployment scenarios.

079 2 RELATED WORK

080 **Diffusion and Flow-Based Generative Models for Planning.** Recent advances in generative
 081 modeling, including diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021)
 082 and flow-matching (Lipman et al., 2023; Albergo & Vanden-Eijnden, 2023), have shown remarkable
 083 performance across various domains such as image generation (Dhariwal & Nichol, 2021; Du et al.,
 084 2020) and language modeling (Liu et al., 2023; Saharia et al., 2022). These generative approaches
 085 have also been successfully applied to data-driven planning, where the model learns to imitate expert
 086 behavior from datasets. For example, some works generate entire state-action trajectories directly
 087 using one (Janner et al., 2022; Zheng et al., 2023) or two separate models (Zhou et al., 2024), while
 088 others predict high-level trajectories and rely on a downstream controller to compute low-level ac-
 089 tions (Chi et al., 2023; Ajay et al., 2023). However, these generative model planners operate without
 090 mechanisms to ensure that generated trajectories respect real-world constraints. In particular, they
 091 lack formal guarantees for satisfying state and action constraints, as well as dynamic consistency.

092 **Constraint-Aware Generative Model Planner.** To address the issue of constraint satisfaction, sev-
 093 eral recent works have explored constraint-aware planning based on generative models. Guidance-
 094 based methods (Dhariwal & Nichol, 2021; Yuan et al., 2023; Kondo et al., 2024; Ma et al., 2025;
 095 Carvalho et al., 2023) incorporate constraints by injecting gradients of auxiliary cost functions into
 096 the sampling process. While this encourages constraint satisfaction, it provides only a soft inductive
 097 bias without formal guarantees. Classifier-free guidance (Ho & Salimans, 2022) leverages infor-
 098 mation, such as constraint violation (Ajay et al., 2023) during training, enabling constraint-aware
 099 generation. However, these approaches require additional labeled data and have limited generaliza-
 100 tion to novel constraints. Similarly, DDAT (Bouvier et al., 2025) incorporates projection into the
 101 feasible set during training and inference but relies on strong assumptions, such as convexity of the
 102 constraint set. Post-processing approaches (Mazé & Ahmed, 2023; Giannone et al., 2023) attempt to
 103 enforce constraints by optimizing generated samples after denoising. However, these modifications
 104 are unaware of the learned data distribution and can produce samples that significantly drift from it.

105 **Control-Theoretic Enforcement in Generative Planning.** Several works have proposed control-
 106 theoretic techniques to enforce constraints. The works (Xiao et al., 2025; Botteghi et al., 2023; Dai
 107 et al., 2025) employ Control Barrier Functions (CBFs) (Ames et al., 2017) to enforce state con-

108 straints during the denoising process. However, these methods neglect action constraints and often
 109 produce trajectories that are not dynamically consistent and suffer from the local trap problem (Xiao
 110 et al., 2025). Control Lyapunov functions (CLFs) (Sontag, 1983) are leveraged together with CBFs
 111 in (Mizuta & Leung, 2024) to improve safety and stability, but formal guarantees are missing, and
 112 action constraints are not addressed. A constrained optimal control layer is integrated into the de-
 113 noising process in (Römer et al., 2025) to enforce state and action constraints. Since it performs non-
 114 convex optimization throughout the entire sampling process, it exhibits a high computational cost
 115 and potentially steers samples prematurely before they reflect meaningful structure (Fan et al., 2025).

117 3 PROBLEM SETTING

119 We formally define the trajectory planning problem with safety, admissibility and dynamic consis-
 120 tency constraints as follows.

121 **System Model.** We consider a nonlinear dynamical system with state $s(k) \in \mathbb{R}^n$ and action $a(k) \in$
 122 \mathbb{R}^m at time k , evolving as

$$123 \quad s(k+1) = f(s(k), a(k)), \quad (1)$$

124 where f is the (possibly unknown) transition function. A trajectory is defined as a sequence of state-
 125 action pairs, $\tau = \{(s(0), a(0)), \dots, (s(H-1), a(H-1))\}$. Given a dataset of expert trajectories
 126 $\mathcal{D} = \{\tau^{(n)}\}_{n=1}^N$, our goal is to learn planning new trajectories that both imitate expert behavior and
 127 respect all deployment constraints.

128 **Constraints.** To ensure reliable real-world execution, every generated trajectory must, at every time
 129 step, satisfy: (1) **Safety**: the state remains within a safe set \mathbb{S} (e.g., avoid collisions; eq. (SC));
 130 (2) **Admissibility**: the action is within an admissible set \mathbb{A} (e.g., satisfy torque or speed limits;
 131 eq. (AC)); and (3) **Dynamic Consistency**: the trajectory obeys the system dynamics (eq. (DC)).
 132 Neglecting any of these leads to unsafe, infeasible, or unrealizable plans: for example, trajectories
 133 may pass through obstacles (violating safety), demand unattainable actions (violating admissibility),
 134 or include state transitions that cannot be executed by the system (violating dynamics).

135 Formally, these requirements are posed on the distribution $p^\theta(\tau)$ of the learned planner as follows:

$$137 \quad \forall \tau \sim p^\theta(\tau) : \quad s(k) \in \mathbb{S}, \quad \forall k = 0, \dots, H-1, \quad (\text{SC})$$

$$138 \quad \forall \tau \sim p^\theta(\tau) : \quad a(k) \in \mathbb{A}, \quad \forall k = 0, \dots, H-1, \quad (\text{AC})$$

$$140 \quad \forall \tau \sim p^\theta(\tau) : \quad s(k+1) = f(s(k), a(k)), \quad \forall k = 0, \dots, H-1. \quad (\text{DC})$$

141 **Objective.** Given expert data \mathcal{D} and constraint sets \mathbb{S} , \mathbb{A} , our goal is to learn a generative model
 142 $p^\theta(\tau)$ that (i) matches the expert trajectory distribution, and (ii) ensures all sampled trajectories sat-
 143 isfy eqs. (SC), (AC), and (DC). This setting motivates methods that can flexibly enforce constraints,
 144 even as requirements change at test time.

146 4 BACKGROUND: LEARNING TO PLAN WITH FLOW MATCHING

148 When a trajectory data set \mathcal{D} of an expert planner is given, Flow Matching (FM) (Lipman et al., 2023;
 149 Zheng et al., 2023) is an effective technique to learn the distribution $p(\tau)$ of the data \mathcal{D} . In FM, the
 150 unknown distribution $p(\tau)$ is considered as the desired endpoint of a probability path $p_t^\theta(\tau)$, $t \in$
 151 $[0, 1]$. The remainder of the probability path $p_t^\theta(\tau)$ is characterized by a time-dependent vector field
 152 $v_t^\theta : [0, 1] \times \mathbb{R}^{(n+m)H} \rightarrow \mathbb{R}^{(n+m)H}$ parameterized by θ that acts on samples $\tau_0 \sim p_0(\tau)$ via the flow

$$153 \quad \dot{\tau}_t = \frac{d}{dt} \tau_t = v_t^\theta(\tau_t), \quad (2)$$

155 whereby the prior $p_0(\tau)$ is typically set to a Gaussian (Lipman et al., 2023). By prescribing a path
 156 from samples τ_0 of $p_0(\tau)$ to data trajectories τ_1 via a scheduled interpolation $\tau_t = \alpha(t)\tau_1 + \beta(t)\tau_0$
 157 with α, β such that $\alpha(0) = 0$, $\beta(1) = 0$ and $\alpha(t) + \beta(t) = 1$, the distribution learning problem is
 158 transformed into the supervised learning problem

$$159 \quad \mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}[0, 1], \tau_t \sim p_t^\theta(\tau), \tau_1 \sim p(\tau)} \|v_t^\theta(\tau_t) - v_t(\tau_0, \tau_1)\|_2^2, \quad (3)$$

160 where $v_t(\tau_0, \tau_1) = \dot{\alpha}(t)\tau_1 + \dot{\beta}(t)\tau_0$ follows from the interpolation. Minimizing this cost function
 161 via stochastic gradient descent, $v_t^\theta(\tau_t)$ can be efficiently trained using the trajectory data set \mathcal{D} .

Given an initial state s_0 and a trained vector field $v_t^\theta(\tau_t)$, we sample a random trajectory τ_0 from the prior distribution p_0 and numerically solve the ordinary differential equation (ODE) in eq. (2) using τ_0 as the initial condition to obtain τ_1 . Thereby, the trajectories effectively become samples $\tau_1 \sim p^\theta(\tau)$, where $p^\theta(\tau)$ is implicitly represented through $v_t^\theta(\tau_t)$. While such samples can be directly used in planning, they generally do not satisfy the safety, admissibility, and dynamic consistency constraints in eqs. (SC), (AC), and (DC).

Remark 4.1. *To allow plans with given initial states s_0 , we only need to condition the initial distribution on $s(0) = s_0$ and ensure $\dot{\tau}_t^{s(0)} = \mathbf{0}$ for all $t \in [0, 1]$. For training, s_0 is chosen as the first state of trajectories in the data set, while arbitrary values can be set when sampling trajectories.*

5 CONTROL AUGMENTED FLOW MATCHING FOR CONSTRAINED PLANNING

Despite the power of FM-based planners for trajectory generation, ensuring safety, admissibility, and dynamic consistency remains challenging, particularly under new test-time constraints. We address this with SAD-Flower, a control-augmented flow matching framework that provides formal guarantees for constraint satisfaction. In Section 5.1, we outline the control-theoretic intuition and its integration with FM. Section 5.2 details how constraint-aware quadratic programming augments sampling, and Section 5.3 presents theoretical guarantees of convergence and constraint satisfaction.

5.1 CONTROL AUGMENTATION FOR SAFETY, ADMISSIBILITY AND DYNAMIC CONSISTENCY

To ensure the satisfaction of safety (eq. (SC)), admissibility (eq. (AC)) and dynamic consistency (eq. (DC)), we extend the formulation in eq. (2) at test time to a controlled dynamical system

$$\dot{\tau}_t = v_t^\theta(\tau_t) + u_t, \quad (4)$$

where the vector field $v_t^\theta(\tau_t)$ represents the drift, while u_t is a control input. By choosing $u_t = \mathbf{0}$, we recover standard flow matching as a special case of this formulation. Framing the problem in this way enables us to view the requirements in eqs. (SC), (AC), and (DC) as system properties, so that their satisfaction becomes a matter of control design with the following specifications.

From State/Action Constraints to Barrier Specifications. State and action constraints are set inclusion conditions, which require the controlled flow in eq. (4) to converge to and subsequently maintain constraint satisfaction. This behavior can be formalized via control barrier functions (CBFs) (Ames et al., 2017), whose level sets can encode the constraint sets \mathbb{A} and \mathbb{S} . Thereby, we express state and action constraints as a condition on the growth of CBFs along the flow.

From Dynamic Consistency to Lyapunov Specifications. Dynamic consistency is an equality condition, whose violation needs to decay to 0 along the flow in eq. (4). This property can be formalized using control Lyapunov functions (CLF) (Sontag, 1983) – energy-like, non-negative functions with a minimum of 0 at the desired equilibrium. Hence, we formulate dynamic consistency as a condition on the decay of a CLF along the flow.

Prescribed-time Specifications. While FM simulates ODEs for time intervals $t \in [0, 1]$, Lyapunov and barrier specifications usually relate to asymptotic guarantees with $t \rightarrow \infty$. This discrepancy necessitates the scheduling of a sufficiently fast growth of CBFs and decrease of CLFs along the flow, which corresponds to a prescribed-time specification for control (Song et al., 2017).

Our approach – SAD-Flower – splits the numerical integration of eq. (4) into two phases as illustrated in Fig. 1. In the uncontrolled phase ($0 \leq t < T_0$), trajectories evolve under the learned FM vector field without intervention ($u_t = \mathbf{0}$) to preserve sample diversity (Fan et al., 2025). In the controlled phase ($T_0 \leq t \leq 1$), the control law u_t satisfying CLF, CBF, and prescribed-time specifications is applied when integrating eq. (4). Thereby, the activation time $T_0 \in (0, 1)$ allows SAD-Flower to effectively balance generative flexibility with formal guarantees on safety (eq. (SC)), admissibility (eq. (AC)), and dynamic consistency (eq. (DC)).

5.2 CONTROL DESIGN USING CONTROL LYAPUNOV AND BARRIER FUNCTIONS

Given the control specifications in Section 5.1, the actual control design problem remains. We first derive dedicated CBF and CLF constraints, which employ scheduling functions to ensure a sufficient growth/decrease rate. These constraints are exploited in an optimization-based control law.

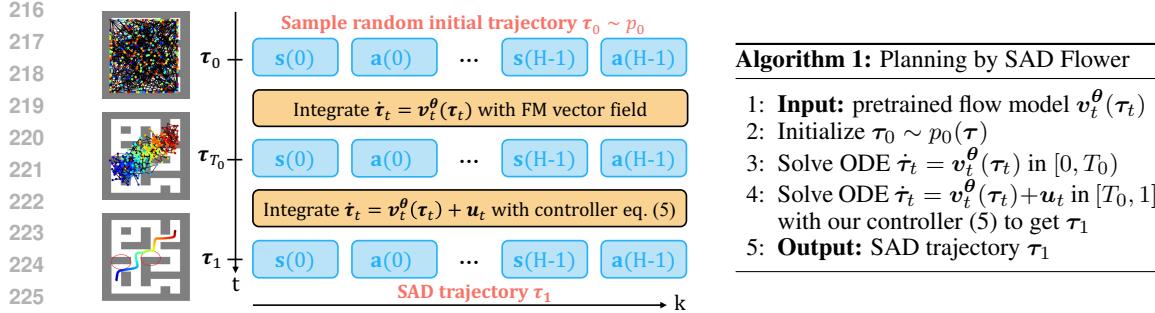


Figure 1: Overview of the trajectory generation using our proposed SAD-Flower.

Control Barrier Constraints. Since state/action constraints are specified separately for each time step, we design state CBFs $h_k^s(\tau_t)$ and action CBFs $h_k^a(\tau_t)$ for each time step $k = 0, \dots, H-1$. Each CBF itself is defined as a signed distance function (SDF) (Park et al., 2019; Long et al., 2021), which measures the distance to the boundary of the set \mathbb{S} and \mathbb{A} , respectively, and assigns a sign based on the inclusion in the constraint set.¹ Thus, these functions are only positive if states $s(k)$ and actions $a(k)$ are inside the sets \mathbb{S} and \mathbb{A} , respectively. Non-negativity of the CBFs can be ensured by constraining the evolution of the CBF values along the flow, which results in the derivative condition

$$\begin{aligned} \dot{h}_k^s(\tau_t) &\geq -\varphi(t)h_k^s(\tau_t), \quad \forall k = 1, \dots, H-1, \\ \dot{h}_k^a(\tau_t) &\geq -\varphi(t)h_k^a(\tau_t), \quad \forall k = 0, \dots, H-1, \end{aligned} \quad \begin{array}{l} \text{(CBF-s)} \\ \text{(CBF-a)} \end{array}$$

where $\dot{h}_k^{s/a}(\tau_t) = \nabla^T h_k^{s/a}(\tau_t) (\mathbf{v}_t^\theta(\tau_t) + \mathbf{u}_t)$ (Ames et al., 2017) and $\varphi(t)$ is a scheduling function that we will design later.

Control Lyapunov Constraints. For defining a suitable CLF, we sum up the squared consistency errors of a trajectory, i.e., $V(\tau_t) = \frac{1}{2} \sum_{k=1}^{H-1} \|s(k) - \mathbf{f}(s(k-1), a(k-1))\|^2$. This function is only 0 if the trajectory τ_t is dynamically consistent, such that we constrain the evolution of its value along the flow to be negative via

$$\dot{V}(\tau_t) \leq -\varphi(t)V(\tau_t), \quad \text{(CLF)}$$

where $\dot{V}(\tau_t) = \nabla^T V(\tau_t) (\mathbf{v}_t^\theta(\tau_t) + \mathbf{u}_t)$ (Sontag, 1983) and $\varphi(t)$ is a scheduling function that we will design later. Computing $\nabla^T V(\tau_t)$ requires access to \mathbf{f} . In some robotic domains this information is available through high fidelity simulators (Gaz et al., 2019; Howell et al., 2022; Acosta et al., 2022). If there exists no accurate physics-based simulator, e.g., for contact-rich manipulation tasks, a dynamics model can also be learned from the trajectory data. The learning errors directly correlate to the magnitude of dynamic consistency violations (Curi et al., 2020), such that practical consistency can still be achieved given sufficiently precise learned model.

Prescribed-time Scheduling. To guarantee that the CBFs are positive and the CLF is 0 at the terminal time $t = 1$ regardless of its state at $t = T_0$, we employ a scheduling function $\varphi(t) = \frac{c}{(1-t)^2}$ with some constant $c > 0$. Due to the steep growth of φ for $t \rightarrow 1$, the constraints in eqs. (CBF-s), (CBF-a) and (CLF) become increasingly more restrictive. This ensures that positivity of CBFs and a vanishing CLF are ensured at some time $t < 1$ (Song et al., 2023).

Constrained Minimum-Norm Optimal Control. To ensure the satisfaction of CLF and CBF constraints, we formulate the minimum-norm optimal control problem

$$\mathbf{u}_t = \min_{\mathbf{u}} \|\mathbf{u}\|^2 \quad \text{s.t. eqs. (CBF-s), (CBF-a) and (CLF) hold,} \quad (5)$$

which ensures them by construction, while simultaneously minimizing the perturbation of the learned vector field $\mathbf{v}_t^\theta(\tau_t)$. Even though this optimization problem often consists of a large number of optimization variables and constraints, it can be solved comparatively efficiently since it is a quadratic program (QP). This renders the numerical integration of eq. (4) with the control law in eq. (5) computationally tractable when using dedicated QP solvers.

¹For a formal definition of the CBFs $h_k^s(\tau_t)$ and $h_k^a(\tau_t)$, see Appendix A

270 5.3 THEORETICAL GUARANTEES FOR CONSTRAINT SATISFACTION AND CONSISTENCY.
271272 Due to the strong theoretical foundations of CBFs and CLFs, strong guarantees for safety, admissi-
273 bility, and dynamic consistency can be provided as shown in the following result.²274 **Theorem 5.1.** *Assume that the QP in eq. (5) is feasible for all $t \in [T_0, 1]$. Then, the solution τ_t of
275 eq. (4) with control law $u_t = \mathbf{0}$ for $t < T_0$ and u_t defined in eq. (5) for $t \geq T_0$ satisfies eqs. (SC),
276 (AC), and (DC) at $t = 1$ for all initial conditions τ_0 .*277 This result shows that the validity of the safety, admissibility, and dynamic consistency guarantees
278 essentially depends on the feasibility of the minimum-norm optimal control law in eq. (5), which is
279 a fundamental requirement for a well-defined solution τ_1 . While feasibility hinges on three tech-
280 nical conditions, it is not an issue in practice for the designed CLF and CBF constraints. Firstly,
281 infeasibility often occurs only on sets of trajectories τ with zero measure, such that it does not occur
282 during the numerical integration of eq. (4). Secondly, many strategies for recovering feasibility ex-
283 ist, e.g., via slack variables (Boyd & Vandenberghe, 2004). If there exists a non-empty time interval
284 including $t = 1$ after employing such a recovery strategy, it follows from Theorem 5.1 that eqs.
285 (SC), (AC), and (DC) remain guaranteed. Hence, feasibility is usually not a practical concern.
286287 6 EXPERIMENT
288289 6.1 EXPERIMENT SETTING
290291 We evaluate all methods in 4 different benchmark tasks: Maze2d, Hopper, and Walker2d from the
292 D4RL problem set (Fu et al., 2020), and Kuka Block-Stacking (Janner et al., 2022).³293 • **Maze2d** (Fu et al., 2020) is a navigation task where a point mass is moved from an initial state to
294 a specified goal. Actions are artificially constrained to $[-0.1, 0.1]$. State constraints are defined
295 for two novel obstacles, which do not block feasible paths and therefore still allow the task to be
296 completed. Training data is generated by a navigation planner for the given maze. We evaluate
297 on two maze configurations: Large and Umaze.298 • **Hopper and Walker2d** (Fu et al., 2020) are locomotion tasks where a one-legged and a bipedal
299 robot must move forward by jumping and walking, respectively. Actions are constrained to the
300 range $[-1, 1]$. State constraints are imposed by requiring the robot’s torso center to remain below
301 a prescribed height (default: 1.6), creating a conflict with the objective of fast forward movement.
302 We evaluate two datasets: demonstrations from a partially trained soft actor-critic policy (Haarnoja
303 et al., 2018) (Medium) and a mixture of expert and partially trained policy data (Med-Expert).304 • **Kuka Block-Stacking** (Janner et al., 2022) is a manipulation task where a 7-DOF robotic arm
305 must stack a set of blocks. Unlike the other tasks, no action constraints or dynamic consistency re-
306 quirements are enforced; only state constraints are applied, allowing us to study a variant of trajec-
307 tory planning focused solely on state feasibility. These constraints ensure that self-collisions of the
308 robot are avoided. Training data is generated using the PDDLStream planner (Garrett et al., 2020).309 **Evaluation Metrics.** We evaluate the safety and admissibility violation of a trajectory via its max-
310 imal distance to the constraint sets \mathbb{S} and \mathbb{A} , which we can effectively express through the CBFs as
311 $-\min_k \min\{h_k^{s,a}(\tau), 0\}$. Thus, a value of 0 means constraint satisfaction, while positive values im-
312 ply a violation. Dynamical consistency is measured using the Lyapunov function, such that large val-
313 ues indicate inconsistency. The performance of planned trajectories is measured by normalized total
314 rewards for D4RL tasks and binary success rewards for stacking, as defined in Janner et al. (2022).315 **Baselines:** We compare our proposed SAD-Flower against the following generative planners:316 • **Diffuser** (Janner et al., 2022) generates trajectories using a diffusion model, without considering
317 safety, admissibility, or dynamic consistency.
318 • **Truncation (Trunc)** (Brockman et al., 2016) enforces constraints by truncating the trajectory
319 generated from the diffusion model.
320 • **Classifier Guidance (CG)** (Dhariwal & Nichol, 2021) augments diffusion-based trajectory gen-
321 eration with constraint-gradient guidance during sampling.322 ²A proof and extended discussion of the theorem’s assumptions can be found in Appendix B.
323 ³Details of the experimental setting are provided in the Appendix. D

324

- 325 • **Flow Matching (FM)** (Feng et al., 2025) trains a model with FM to generate trajectories, which
- 326 is used as a baseline that does not incorporate safety, admissibility, or dynamic consistency.
- 327
- 328 • **SafeDiffuser (S-Diffuser)** (Xiao et al., 2025) generates trajectories via diffusion, while projecting
- 329 states onto the constraint sets using a CBF at each sampling step.
- 330 • **Decision Diffuser (D-Diffuser)** (Ajay et al., 2023) trains a conditional diffusion model to
- 331 generate state trajectories conditioned on task information and constraints, with the corresponding
- 332 actions inferred from a learned inverse dynamics model.

333

334 6.2 CONSTRAINED-PLANNING PERFORMANCE ACROSS BENCHMARKS

335

336 As shown in Table 1, our method consistently satisfies safety and admissibility constraints while

337 matching the planning performance of other methods across tasks. Although perfect dynamical

338 consistency is not reached, the remaining violations are minor and mainly due to numerical integration

339 of eq. (4). These results demonstrate the effectiveness of our control-theoretic mechanism, which

340 guarantees constraint satisfaction at test time. **SAD-Flower achieves this performance at the cost of**

341 **merely a small computation time increase compared to existing methods (Appendix, Table 13).**

342

343 In the navigation tasks of Maze2D, SAD-Flower achieves high rewards while maintaining complete

344 constraint satisfaction. Since the task goals and imposed constraints are not in conflict, our method

345 effectively balances planning performance and constraint adherence. Among baselines, SafeDiffuser

346 employs a constraint-following mechanism during the diffusion sampling process to enforce state

347 constraints, but it fails to guarantee admissibility. In contrast, Diffuser and FM achieve high rewards

348 at the cost of violations in both safety and admissibility.

349

350 Table 1: Performance of the proposed SAD-Flower and baselines across navigation, locomotion, and

351 manipulation tasks. The methods are compared on the maximum safety and admissibility constraint

352 violations of planned trajectories, the magnitude of the dynamic consistency violation, and the model

353 accuracy expressed through the reward. Truncate is not applicable in Maze2d (Umaze) due to more

354 complex safety constraints, such that truncation becomes non-trivial (Xiao et al., 2025).

355

356 Experiment	357 Metric	358 Diffuser	359 Trunc	360 CG	361 FM	362 S-Diffuser	363 D-Diffuser	364 Ours
365 Maze2d (Large)	safety	0.43±0.39	0.10±0.25	0.27±0.37	0.37±0.39	0.00±0.00	0.83±0.24	0.00±0.00
	admissib.	0.89±0.01	0.88±0.07	0.87±0.02	0.90±0.00	0.91±0.02	0.90±0.03	0.00±0.00
	dyn. consist.	0.06±0.03	0.06±0.03	0.44±0.16	0.02±0.00	0.09±0.06	2.78±0.06	0.01±0.01
	reward	1.40±0.26	1.39±0.26	0.40±0.33	1.43±0.20	1.20±0.06	0.38±0.18	1.42±0.52
366 Maze2d (Umaze)	safety	0.04±0.20	—	0.51±0.33	0.11±0.22	0.87±3.77	0.05±0.15	0.00±0.00
	admissib.	0.90±0.01	—	0.88±0.03	0.01±0.00	0.90±0.02	0.89±0.02	0.00±0.00
	dyn. consist.	0.05±0.02	—	0.68±0.18	0.01±0.01	0.10±0.10	1.80±0.11	0.01±0.01
	reward	1.11±0.44	—	0.06±0.32	2.62±1.09	1.06±0.35	0.60±0.33	2.66±0.88
367 Hopper (Med-Expert)	safety	0.01±0.02	0.05±0.04	0.07±0.03	0.11±0.08	0.05±0.04	0.10±0.02	0.00±0.00
	admissib.	0.21±0.05	0.18±0.04	0.26±0.07	0.17±0.05	0.18±0.04	0.00±0.00	0.00±0.00
	dyn. consist.	0.38±0.03	0.41±0.04	0.79±0.10	0.23±0.02	0.36±0.06	0.16±0.01	0.01±0.01
	reward	1.06±0.18	0.50±0.12	0.73±0.22	1.02±0.20	0.53±0.19	1.12±0.01	0.93±0.23
368 Hopper (Medium)	safety	0.01±0.02	0.05±0.03	0.00±0.00	0.39±0.13	0.01±0.01	0.15±0.01	0.00±0.00
	admissib.	0.21±0.05	0.18±0.04	0.16±0.03	0.32±0.21	0.18±0.04	0.00±0.00	0.00±0.00
	dyn. consist.	0.46±0.01	0.47±0.01	0.95±0.02	0.42±0.39	0.47±0.01	0.18±0.01	0.01±0.01
	reward	0.44±0.05	0.45±0.06	0.39±0.03	0.49±0.05	0.45±0.06	0.48±0.08	0.34±0.03
369 Walker2d (Med-Expert)	safety	0.06±0.04	0.06±0.05	0.02±0.03	0.40±0.11	0.09±0.07	0.04±0.04	0.00±0.00
	admissib.	0.67±0.18	0.62±0.14	0.72±0.18	0.15±0.02	0.58±0.20	0.00±0.00	0.00±0.00
	dyn. consist.	0.71±0.05	0.72±0.05	0.83±0.91	0.44±0.01	0.79±0.03	0.69±0.05	0.04±0.04
	reward	1.06±0.23	0.56±0.29	0.39±0.19	1.07±0.01	0.59±0.21	0.95±0.24	0.89±0.32
370 Walker2d (Medium)	safety	0.03±0.03	0.02±0.01	0.02±0.02	0.21±0.15	0.02±0.02	0.09±0.04	0.00±0.00
	admissib.	0.56±0.10	0.44±0.19	0.54±0.14	0.48±0.06	0.52±0.12	0.00±0.00	0.00±0.00
	dyn. consist.	0.68±0.08	0.65±0.08	0.72±0.38	0.40±0.06	0.64±0.07	1.52±0.05	0.07±0.15
	reward	0.57±0.26	0.50±0.26	0.55±0.28	0.73±0.15	0.49±0.23	0.76±0.16	0.42±0.23
371 KUKA Block Stacking	safety	0.23±0.09	0.00±0.00	0.22±0.09	0.02±0.04	0.00±0.00	0.14±0.13	0.00±0.00
	reward	0.46±0.23	0.45±0.21	0.45±0.23	0.44±0.20	0.49±0.23	0.55±0.26	0.45±0.21

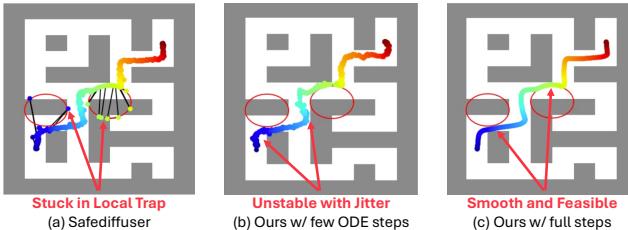


Figure 2: (a) Without enforcing dynamic consistency, applying state and action constraints can result in significant outliers, known as the local trap problem (Xiao et al., 2025). (b) Our method satisfies constraints, but using too few numerical integration steps for the ODE introduces jitter in the trajectory. (c) With sufficient integration steps, our method produces dynamically consistent trajectories while satisfying unseen constraints (red ellipses).

In locomotion tasks, where dynamics are more complex, SAD-Flower is the only method that satisfies all constraints while maintaining strong dynamical consistency. Its slightly lower rewards compared to baselines stem from conflicts between the imposed height constraint and the jumping or walking behaviors required for high returns, making safety violations directly correlated with the rewards that reflect planning performance. For instance, CG achieves planning performance comparable to SAD-Flower in Hopper (Med-Expert) when both methods avoid safety violations. While some methods, such as Decision Diffuser, consistently ensure admissibility, all struggle with dynamical consistency due to the complexity of the robot dynamics.

In the Kuka Block-Stacking task, which excludes admissibility and dynamical consistency requirements, SAD-Flower leverages the simplified setting to guarantee safety while achieving rewards competitive with other safety-enforcing baselines. This highlights both the effectiveness and flexibility of our proposed approach.

6.3 WHY DOES SAD-FLOWER WORK EFFECTIVELY?

We analyze three key properties of SAD-Flower that contribute to its effectiveness.

Dynamic consistency prevents local traps. Safety and admissibility can be enforced by projecting states or actions back into the constraint sets, but this can introduce sharp discontinuities in the trajectory, as illustrated in Fig. 2(a) for SafeDiffuser in Maze2D. This phenomenon, known as the local trap problem (Xiao et al., 2025), arises because constraints are treated independently at each trajectory step. In contrast, SAD-Flower enforces dynamic consistency by coupling consecutive states and actions through the CLF, ensuring coherent evolution during integration of the flow. This coupling prevents misaligned guidance and eliminates the risk of local traps, as demonstrated in Fig. 2(c).

Delayed control activation avoids premature interventions. At the beginning of the flow sampling process, sample distributions are relatively unstructured, and applying control signals too early can push trajectories away from behaviors captured by the learned generative planner. Perturbations introduced in this phase are propagated throughout the flow in eq. (4), potentially degrading performance. SAD-Flower mitigates this issue by activating guidance only after a prescribed flow time T_0 . As shown in Fig. 3 (left) for Maze2D, even a small T_0 suffices to achieve high rewards, while a larger T_0 can further improve performance.

Dynamic consistency violations increase only marginally with larger T_0 . Safety and admissibility are satisfied across all T_0 values in our experiments (not shown due to space). These results highlight the effectiveness of our control-theoretic formulation, which allows delayed activation of CLF and CBF guidance without compromising constraints or dynamic consistency.

QP vs. non-convex optimization. To assess the benefits of our QP-based formulation, we compare SAD-Flower with DPCC (Römer et al., 2025), a recent method

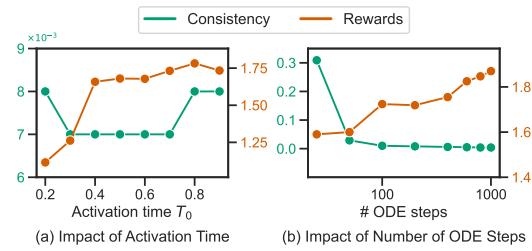


Figure 3: Effect of increasing the number of ODE steps for solving eq. (4) and the activation time T_0 of controller eq. (5) on the consistency and performance of trajectories.

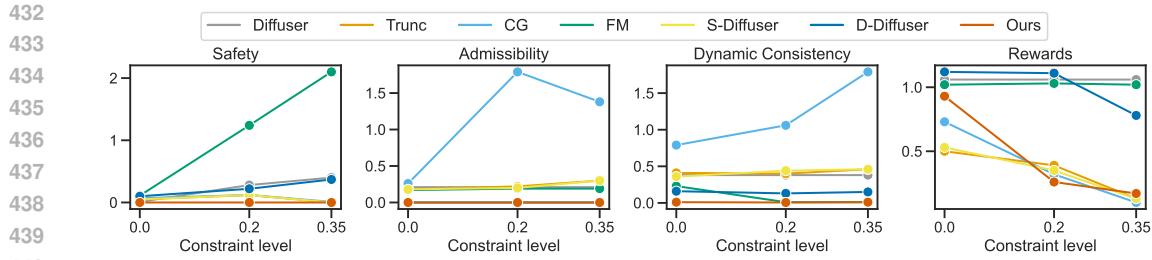


Figure 4: Performance of the proposed SAD-Flower and baselines depending on level of constraint tightening for the Hopper (Med-Expert) benchmark.

that also integrates optimal control into generative planning to enforce safety, admissibility, and dynamic consistency. DPCC applies control from the earliest stages and relies on full-system non-convex optimization at every sampling step.

As shown in Table 2 for Hopper (Med-Expert), both methods achieve comparable levels of admissibility and dynamic consistency, but DPCC incurs much higher computation per sampling step and lower rewards. This highlights the efficiency of SAD-Flower gained without any loss of constraint satisfaction by enabling the usage of QP solvers compared to general non-convex optimization methods. Note though that the enforcement of constraints does still not come for free with QPs since the integration steps of the uncontrolled ODE merely takes 0.01s on average.

6.4 HOW RELIABLE IS SAD-FLOWER?

To highlight the reliability of SAD-Flower, we demonstrate its behavior when approaching the extremes in terms of approximations for implementation and problem difficulty.

Reducing integration accuracy. We first illustrate the reliability of SAD-Flower when the accuracy of integrating the controlled ODE in eq. (4) is small, which we measure through the number of ODE discretization steps. As illustrated on the right of Fig. 3, increasing the number of ODE steps improves dynamic consistency since the controller in eq. (5) has more opportunities to intervene. Even with as few as 25 ODE steps, the inconsistency is relatively small and corresponds only to minor jittering behavior in the trajectories as illustrated in Fig. 2. With ≈ 100 ODE steps, dynamic consistency is practically achieved. Safety and admissibility are ensured for the whole range of considered ODE steps (not depicted). In contrast, the rewards continue to grow for finer discretizations, which cause higher computational complexity. Thus, SAD-Flower allows to trade-off performance and complexity, while simultaneously ensuring constraint satisfaction and dynamic consistency.

Handling stricter test-time constraints. To evaluate the robustness of our method under unseen conditions, we test it with increasingly restrictive constraints. For this, we reduce the allowed torso height by 0.2 and 0.35 in the Hopper environment. As shown in Fig. 4, SAD-Flower maintains perfect constraint satisfaction even under significantly stricter settings. Although the rewards achieved by SAD-Flower decrease with tighter constraints, this effect is caused by the conflict of constraints and objective. Thus, other methods enforcing state constraints, e.g., SafeDiffuser, also exhibit this behavior. Notably, the rewards of SAD-Flower are generally at least on par with these methods, even though the baselines exhibit significant admissibility and dynamic consistency violations. This clearly demonstrates the strong robustness of SAD-Flower in handling unseen test-time constraints.

Scalability in high-dimensional tasks. To evaluate the scalability of our approach in complex, contact-rich manipulation settings, we additionally consider the D4RL Adroit Relocate task, which involves a 39-dimensional state space and a 30-dimensional action space. This dexterous grasping environment poses significant challenges due to the need for precise contact interactions and high-dimensional control. We define the state constraint via

Table 2: Benefits of QPs in SAD-flower compared to general non-convex optimization of DPCC (Römer et al., 2025).

Experiment	Metric	DPCC	Ours
Hopper (Med-Expert)	safety	0.01 \pm 0.03	0.00 \pm 0.00
	admissib.	0.00 \pm 0.00	0.00 \pm 0.00
	dyn. consist.	0.01 \pm 0.01	0.01 \pm 0.01
	reward	0.61 \pm 0.07	0.93 \pm 0.23
	comp. time [s]	4.24 \pm 1.64	0.06 \pm 0.00

Table 3: Performance of SAD-Flower and FM in a dexterous grasping scenario.

Experiment	Metric	FM	Ours
Adroit-Hand (Relocate)	safety	0.15 \pm 0.21	0.00 \pm 0.00
	admissib.	0.62 \pm 0.19	0.00 \pm 0.00
	dyn. consist.	0.07 \pm 0.04	0.06 \pm 0.09
	reward	1.07 \pm 0.08	1.05 \pm 0.05

486 the arm’s position to avoid collisions and the action constraint through actuator limits. As shown in
 487 Table 3, standard flow matching achieves moderate task success but suffers from frequent violations
 488 of safety and admissibility constraints due to unconstrained sampling. In contrast, SAD-Flower
 489 consistently satisfies both state and action constraints, demonstrating its ability to scale to high-
 490 dimensional robotic systems while maintaining constraint enforcement. This result underscores the
 491 robustness and generality of our control-augmented framework in realistic manipulation tasks.

492 **Sensitivity analysis of constraint-enforcement strength.** We conduct an ablation study on the
 493 LargeMaze task to assess the sensitivity of SAD-Flower to the hyperparameter c , which modu-
 494 lates the strength of constraint enforcement. While the main experiments use $c = 0.5$, we vary
 495 it from 1.0 to 0.2. As shown in Table 4, SAD-Flower maintains perfect safety and admissibility
 496 across all tested values. Violations of dynamic consistency remain negligible and stable, while task
 497 performance varies moderately. These results demonstrate that SAD-Flower achieves constraint
 498 satisfaction without requiring precise hyperparameter tuning, highlighting its practical reliability.

499 **Robustness under imperfect dynamics models.** To evaluate the robustness of our method un-
 500 der model approximation errors, we simulate degraded dynamics by training the forward model on
 501 progressively smaller subsets of the Hopper-Medium dataset. As shown in Table 5, SAD-Flower
 502 maintains full constraint satisfaction even with only 10% of the original data, and dynamic consis-
 503 tency remains stable across these settings. Degradation becomes apparent only when the dataset
 504 is reduced to 0.01%, at which point safety violations emerge. These results demonstrate that SAD-
 505 Flower is tolerant to moderately inaccurate dynamics models and that training a sufficiently accurate
 506 model is feasible in practice.

507 Table 4: Sensitivity of SAD-Flower to the hyperparameter c .

Experiment	Metric	1.0	0.8	0.6	0.4	0.2
Maze2d (Large)	safety	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
	admissib.	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
	dyn. consist.	0.01±0.01	0.01±0.01	0.01±0.01	0.01±0.01	0.02 ± 0.01
	reward	1.38±0.59	1.54±0.59	1.46±0.52	1.47±0.53	1.29 ± 0.73

515 Table 5: Effect of training dataset size on dynamic consistency and constraint satisfaction.

Experiment	Metric	90%	70%	50%	30%	10%	0.01%
Hopper (Medium)	safety	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.02±0.09
	admissib.	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
	dyn. consist.	0.01±0.02	0.01±0.01	0.01±0.01	0.01±0.02	0.01±0.01	0.04±0.02
	reward	0.35±0.01	0.33±0.04	0.35±0.07	0.39±0.07	0.38±0.03	0.36±0.03

524 7 CONCLUSION

526 We presented SAD-Flower, a control-augmented flow matching framework that ensures safe, admis-
 527 sible, and dynamically consistent trajectory planning. By reformulating flow matching as a control-
 528 lable dynamical system with a virtual control input, our method leverages Control Barrier Function
 529 and Control Lyapunov Function conditions scheduled using prescribed-time control principles to
 530 enforce constraints at test time without retraining. Experiments across navigation, locomotion, and
 531 manipulation tasks show that SAD-Flower achieves perfect constraint satisfaction, avoids local traps,
 532 and maintains competitive task performance. These results establish it as a practical, theoretically
 533 grounded solution for real-world deployment. Looking ahead, extending SAD-Flower to stochastic
 534 dynamics, complex constraint geometries, and online replanning offers promising directions for safe
 535 and reliable generative trajectory planning.

536 537 REPRODUCIBILITY STATEMENT

538 **Code.** Our method is implemented based on the publicly available codebase from (Feng et al.,
 539 2025) at https://github.com/AI4Science-WestlakeU/flow_guidance. The code

540 for our method will be released upon acceptance. Implementation and experimental details are
 541 provided in section 6 of the main text, with further information in the Appendix.
 542

543 **Theory.** Our theoretical claims in Section 5.3 are supported by complete proofs in Appendix B.
 544

545 **Datasets.** We evaluate our method on a variety of public datasets across locomotion, navigation,
 546 and manipulation tasks. Full descriptions of the datasets can be found in Section 6 and Appendix D.
 547

548 **Compute.** All experiments were conducted using NVIDIA Tesla P100 GPUs. Additional information
 549 on hardware and computational resources is provided in Appendix G.
 550

551 BIBLIOGRAPHY

552 Brian Acosta, William Yang, and Michael Posa. Validating robotics simulators on real-world im-
 553 pacts. *IEEE Robotics and Automation Letters*, 7(3):6471–6478, 2022.

554 Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is
 555 conditional generative modeling all you need for decision-making? *The Eleventh International
 556 Conference on Learning Representations*, 2023.

557 Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic inter-
 558 polants. *The Eleventh International Conference on Learning Representations*, 2023.

559 Aaron D. Ames, Xiangru Xu, Jessy W. Grizzle, and Paulo Tabuada. Control barrier function based
 560 quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):
 561 3861–3876, 2017.

562 Nicolò Botteghi, Federico Califano, Mannes Poel, and Christoph Brune. Trajectory generation, con-
 563 trol, and safety with denoising diffusion probabilistic models. *arXiv preprint arXiv:2306.15512*,
 564 2023.

565 Jean-Baptiste Bouvier, Kanghyun Ryu, Kartik Nagpal, Qiayuan Liao, Koushil Sreenath, and Ne-
 566 gar Mehr. Ddat: Diffusion policies enforcing dynamically admissible robot trajectories. *arXiv
 567 preprint arXiv:2502.15043*, 2025.

568 Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

569 Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and
 570 Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

571 Joao Carvalho, An T Le, Mark Baierl, Dorothea Koert, and Jan Peters. Motion planning diffusion:
 572 Learning and planning of robot motions with diffusion models. In *2023 IEEE/RSJ International
 573 Conference on Intelligent Robots and Systems (IROS)*, pp. 1916–1923. IEEE, 2023.

574 Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake,
 575 and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The Inter-
 576 national Journal of Robotics Research*, pp. 02783649241273668, 2023.

577 Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games,
 578 robotics and machine learning, 2016.

579 John J Craig. *Introduction to robotics: mechanics and control*, 3/E. Pearson Education India, 2009.

580 Sebastian Curi, Felix Berkenkamp, and Andreas Krause. Efficient model-based reinforcement learn-
 581 ing through optimistic policy search and planning. *Advances in Neural Information Processing
 582 Systems*, 33:14156–14170, 2020.

583 Xiaobing Dai, Zewen Yang, Dian Yu, Shanshan Zhang, Hamid Sadeghian, Sami Haddadin, and
 584 Sandra Hirche. Safe flow matching: Robot motion planning with control barrier functions. *arXiv
 585 preprint arXiv:2504.08661*, 2025.

586 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances
 587 in neural information processing systems*, 34:8780–8794, 2021.

594 Zibin Dong, Jianye Hao, Yifu Yuan, Fei Ni, Yitian Wang, Pengyi Li, and Yan Zheng. Diffuserlite:
 595 Towards real-time diffusion planning. *Advances in Neural Information Processing Systems*, 37:
 596 122556–122583, 2024.

597 Yilun Du, Shuang Li, and Igor Mordatch. Compositional visual generation with energy based mod-
 598 els. *Advances in Neural Information Processing Systems*, 33:6637–6647, 2020.

600 Weichen Fan, Amber Yijia Zheng, Raymond A Yeh, and Ziwei Liu. Cfg-zero*: Improved classifier-
 601 free guidance for flow matching models. *arXiv preprint arXiv:2503.18886*, 2025.

603 Lang Feng, Pengjie Gu, Bo An, and Gang Pan. Resisting stochastic risks in diffusion planners with
 604 the trajectory aggregation tree. *arXiv preprint arXiv:2405.17879*, 2024.

605 Ruiqi Feng, Chenglei Yu, Wenhao Deng, Peiyan Hu, and Tailin Wu. On the guidance of flow
 606 matching. In *Forty-second International Conference on Machine Learning*, 2025.

608 Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep
 609 data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.

611 Caelan Reed Garrett, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Pddlstream: Integrating
 612 symbolic planners and blackbox samplers via optimistic adaptive planning. In *Proceedings of the*
 613 *international conference on automated planning and scheduling*, volume 30, pp. 440–448, 2020.

614 Claudio Gaz, Marco Cognetti, Alexander Oliva, Paolo Robuffo Giordano, and Alessandro De Luca.
 615 Dynamic identification of the franka emika panda robot with retrieval of feasible parameters using
 616 penalty-based optimization. *IEEE Robotics and Automation Letters*, 4(4):4147–4154, 2019.

618 Giorgio Giannone, Akash Srivastava, Ole Winther, and Faez Ahmed. Aligning optimization trajec-
 619 tories with diffusion models for constrained design generation. *Advances in neural information*
 620 *processing systems*, 36:51830–51861, 2023.

621 David Gilbarg and Neil S Trudinger. *Elliptic partial differential equations of second order*, volume
 622 224. Springer, 1977.

624 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
 625 maximum entropy deep reinforcement learning with a stochastic actor. In *International confer-
 626 ence on machine learning*, pp. 1861–1870. Pmlr, 2018.

627 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
 628 *arXiv:2207.12598*, 2022.

630 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
 631 *neural information processing systems*, 33:6840–6851, 2020.

633 Taylor A Howell, Simon Le Cleac'h, Jan Brüdigam, J Zico Kolter, Mac Schwager, and Zachary
 634 Manchester. Dojo: A differentiable physics engine for robotics. *arXiv preprint arXiv:2203.00806*,
 635 2022.

636 Tzu-Yuan Huang, Sihua Zhang, Xiaobing Dai, Alexandre Capone, Velimir Todorovski, Stefan Sos-
 637 nowski, and Sandra Hirche. Learning-based prescribed-time safety for control of unknown sys-
 638 tems with control barrier functions. *IEEE Control Systems Letters*, 8:1817–1822, 2024.

639 Tzu-Yuan Huang, Armin Lederer, Nicolas Hoischen, Jan Brüdigam, Xuehua Xiao, Stefan Sos-
 640 nowski, and Sandra Hirche. Toward near-globally optimal nonlinear model predictive control
 641 via diffusion models. *Learning for Dynamics and Control*, 2025.

643 Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for
 644 flexible behavior synthesis. *International Conference on Machine Learning*, 2022.

646 Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. Stomp:
 647 Stochastic trajectory optimization for motion planning. In *2011 IEEE international conference*
 on robotics and automation, pp. 4569–4574. IEEE, 2011.

648 Nan Rosemary Ke, Amanpreet Singh, Ahmed Touati, Anirudh Goyal, Yoshua Bengio, Devi Parikh,
 649 and Dhruv Batra. Modeling the long term future in model-based reinforcement learning. In
 650 *International Conference on Learning Representations*, 2019.

651

652 Matthew Kelly. An introduction to trajectory optimization: How to do your own direct collocation.
 653 *SIAM review*, 59(4):849–904, 2017.

654

655 Kota Kondo, Andrea Tagliabue, Xiaoyi Cai, Claudius Tewari, Olivia Garcia, Marcos Espitia-
 656 Alvarez, and Jonathan P How. Cgd: Constraint-guided diffusion policies for uav trajectory plan-
 657 ning. *arXiv preprint arXiv:2405.01758*, 2024.

658

659 Kywoon Lee and Jaesik Choi. Local manifold approximation and projection for manifold-aware
 660 diffusion planning. *arXiv preprint arXiv:2506.00867*, 2025.

661

662 Kywoon Lee, Seongun Kim, and Jaesik Choi. Refining diffusion planner for reliable behavior
 663 synthesis by automatic detection of infeasible plans. *Advances in Neural Information Processing
 664 Systems*, 36:24223–24246, 2023.

665

666 Jinhao Liang, Jacob K Christopher, Sven Koenig, and Ferdinando Fioretto. Simultaneous multi-
 667 robot motion planning with projected diffusion models. *arXiv preprint arXiv:2502.03607*, 2025.

668

669 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 670 for generative modeling. *The Eleventh International Conference on Learning Representations*,
 671 2023.

672

673 Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and
 674 Mark D Plumbley. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv
 675 preprint arXiv:2301.12503*, 2023.

676

677 Kehan Long, Cheng Qian, Jorge Cortés, and Nikolay Atanasov. Learning barrier functions with
 678 memory for robust safe navigation. *IEEE Robotics and Automation Letters*, 6(3):4931–4938,
 679 2021.

680

681 Haofei Lu, Yifei Shen, Dongsheng Li, Junliang Xing, and Dongqi Han. Habitizing diffusion plan-
 682 ning for efficient and effective decision making. *arXiv preprint arXiv:2502.06401*, 2025.

683

684 Hao Ma, Sabrina Bodmer, Andrea Carron, Melanie Zeilinger, and Michael Muehlebach. Constraint-
 685 aware diffusion guidance for robotics: Real-time obstacle avoidance for autonomous racing. *arXiv
 686 preprint arXiv:2505.13131*, 2025.

687

688 François Mazé and Faez Ahmed. Diffusion models beat gans on topology optimization. In *Proceed-
 689 ings of the AAAI conference on artificial intelligence*, volume 37, pp. 9108–9116, 2023.

690

691 Kazuki Mizuta and Karen Leung. Cobl-diffusion: Diffusion-based conditional robot planning in
 692 dynamic environments using control barrier and lyapunov functions. In *2024 IEEE/RSJ Interna-
 693 tional Conference on Intelligent Robots and Systems (IROS)*, pp. 13801–13808. IEEE, 2024.

694

695 Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove.
 696 Deep sdf: Learning continuous signed distance functions for shape representation. In *Proceedings
 697 of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 165–174, 2019.

698

699 Michael Posa, Cecilia Cantu, and Russ Tedrake. A direct method for trajectory optimization of rigid
 700 bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, 2014.

701

702 Ralf Römer, Alexander von Rohr, and Angela P Schoellig. Diffusion predictive control with con-
 703 straints. *Learning for Dynamics and Control*, 2025.

704

705 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar
 706 Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic
 707 text-to-image diffusion models with deep language understanding. *Advances in neural informa-
 708 tion processing systems*, 35:36479–36494, 2022.

709

710 Peiyao Shen, Xuebo Zhang, and Yongchun Fang. Complete and time-optimal path-constrained
 711 trajectory planning with torque and velocity constraints: Theory and applications. *IEEE/ASME
 712 Transactions on Mechatronics*, 23(2):735–746, 2018.

702 Andreas Sochopoulos, Nikolay Malkin, Nikolaos Tsagkas, João Moura, Michael Gienger, and Sethu
 703 Vijayakumar. Fast flow-based visuomotor policies via conditional optimal transport couplings.
 704 *arXiv preprint arXiv:2505.01179*, 2025.

705 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
 706 learning using nonequilibrium thermodynamics. In *International conference on machine learn-
 707 ing*, pp. 2256–2265. pmlr, 2015.

708 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 709 Poole. Score-based generative modeling through stochastic differential equations. *International
 710 Conference on Learning Representations*, 2021.

711 Yongduan Song, Yujuan Wang, John Holloway, and Miroslav Krstic. Time-varying feedback for
 712 regulation of normal-form nonlinear systems in prescribed finite time. *Automatica*, 83:243–251,
 713 2017. ISSN 0005-1098.

714 Yongduan Song, Yujuan Wang, and Miroslav Krstic. Time-varying feedback for stabilization in
 715 prescribed finite time. *International Journal of Robust and Nonlinear Control*, 29(3):618–633,
 716 2019.

717 Yongduan Song, Hefu Ye, and Frank L Lewis. Prescribed-time control and its latest developments.
 718 *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(7):4102–4116, 2023.

719 Eduardo D. Sontag. A lyapunov-like characterization of asymptotic controllability. *SIAM Journal
 720 on Control and Optimization*, 21(3):462–471, 1983.

721 Erik Talvitie. Model regularization for stable sample rollouts. In *UAI*, pp. 780–789, 2014.

722 Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control.
 723 In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033,
 724 2012. doi: 10.1109/IROS.2012.6386109.

725 Kim P Wabersich, Andrew J Taylor, Jason J Choi, Koushil Sreenath, Claire J Tomlin, Aaron D
 726 Ames, and Melanie N Zeilinger. Data-driven safety filters: Hamilton-jacobi reachability, control
 727 barrier functions, and predictive methods for uncertain systems. *IEEE Control Systems Magazine*,
 728 43(5):137–177, 2023.

729 Han Wang, Kostas Margellos, and Antonis Papachristodoulou. Relaxed compatibility between con-
 730 trol barrier and lyapunov functions. In *2024 UKACC 14th International Conference on Control
 731 (CONTROL)*, pp. 125–126, 2024. doi: 10.1109/CONTROL60310.2024.10531901.

732 Yanwei Wang, Lirui Wang, Yilun Du, Balakumar Sundaralingam, Xuning Yang, Yu-Wei Chao,
 733 Claudia Pérez-D'Arpino, Dieter Fox, and Julie Shah. Inference-time policy steering through
 734 human interactions. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*,
 735 pp. 15626–15633. IEEE, 2025.

736 Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, Ramin Hasani, Mathias Lechner, and Daniela Rus.
 737 Safediffuser: Safe planning with diffusion probabilistic models. In *The Fifteenth International
 738 Conference on Learning Representations*, 2025.

739 Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human
 740 motion diffusion model. In *Proceedings of the IEEE/CVF international conference on computer
 741 vision*, pp. 16010–16021, 2023.

742 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided
 743 flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023.

744 Guangyao Zhou, Sivaramakrishnan Swaminathan, Rajkumar Vasudeva Raju, J Swaroop Guntupalli,
 745 Wolfgang Lehrach, Joseph Ortiz, Antoine Dedieu, Miguel Lázaro-Gredilla, and Kevin Murphy.
 746 Diffusion model predictive control. *arXiv preprint arXiv:2410.05364*, 2024.

747

748

749

750

751

752

753

754

755

APPENDIX

Table of Contents

764	A	Signed Distance Functions	15
765	B	Proofs of Theoretical Results	15
766	B.1	Feasibility of CBF constraints	15
767	B.2	Feasibility of CLF Constraints	16
768	B.3	Proof of Theorem 5.1	17
770	C	Additional Experiment Details and Results	17
771	D	Environment details	18
772	E	Constraint Settings in Each Tasks	19
773	F	Training detail	20
774	G	Computational resources	21
775	H	The Use of Large Language Models (LLMs)	21
776	I	Computation Analysis	22
777	J	Additional Related Work	22
778	K	Deployment on Real-World Robotic Platforms	23
779	L	Comparison with constraint-aware baselines	23

A SIGNED DISTANCE FUNCTIONS

Given sets \mathbb{S} and \mathbb{A} , the signed distance functions are defined as the minimal distance to a point on the boundary $\partial\mathbb{S}$ and $\partial\mathbb{A}$ of the sets, with the sign indicating if the state/action along the trajectory is in the corresponding sets. This formally leads to the following functions

$$h_k^s(\tau) = \begin{cases} \min_{s \in \partial\mathbb{S}} \|s(k) - s\| & \text{if } s \in \mathbb{S} \\ -\min_{s \in \partial\mathbb{S}} \|s(k) - s\| & \text{else} \end{cases} \quad h_k^a(\tau) = \begin{cases} \min_{a \in \partial\mathbb{A}} \|a(k) - a\| & \text{if } a \in \mathbb{A} \\ -\min_{a \in \partial\mathbb{A}} \|a(k) - a\| & \text{else.} \end{cases} \quad (6)$$

B PROOFS OF THEORETICAL RESULTS**B.1 FEASIBILITY OF CBF CONSTRAINTS**

Theorem B.1. Assume that $\varphi(t) = \frac{c}{(t-1)^2}$ with constant $c \in \mathbb{R}_+$. If h_k^s and h_k^a are differentiable, then, for each τ , there exists a \mathbf{u} jointly satisfying eq. (CBF-s) and eq. (CBF-a).

Proof. Since the constraints defined in eqs. (CBF-s) and (CBF-a) concern independent variables, individual feasibility of each constraint implies joint feasibility of all of them. Due to the assumed differentiability of $h_k^{s,a}$ and the definition of SDFs implying $\|\nabla h_k^{s,a}(\tau)\| = 1$, \mathbf{u} can always be chosen such $\nabla^T h_k^{s,a}(\tau) \mathbf{u}$ takes an arbitrary value. Thus, the constraints defined in eqs. (CBF-s) and (CBF-a) are always feasible. \square

The assumption on the differentiability of h_k^s and h_k^a is required as the signed distance functions (SDF) eq. (6) are generally not smooth. However, they are differentiable almost everywhere for sets with smooth boundary under weak assumptions (Gilbarg & Trudinger, 1977). Non-differentiable set boundaries can be overcome by increasing the number of CBF constraints via a decomposition of $\bar{\mathbb{S}}$ and $\bar{\mathbb{A}}$ into suitable subsets $\bar{\mathbb{S}}_i$ and $\bar{\mathbb{A}}_i$, such that the SDFs can be computed with respect to the subsets instead. For example, this approach immediately yields linear functions of the form $h_{k,i}^a(\boldsymbol{\tau}) = \bar{a} \pm \mathbf{a}_i(k)$ for individual elements $\mathbf{a}_i(k)$ of actions when defining suitable subsets $\bar{\mathbb{A}}_i$ for the the commonly employed box constraints $\|\mathbf{a}_i(k)\|_\infty \leq \bar{a}$ with some constant $a \in \mathbb{R}_+$. Thus the differentiability assumption in Theorem B.1 is rather technical and often not relevant in practice for the usage of SDFs in CBF constraints (Long et al., 2021).

B.2 FEASIBILITY OF CLF CONSTRAINTS

Theorem B.2. *Assume that $\varphi(t) = \frac{c}{(t-1)^2}$ with constant $c \in \mathbb{R}_+$. If $\{(s, a) : \det(\frac{\partial}{\partial s} \mathbf{f}(s, a)) = 0 \wedge \det(\frac{\partial}{\partial a} \mathbf{f}(s, a)) = \mathbf{0}\} = \emptyset$, then, for each $\boldsymbol{\tau}$, there exists a \mathbf{u} satisfying eq. (CLF).*

Proof. Let

$$\ell_k(\boldsymbol{\tau}) = \left(\boldsymbol{\tau}^{s(k+1)} - \mathbf{f}(\boldsymbol{\tau}^{s(k)}, \boldsymbol{\tau}^{a(k)}) \right). \quad (7)$$

Then, the gradient of V is given by

$$\nabla V(\boldsymbol{\tau}) = \begin{bmatrix} -\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(0)}, \boldsymbol{\tau}^{a(0)})}{\partial \boldsymbol{\tau}^{s(0)}} \ell_0(\boldsymbol{\tau}) \\ -\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(0)}, \boldsymbol{\tau}^{a(0)})}{\partial \boldsymbol{\tau}^{a(0)}} \ell_0(\boldsymbol{\tau}) \\ \ell_0(\boldsymbol{\tau}) - \frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(1)}, \boldsymbol{\tau}^{a(1)})}{\partial \boldsymbol{\tau}^{s(1)}} \ell_1(\boldsymbol{\tau}) \\ -\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(1)}, \boldsymbol{\tau}^{a(1)})}{\partial \boldsymbol{\tau}^{a(1)}} \ell_1(\boldsymbol{\tau}) \\ \vdots \\ \ell_{H-2}(\boldsymbol{\tau}) - \frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(H-1)}, \boldsymbol{\tau}^{a(H-1)})}{\partial \boldsymbol{\tau}^{s(H-1)}} \ell_{H-1}(\boldsymbol{\tau}) \\ -\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(H-1)}, \boldsymbol{\tau}^{a(H-1)})}{\partial \boldsymbol{\tau}^{a(H-1)}} \ell_{H-1}(\boldsymbol{\tau}) \\ \ell_{H-1}(\boldsymbol{\tau}) \\ \mathbf{0} \end{bmatrix}. \quad (8)$$

In the following, we will show by contradiction that $\nabla V(\boldsymbol{\tau}) = \mathbf{0}$ if and only if $\ell_k(\boldsymbol{\tau}) = \mathbf{0}$ for all $k = 0, \dots, H-1$. For this purpose, assume that $\nabla V(\boldsymbol{\tau}) = \mathbf{0}$ and $\ell_i \neq \mathbf{0}$ for some $i = 1, \dots, H-1$. If $\text{rank}(\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(i)}, \boldsymbol{\tau}^{a(i)})}{\partial \boldsymbol{\tau}^{a(i)}}) \neq 0$, $\nabla V(\boldsymbol{\tau}) \neq \mathbf{0}$ is trivially contradicted. If $\text{rank}(\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(i)}, \boldsymbol{\tau}^{a(i)})}{\partial \boldsymbol{\tau}^{s(i)}}) \neq 0$, $\nabla V(\boldsymbol{\tau}) = \mathbf{0}$ requires $\ell_{i-1} \neq \mathbf{0}$. Hence, we can consider the same two cases as for $\ell_i \neq \mathbf{0}$. By repeating this procedure and always considering the case $\text{rank}(\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(k)}, \boldsymbol{\tau}^{a(k)})}{\partial \boldsymbol{\tau}^{s(k)}}) \neq 0$, we eventually end up with the condition

$$\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(0)}, \boldsymbol{\tau}^{a(0)})}{\partial \boldsymbol{\tau}^{s(0)}} \ell_0(\boldsymbol{\tau}) = \mathbf{0} \quad \wedge \quad \text{rank}(\frac{\partial \mathbf{f}(\boldsymbol{\tau}^{s(0)}, \boldsymbol{\tau}^{a(0)})}{\partial \boldsymbol{\tau}^{a(0)}}) \neq 0 \quad \wedge \quad \ell_0 \neq \mathbf{0}, \quad (9)$$

which cannot be satisfied. Thus, $\nabla V(\boldsymbol{\tau}) \neq \mathbf{0}$ holds if $\ell_k(\boldsymbol{\tau}) \neq \mathbf{0}$ for some k . Consequently, there always exists a \mathbf{u} such that eq. (CLF) is satisfied rendering the constraint feasible. \square

To ensure the feasibility of CLF constraints, we again need one technical assumption: Through infinitesimal changes of states or actions, the dynamics \mathbf{f} can be changed in arbitrary directions. If this property is not satisfied, $\nabla V(\boldsymbol{\tau})$ does not necessarily provide information about directions for reducing the violation of eq. (DC). Note that matrices with rank deficiency have zero measure among all matrices, such that infeasibilities related to a violation of the rank condition in Theorem B.2 occur only at isolated states (except for special cases of dynamics, e.g., piecewise constant dynamics). Thus, the rank condition is usually satisfied almost everywhere for many relevant systems, which is sufficient for the feasibility of the CLF constraint in eq. (CLF) in practice.

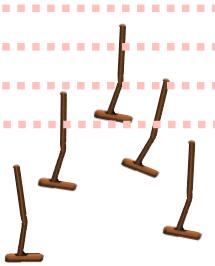


Figure 5: Performance under progressively tightened constraints in the locomotion task. As the allowed torso height decreases, the admissibility constraint becomes stricter, creating a conflict with the task objective of jumping forward.

B.3 PROOF OF THEOREM 5.1

Proof of Theorem 5.1. Due to the assumed feasibility of constraints, eq. (4) controlled by eq. (5) satisfies eq. (CBF-s), eq. (CBF-a), and eq. (CLF), i.e., prescribed-time CBF and CLF conditions are satisfied. As $h_k^{s,a}$ is positive inside \mathbb{S}/\mathbb{A} , negative outside, and zero on the boundaries of these sets, it corresponds to a prescribed-time control barrier function (Huang et al., 2024). Therefore, it immediately follows from (Huang et al., 2024, Theorem 1) that $h_k^{s,a}(\tau_1) > 0$ at $t = 1$, which implies satisfaction of safety (eq. (SC)) and admissibility (eq. (AC)) by construction. Since V is positive definite, it corresponds to a prescribed-time control Lyapunov function (Song et al., 2019). Hence, it immediately follows from (Song et al., 2017, Theorem 2) that $V(\tau_1) = 0$, which implies satisfaction of eq. (DC) by construction. \square

While the two constraints in eq. (CBF-s) and eq. (CBF-a) concern independent variables, the addition of constraint eq. (CLF) introduces a coupling between the constraints. This coupling can cause infeasibility at some trajectories τ in general, but technical conditions exist that exclude them (Wang et al., 2024). Moreover, these infeasibilities are not an issue from a practical perspective since they often affect isolated τ or small subsets of trajectories usually not occurring while numerically integrating eq. (4).

C ADDITIONAL EXPERIMENT DETAILS AND RESULTS

Robustness to Stricter Test-Time Constraints. We evaluate the robustness of SAD-Flower by tightening the test-time constraint on torso height (Fig. 5) in the Hopper environment, decreasing the upper bound from 1.6 to 1.4 and 1.25 settings not encountered during training. As shown in Table 6, our method maintains perfect satisfaction of both safety and admissibility constraints across all levels of difficulty, despite the increasingly restrictive conditions. This demonstrates SAD-Flower’s strong generalization to stricter, unseen constraints at test time. While task rewards decrease modestly—as expected due to the heightened challenge—our method consistently preserves formal guarantees of constraint satisfaction, validating the effectiveness of prescribed-time control as a flexible enforcement mechanism even under distributional shifts.

Comparison with DPCC (Römer et al., 2025). To provide a more comprehensive comparison with optimization-based generative planning, we evaluate SAD-Flower against three DPCC variants (Römer et al., 2025): DPCC-T, which selects samples with minimal deviation from the previous timestep; DPCC-C, which minimizes cumulative projection cost; and DPCC-R, which selects samples randomly after projection. All DPCC variants leverage full-system nonlinear optimization at each denoising step to enforce constraints. As shown in Table 7, while all DPCC variants succeed in satisfying admissibility and achieving dynamic consistency, they exhibit consistently higher safety violations (up to 0.02) even when we apply the constraint tightening technique (using height=1.5 under the constraint with height=1.6) and incur significantly greater computational cost—up to 25× slower than SAD-Flower. These issues likely stem from performing optimization over early-stage noisy samples, which may be difficult or even infeasible to correct without distorting future structure. Furthermore, the lower planning reward across all DPCC variants (0.61–0.65 vs. 0.93) suggest that

Table 6: Performance of the proposed SAD-Flower and baselines depending on constraint tightness. Constraints for the Hopper are tightened via lower admissible heights.

Experiment	Metric	Diffuser	Truncate	CG	FM	S-Diffuser	D-Diffuser	Ours
Hopper (height=1.6)	safety	0.01±0.02	0.05±0.04	0.07±0.03	0.11±0.08	0.05±0.04	0.10±0.02	0.00±0.00
	admissib.	0.21±0.05	0.18±0.04	0.26±0.07	0.17±0.05	0.18±0.04	0.00±0.00	0.00±0.00
	dyn. consist.	0.38±0.03	0.41±0.04	0.79±0.10	0.23±0.02	0.36±0.06	0.16±0.01	0.01±0.01
	reward	1.06±0.18	0.50±0.12	0.73±0.22	1.02±0.20	0.53±0.19	1.12±0.01	0.93±0.23
Hopper (height=1.4)	safety	0.28±0.08	0.12±0.09	0.04±0.03	1.24±0.07	0.12±0.09	0.22±0.03	0.00±0.00
	admissib.	0.21±0.05	0.22±0.09	1.79±0.33	0.19±0.10	0.20±0.08	0.00±0.00	0.00±0.00
	dyn. consist.	0.38±0.03	0.40±0.08	1.06±0.03	0.01±0.01	0.44±0.06	0.13±0.01	0.01±0.01
	reward	1.06±0.18	0.39±0.17	0.32±0.15	1.03±0.17	0.35±0.10	1.11±0.02	0.26 ± 0.02
Hopper (height=1.25)	safety	0.40±0.04	0.00±0.00	0.01±0.01	2.10±0.07	0.00±0.00	0.37±0.02	0.00±0.00
	admissib.	0.21±0.05	0.30±0.05	1.38±0.27	0.05±0.10	0.30±0.05	0.00±0.00	0.00±0.00
	dyn. consist.	0.38±0.03	0.46±0.21	1.79±0.08	0.01±0.01	0.46±0.21	0.15±0.01	0.01±0.01
	reward	1.06±0.18	0.13±0.02	0.10±0.05	1.02±0.16	0.13±0.02	0.78±0.02	0.17±0.00

Table 7: The performance of the proposed SAD-Flower with QP and DPCC (Römer et al., 2025) with nonlinear optimization. Constraint for the Hopper is the same as the admissible heights in Table 1. Three DPCCs are proposed in their work. DPCC-T (Temporal consistency) selects the trajectory that deviates the least from the previous timestep, DPCC-C (Cumulative projection cost) selects the trajectory that has been modified the least by the projection operation, and DPCC-R (Random) selects the trajectory randomly.

Experiment	Metric	DPCC-T	DPCC-C	DPCC-R	Ours
Hopper (Med-Expert)	safety	0.01±0.03	0.02±0.02	0.02±0.02	0.00±0.00
	admissib.	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00
	dyn. consist.	0.01±0.01	0.01±0.00	0.01±0.00	0.01±0.01
	reward	0.61±0.07	0.65±0.16	0.64±0.25	0.93±0.23
	time	4.24±1.64	4.56±1.72	6.21±2.61	0.06±0.00

early and aggressive control leads to sample deviation from the distribution learned by the generative model. In contrast, SAD-Flower activates control later using a lightweight QP-based formulation and preserves both safety and performance, highlighting the advantages of flexible, prescribed-time constraint enforcement.

D ENVIRONMENT DETAILS

We evaluate our method across a variety of trajectory planning domains, summarized in Table 8. The benchmark includes navigation (Maze2d-Umaze, Maze2d-Large), locomotion (Hopper, Walker2d), and robotic manipulation (Kuka Block-Stacking). Environments are simulated using MuJoCo (Todorov et al., 2012) or PyBullet (Coumans & Bai, 2016), and cover a range of state and action dimensions.

We use offline datasets to train the generative and control models for each benchmark. As shown in Table 9, the datasets for locomotion and navigation tasks are retrieved from the D4RL benchmark suite (Fu et al., 2020), with varying sizes (e.g., medium vs. medium-expert settings). For locomotion environments, datasets are collected using soft actor-critic (SAC) policies (Haarnoja et al., 2018), either partially trained or mixed with expert-level demonstrations. For the Kuka block-stacking task (Janner et al., 2022), data is generated via the PDDLStream planner (Garrett et al., 2020), which provides feasible robotic manipulation trajectories in structured stacking scenarios.

972
973
974
975
976
977
978
979
Table 8: Settings for each tasks.
980

Environment	Simulator	Obs. Dim.	Action Dim.
Maze2d-Umaze-v1	MuJoCo	4	2
Maze2d-Large-v1	MuJoCo	4	2
Hopper	MuJoCo	11	3
Walker2d	MuJoCo	17	6
Kuka Block-Stacking	PyBullet	39	-

980
981
982
Table 9: Dataset details for each benchmark environment, including the number of trajectories and
the source or algorithm used to generate the data.
983

Environment	# of Trajectories	Source / Generation Method
Maze2d-Umaze-v1	10^6	D4RL (Fu et al., 2020)
Maze2d-Large-v1	4×10^6	D4RL (Fu et al., 2020)
Hopper-medium	10^6	Partially trained SAC (Haarnoja et al., 2018)
Hopper-medium-expert	2×10^6	Mixture of expert and partial SAC
Walker2d-medium	10^6	Partially trained SAC (Haarnoja et al., 2018)
Walker2d-medium-expert	2×10^6	Mixture of expert and partial SAC
Kuka Block-stacking	10,000	PDDLStream planner (Garrett et al., 2020)

991
992
993
994
995
996
997
E CONSTRAINT SETTINGS IN EACH TASKS

We introduce task-specific state and action constraints to evaluate the ability of generative planners to handle safety and admissibility under diverse and challenging conditions. The test-time constraints are deliberately chosen to be more restrictive and often unseen during training, making constraint satisfaction a non-trivial requirement.

Maze2d. The goal in Maze2d environments is to generate a feasible trajectory from a randomly sampled initial position to a randomly sampled goal location. The generative model is conditioned on both endpoints and produces full trajectories through the maze.

To evaluate constraint satisfaction, we introduce two novel, unseen obstacles at test time that do not completely block the path but add non-trivial planning constraints. The first is a superellipse-shaped obstacle defined as:

$$\left(\frac{x - x_0}{a}\right)^2 + \left(\frac{y - y_0}{b}\right)^2 \geq 1, \quad (10)$$

where $(x, y) \in \mathbb{R}^2$ is a trajectory state and $(x_0, y_0) \in \mathbb{R}^2$ is the center of the obstacle. The parameters $a > 0$ and $b > 0$ control the width and height of the obstacle.

The second is a higher-order polynomial barrier defined as:

$$\left(\frac{x - x_0}{a}\right)^4 + \left(\frac{y - y_0}{b}\right)^4 \geq 1. \quad (11)$$

This formulation results in sharper obstacle boundaries and makes naive post-processing or trajectory truncation ineffective due to the nonlinearity of the feasible region.

The action constraint is also tightened during test time by imposing a box constraint on the control inputs:

$$\mathbf{a} \in [-0.1, 0.1]^2. \quad (12)$$

Both Maze2d-Umaze-v1 and Maze2d-Large-v1 contain these two novel obstacles at test time to assess generalization to unseen constraints.

Hopper and Walker2d. In the locomotion tasks, we impose test-time constraints that limit the robot's vertical motion to avoid collisions with overhead obstacles. Specifically, the height of the robot's torso must remain below a fixed roof height z , defined as:

$$s < z, \quad (13)$$

1026 where s is the vertical position of the torso. We evaluate this constraint under increasingly restrictive
 1027 settings, such as $z = 1.6, 1.4$, and 1.25 , to assess robustness to constraint tightening.
 1028

1029 This state constraint is often in conflict with the task objective, which rewards forward jumping
 1030 or walking. As a result, satisfying the constraint typically requires sacrificing task performance.
 1031 Additionally, the control inputs are constrained by an action box constraint:

$$1032 \quad \mathbf{a} \in [-1, 1]^d, \quad (14)$$

1033 where d is the dimensionality of the action space (3 for Hopper and 6 for Walker2d).
 1034

1035 **Kuka Block-Stacking.** For the Kuka block-stacking task, the generative model is conditioned on
 1036 object locations and outputs joint trajectories for manipulation. To simulate partial system degra-
 1037 dation or workspace reconfiguration, we apply a tighter state constraint on the robot’s joint positions.
 1038 Specifically, we scale the original feasible joint limits by a factor of 0.9:

$$1039 \quad \mathbf{q} \in 0.9 \cdot [\mathbf{q}_{\min}, \mathbf{q}_{\max}], \quad (15)$$

1040 where \mathbf{q}_{\min} and \mathbf{q}_{\max} represent the original lower and upper bounds of each joint. No explicit action
 1041 constraint is applied in this task.

1042 These constraints, particularly when unseen during training, pose a significant challenge for plan-
 1043 ning and provide a rigorous benchmark for evaluating constraint satisfaction and generalization
 1044 capabilities.
 1045

1046 F TRAINING DETAIL

1047 We provide implementation details of all baseline methods and our proposed approach, including
 1048 training configurations, hyperparameters, and model architectures.
 1049

1050 All baseline methods are trained using their official or publicly available codebases: Diffuser follows
 1051 the implementation of Janner et al. (2022), Truncation, Classifier Guidance, and SafeDiffuser are
 1052 implemented based on Xiao et al. (2025), Decision Diffuser follows Ajay et al. (2023), and Flow
 1053 Matching uses the codebase from Feng et al. (2025). Our method is built on top of the same Flow
 1054 Matching framework from Feng et al. (2025).
 1055

1056 **Hyperparameters.** Table 10 shows the training hyperparameters for all baselines and our method
 1057 across different environments. All methods use the same horizon and batch size per environment.
 1058 Note that our method shares training settings with the flow matching baseline to ensure a fair com-
 1059 parison.
 1060

1061 **Forward Dynamics Model Training.** To support constraint enforcement in SAD-Flower, we train
 1062 a forward dynamics model for each environment using the same dataset used to train the generative
 1063 models. This ensures a fair comparison without introducing additional supervision. The training
 1064 details are summarized in Table 11.

1065 For the Maze2d environments, a known analytical dynamics model can also be derived:

$$1066 \quad \begin{bmatrix} x(k+1) \\ y(k+1) \\ v_x(k+1) \\ v_y(k+1) \end{bmatrix} = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x(k) \\ y(k) \\ v_x(k) \\ v_y(k) \end{bmatrix} + \begin{bmatrix} 0.5\alpha dt^2 & 0 \\ 0 & 0.5\alpha dt^2 \\ \alpha dt & 0 \\ 0 & \alpha dt \end{bmatrix} \begin{bmatrix} u_x(k) \\ u_y(k) \end{bmatrix} \quad (16)$$

1067 where $[x_k, y_k, v_{x,k}, v_{y,k}]^T$ is the system state representing position and velocity, $[u_{x,k}, u_{y,k}]^T$ is the
 1068 input force, dt is the simulation time step, and α is the gear ratio divided by mass (due to primitive
 1069 joint control).
 1070

1071 **Model Architecture.** Diffusion-based models (Diffuser, Truncation, Classifier Guidance, SafeD-
 1072 iffuser) use the U-Net architecture with residual temporal convolutions, group normalization, and
 1073 Mish nonlinearities, as described in (Janner et al., 2022). Flow Matching and SAD-Flower use the
 1074 Transformer-based backbone proposed by (Feng et al., 2025), consisting of 8 layers with a hidden
 1075 dimension of 256. The forward model in SAD-Flower is a feedforward neural network with 3 hid-
 1076 den layers of size 512, where the input dimension is the sum of the observation dimension and the
 1077 action dimension, and the output is the observation dimension.
 1078

Table 10: Training hyperparameters for each method across environments.

Maze2d-Umaze	Diffuser	Truncation	CG	S-Diffuser	FM	SAD-Flower
Batch Size	32	32	32	32	32	32
Learning Rate	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$
Steps	$2e^6$	$2e^6$	$2e^6$	$1e^6$	$1e^6$	$1e^6$
Maze2d-Large	Diffuser	Truncation	CG	S-Diffuser	FM	SAD-Flower
Batch Size	32	32	32	32	32	32
Learning Rate	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$
Steps	$2e^6$	$2e^6$	$2e^6$	$1e^6$	$1e^6$	$1e^6$
Hopper	Diffuser	Truncation	CG	S-Diffuser	FM	SAD-Flower
Batch Size	32	32	32	32	32	32
Learning Rate	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$
Steps	$2e^6$	$2e^6$	$2e^6$	$1e^6$	$1e^6$	$1e^6$
Walker2d	Diffuser	Truncation	CG	S-Diffuser	FM	SAD-Flower
Batch Size	32	32	32	32	32	32
Learning Rate	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$
Steps	$2e^6$	$2e^6$	$2e^6$	$1e^6$	$1e^6$	$1e^6$
Kuka Block-Stacking	Diffuser	Truncation	CG	S-Diffuser	FM	SAD-Flower
Batch Size	32	32	32	32	32	32
Learning Rate	$2e^{-5}$	$2e^{-5}$	$2e^{-5}$	$2e^{-5}$	$2e^{-4}$	$2e^{-4}$
Steps	$7e^5$	$7e^5$	$7e^5$	$7e^5$	$7e^5$	$7e^5$

Table 11: Training settings for the forward dynamics models used in SAD-Flower.

Environment	Batch Size	Learning Rate	Steps
Maze2d (Umaze/Large)	256	1×10^{-3}	1×10^6
Hopper	256	1×10^{-3}	1×10^6
Walker2d	256	1×10^{-3}	1×10^6

G COMPUTATIONAL RESOURCES

All experiments were conducted using four high-performance workstations with identical or near-identical configurations. Each workstation is equipped with an AMD EPYC series CPU, an NVIDIA Tesla P100 GPU, and 16 GB of GPU memory. The specific hardware details are summarized in Table 12.

Table 12: Hardware specifications of workstations used for training and evaluation.

Workstation	CPU	GPU	GPU RAM
1	AMD EPYC 7542	NVIDIA Tesla P100	16 GB
2	AMD EPYC 7542	NVIDIA Tesla P100	16 GB
3	AMD EPYC 7742	NVIDIA Tesla P100	16 GB
4	AMD EPYC 7542	NVIDIA Tesla P100	16 GB

H THE USE OF LARGE LANGUAGE MODELS (LLMs)

Large language models (LLMs), e.g., ChatGPT, were employed to assist with editing, rephrasing, and LaTeX formatting during the paper writing process. No LLMs were involved in formulating research ideas, theoretical developments, or core experimental designs.

1134 I COMPUTATION ANALYSIS

1135
 1136 We report the detailed computational costs of SAD-Flower and all baseline methods across our
 1137 benchmark tasks in Table 13. As expected, SAD-Flower incurs greater computational cost compared
 1138 to unconstrained generative planners such as Diffuser and FM, due to the overhead introduced by
 1139 enforcing multiple constraints during sampling. Nonetheless, in some domains—such as Hopper-
 1140 Medium-Expert and Kuka Block-Stacking—SAD-Flower achieves lower runtime than SafeDiffuser,
 1141 highlighting the efficiency gains enabled by our prescribed-time control formulation. Importantly,
 1142 SAD-Flower is the only method that consistently satisfies all three constraint classes—state, action,
 1143 and dynamic consistency—across all evaluated domains. No faster baseline provides this level of
 1144 safety and reliability.

1145 Table 13: Computation analysis of the proposed SAD-Flower and baselines across navigation, loco-
 1146 motion, and manipulation tasks.

Experiment	Metric	Diffuser	Trunc	CG	FM	S-Diffuser	D-Diffuser	Ours
Maze2d (Large)	comp. time	0.01±0.01	0.02±0.01	0.04±0.01	0.06±0.02	0.06±0.01	0.01±0.01	0.13±0.01
Maze2d (Umaze)	comp. time	0.01±0.01	—	0.04±0.01	0.02±0.01	0.06±0.02	0.01±0.01	0.08±0.02
Hopper (Med-Expert)	comp. time	0.05±0.01	0.06±0.02	0.06±0.01	0.01±0.01	0.11±0.02	0.02±0.01	0.06±0.01
Hopper (Medium)	comp. time	0.05±0.01	0.06±0.01	0.06±0.02	0.01±0.01	0.09±0.02	0.02±0.01	0.12±0.02
Walker2D (Med-Expert)	comp. time	0.06±0.02	0.06±0.01	0.05±0.01	0.01±0.01	0.09±0.02	0.02±0.01	0.09±0.01
Walker2D (Medium)	comp. time	0.04±0.01	0.13±0.03	0.21±0.02	0.01±0.01	0.10±0.02	0.02±0.01	0.09±0.01
KUKA Block Stacking	comp. time	0.70±0.01	0.84±0.01	0.86±0.01	0.47±0.06	0.78±0.01	0.04±0.01	0.74±0.08

1157 J ADDITIONAL RELATED WORK

1158
 1159 **Diffusion planning with reliability and constraint enforcement.** Recent works have sought to
 1160 improve the reliability of diffusion-based planners by modifying sampling strategies or adding post-
 1161 hoc corrections, without addressing dynamic consistency or unseen constraints. Lee *et al.* propose
 1162 *Refining Diffusion Planner for Reliable Behavior Synthesis* and introduce a restoration-gap metric
 1163 predicted by a gap predictor to identify error-prone plans (Lee et al., 2023). The restoration gap acts
 1164 as refining guidance, and an attribution-map regularizer prevents adversarial guidance, improving
 1165 feasibility without altering the core diffusion model.

1166 Feng *et al.* address stochastic failure modes by aggregating multiple trajectory samples. Their
 1167 *Trajectory Aggregation Tree* (TAT) combines historical and current trajectories into a tree struc-
 1168 ture where each branch corresponds to a trajectory and nodes correspond to individual states (Feng
 1169 et al., 2024). Unreliable states are marginalized and the most impactful nodes prioritized, yielding a
 1170 training-free module deployable without modifying the original diffusion planner.

1171 Wang *et al.* examine human-in-the-loop alignment. Their *Inference-Time Policy Steering* (ITPS)
 1172 framework injects human interactions into the sampling process of a pretrained policy rather than
 1173 fine-tuning on new data Wang et al. (2025). ITPS evaluates various forms of human interaction
 1174 and sampling strategies to bias the denoising process toward user-desired subgoals while mitigating
 1175 distribution shift.

1176 LoMAP, introduced by Lee and Choi, targets off-manifold drift in guidance. The authors derive a
 1177 lower bound on the guidance gap that quantifies manifold deviation and propose a training-free *Lo-
 1178 cal Manifold Approximation and Projection* method Lee & Choi (2025). By projecting the guided
 1179 sample onto a low-rank subspace approximated from offline data, LoMAP prevents infeasible tra-
 1180 jectory generation and can be incorporated as a plug-in module for hierarchical diffusion planners.

1181 Finally, Liang *et al.* extend diffusion planning to multi-robot settings. Their *Simultaneous Multi-
 1182 Robot Motion Planning with Projected Diffusion Models* integrates constrained optimization into
 1183 every denoising step, producing collision-free and kinematically feasible trajectories in dense multi-
 1184 robot environments Liang et al. (2025). The authors introduce a comprehensive benchmark and
 1185 report higher success rates compared with classical and learning-based planners.

1186 These methods improve reliability by detecting artifacts, aggregating multiple candidates, steer-
 1187 ing policies with human input, projecting onto data manifolds, or enforcing per-step feasibility.

1188 However, none simultaneously enforce safety, admissibility, and dynamic consistency, nor guarantee
 1189 adherence to constraints absent from training data. By contrast, *SAD-Flower* formulates diffusion
 1190 sampling as a controlled dynamical system. A quadratic-programming controller derived from
 1191 control-barrier and control-Lyapunov conditions injects guidance at each denoising step, ensuring
 1192 constraint satisfaction and dynamic consistency throughout the entire sampling horizon.
 1193

1194 K DEPLOYMENT ON REAL-WORLD ROBOTIC PLATFORMS

1196 Recent works have demonstrated the successful deployment of diffusion- and flow-based models on
 1197 physical robotic systems (Chi et al., 2023; Sochopoulos et al., 2025), highlighting the feasibility of
 1198 integrating such models into real-world robotic pipelines that involve perception, state estimation,
 1199 and control. These results indicate that generative trajectory models can operate effectively outside
 1200 of simulation.

1201 *SAD-Flower* is designed as an open-loop planner that generates full trajectories prior to execution,
 1202 without requiring real-time feedback control. This formulation avoids the need to meet strict control-
 1203 frequency constraints, making it amenable to practical deployment. Although *SAD-Flower* solves
 1204 a small quadratic program (QP) at each sampling step to enforce constraints, this overhead remains
 1205 moderate and does not hinder offline trajectory generation. Moreover, recent advances in efficient
 1206 generative planning—such as DiffuserLite (Dong et al., 2024) and Habi (Lu et al., 2025)—can be
 1207 directly incorporated into our framework to further accelerate sampling, enabling tighter planning
 1208 budgets without sacrificing constraint guarantees.

1209 While our current implementation is focused on simulation, the modular structure of *SAD-Flower*
 1210 makes it compatible with standard deployment pipelines. The forward model used for dynamics
 1211 consistency can be learned directly from onboard sensor data, and constraint formulations (e.g.,
 1212 actuator limits, workspace boundaries) can be adapted to real-world specifications. As with all
 1213 generative models, care must be taken to address sensor noise and execution uncertainty, but our
 1214 control-theoretic formulation inherently provides robustness by enforcing constraints throughout
 1215 the trajectory.

1216 In summary, *SAD-Flower* introduces only modest additional computation to existing generative
 1217 planners while providing formal guarantees for safety, admissibility, and dynamic consistency.
 1218 Given the demonstrated deployability of similar frameworks and the compatibility of our method
 1219 with efficient sampling and standard control pipelines, we believe *SAD-Flower* can be incorporated
 1220 into real-world robotic systems with limited adaptation. Investigating such deployment remains an
 1221 important direction for future work.

1223 L COMPARISON WITH CONSTRAINT-AWARE BASELINES

1225 We further compare *SAD-Flower* against two additional constraint-aware baselines: a reject-
 1226 sampling method, which offers a simple heuristic for constraint handling, and CoBL-Diffusion
 1227 (Mizuta & Leung, 2024), which leverages control-theoretic rewards to encourage constraint sat-
 1228 isfaction.

1229 The reject-sampling approach first trains a flow matching model on the original dataset, which in-
 1230 cludes trajectories that violate test-time constraints, since those constraints are unseen during train-
 1231 ing. At inference, it samples a batch of trajectory candidates and discards those violating constraints,
 1232 selecting from the remaining set the trajectory with the least violation. As shown in Table 14, this ap-
 1233 proach consistently fails to satisfy constraints across all tasks, including Maze2d, Hopper, Walker2d,
 1234 and Kuka Block-Stacking. Both safety and admissibility violations remain non-zero. This con-
 1235 firms that post-hoc filtering alone is insufficient when the generative model has no knowledge of
 1236 constraint structure—especially under novel test-time constraints. Even though flow matching can
 1237 model multimodal behavior patterns, it cannot reliably sample within unseen constraint sets unless
 1238 such constraints are explicitly incorporated during inference.

1239 CoBL-Diffusion injects physical constraints into a diffusion model by shaping its denoising pro-
 1240 cess with auxiliary rewards derived from Control Barrier Functions (CBFs) and Control Lyapunov
 1241 Functions (CLFs). This allows it to softly bias generation toward constraint-adherent behaviors, but
 without any formal guarantees. We evaluate CoBL-Diffusion in the LargeMaze environment, using

the same horizon and constraints as in other baselines. As shown in Table 15, the method fails to fully enforce safety and admissibility constraints. This is expected given several key limitations: (1) it predicts only actions, whereas SAD-Flower jointly generates states and actions, which allows for more precise control of trajectory behavior; (2) its CBFs target only state constraints, while ours apply to both state and action spaces; (3) it uses CLFs for general stability, not for ensuring formal dynamic consistency as we do; and (4) it lacks prescribed-time scheduling, making it difficult to guarantee constraint satisfaction at the final step. These design choices limit CoBL-Diffusion’s ability to robustly handle complex, temporally structured constraints. In contrast, SAD-Flower provides a framework to reliably enforce constraints even under novel test-time scenarios.

Table 14: Performance of SAD-Flower and reject-sample method across all tasks.

Experiment	Metric	Reject-sample	Ours
Maze2d (Large)	safety	0.16 ± 0.29	0.00 ± 0.00
	admissib.	0.90 ± 0.01	0.00 ± 0.00
	dyn. consist.	0.02 ± 0.01	0.01 ± 0.01
	reward	1.52 ± 0.28	1.42 ± 0.52
Maze2d (Umaze)	safety	0.04 ± 0.09	0.00 ± 0.00
	admissib.	0.91 ± 0.01	0.00 ± 0.00
	dyn. consist.	0.01 ± 0.01	0.01 ± 0.01
	reward	2.91 ± 0.65	2.66 ± 0.88
Hopper (Med-Expert)	safety	0.14 ± 0.18	0.00 ± 0.00
	admissib.	0.05 ± 0.02	0.00 ± 0.00
	dyn. consist.	0.02 ± 0.01	0.01 ± 0.01
	reward	0.49 ± 0.33	0.93 ± 0.23
Hopper (Medium)	safety	0.03 ± 0.04	0.00 ± 0.00
	admissib.	0.06 ± 0.03	0.00 ± 0.00
	dyn. consist.	0.02 ± 0.02	0.01 ± 0.01
	reward	0.24 ± 0.14	0.34 ± 0.03
Walker2d (Med-Expert)	safety	0.08 ± 0.10	0.00 ± 0.00
	admissib.	0.08 ± 0.08	0.00 ± 0.00
	dyn. consist.	0.04 ± 0.06	0.04 ± 0.04
	reward	0.97 ± 0.22	0.89 ± 0.32
Walker2d (Medium)	safety	0.18 ± 0.15	0.00 ± 0.00
	admissib.	0.16 ± 0.10	0.00 ± 0.00
	dyn. consist.	0.07 ± 0.17	0.07 ± 0.15
	reward	0.68 ± 0.19	0.42 ± 0.23
KUKA Block Stacking	safety	0.01 ± 0.01	0.00 ± 0.00
	reward	0.44 ± 0.67	0.45 ± 0.21

Table 15: Performance of SAD-Flower and CoBL-Diffusion in the navigation task.

Experiment	Metric	CoBL-Diffusion	Ours
Maze2d (Large)	safety	0.01 ± 0.04	0.00 ± 0.00
	admissib.	0.23 ± 0.33	0.00 ± 0.00
	dyn. consist.	0.00 ± 0.00	0.01 ± 0.01
	reward	0.15 ± 0.18	1.42 ± 0.52