Aerial Inspection of Marine Infrastructures via Diffusion-Driven Safe Model Predictive Control

Erdi Sayar*, Van Huyen Dang*, Ersin Daş†, and Erdal Kayacan*
*Paderborn University, Germany. Emails: {erdi.sayar, van.huyen.dang, erdal.kayacan}@uni-paderborn.de
†Illinois Institute of Technology, USA. Email: edas2@illinoistech.edu

***Abstract*—Autonomous aerial inspection of marine infrastructures requires navigating complex, safety-critical environments where strict structural avoidance is paramount. While diffusion models excel at generating high-quality multi-modal trajectories, their inability to natively enforce hard safety and dynamical bounds often results in physically infeasible outputs. Hybrid approaches coupling diffusion with model predictive control (MPC) address this by projecting trajectories onto feasible sets; however, poor initializations in constrained inspection spaces frequently cause MPC failure. To overcome these challenges, we propose D-SafeMPC (*Diffusion-Driven Safe MPC*), a novel trajectory planning framework tailored for structural inspections. Our method guides the reverse diffusion process using discrete-time control barrier functions (CBFs) and control Lyapunov functions (CLFs), providing the MPC with probabilistically safe, well-initialized trajectories. An iterative-projection scheme rigorously refines the trajectory at each denoising step, ensuring strict collision avoidance and dynamic feasibility. In simulations on a drone across four highly constrained scenarios, D-SafeMPC demonstrates superior safety, task success rates, and real-time planning efficiency over state-of-the-art baselines, making it highly suitable for safety-critical marine inspections.**

I. Introduction

Autonomous aerial inspection of marine infrastructures, such as offshore wind turbines and oil rigs, requires highly reliable and agile trajectory planning. In these complex scenarios, unmanned aerial vehicles (UAVs) must navigate tightly constrained, dynamic environments while strictly avoiding structural collisions. Diffusion models [1] have recently demonstrated impressive capabilities in learning multimodal robotic policies from demonstrations [2]–[5]. However, their stochastic formulation lacks an intrinsic mechanism to enforce explicit hard constraints, such as structural avoidance or precise safety margins [6], frequently resulting in trajectories that violate physical system limits [7].

While some methods use control barrier functions (CBFs) [8] and control Lyapunov functions (CLFs) [9] to guide diffusion models away from obstacles [10], this provides only soft guidance, rendering trajectories probabilistically safer but lacking hard guarantees. Alternatively, integrating diffusion models with model predictive control (MPC) [11], which is widely used for complex dynamics [12], [13], can enforce strict safety by projecting trajectories onto feasible sets [14], [15]. However, in the highly constrained spaces typical of marine inspections, a poor initial trajectory generated by the diffusion model can prevent the MPC from converging, leading to critical failures.

To address these challenges in safety-critical inspection domains, we propose D-SafeMPC. Our framework seamlessly integrates diffusion-based trajectory generation with CBF/CLF guidance and MPC (Fig. 1). In each denoising iteration, discrete-time CBFs and CLFs steer the generated samples toward safe, goal-directed regions. This provides the MPC with well-initialized, probabilistically safe trajectories, ensuring highly stable and efficient convergence to strictly safe operational boundaries [14] within real-time computational budgets. The structurally refined trajectories from the MPC are then used to initialize the subsequent denoising iteration.

II. Trajectory Planning with Conditional Diffusion models

We address trajectory planning for a robot (e.g., a quadrotor drone inspecting marine infrastructure, modeled here as a velocity-controlled point mass) navigating environments with static or dynamic obstacles. The discrete-time system dynamics are denoted as $\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k)$. Our objective is generating a sequence of control inputs $\mathbf{u}_{0:T}$ that guides the robot from $\mathbf{x}_0$ to a goal $g \in \mathcal{G}$, ensuring collision avoidance at all times. To do this, we combine offline-trained conditional diffusion models $p_\theta(\tau \mid y)$ (where y is conditioning information like current state and goal) with control constraints.

A. CBF and CLF Guidance in Diffusion

Diffusion models learn unconstrained dynamics from offline data, but during aerial inspection, encountering unseen structural obstacles requires online safety enforcement. We construct a unified safety function $h(\bar{\mathbf{x}})$ mapping the minimum distance to all obstacles into a smooth scalar value (e.g., via a Log-Sum-Exp under-approximation). A discrete-time CBF ensures forward invariance (safety) if $\psi_h(\bar{\mathbf{x}}_k, \mathbf{u}_k) \geq 0$. Similarly, a CLF $V(\bar{\mathbf{x}}) = \|\bar{\mathbf{x}} - g\|^2$ ensures goal convergence if $\Phi_V(\bar{\mathbf{x}}_k, \mathbf{u}_k) \leq 0$.

Rather than retraining the model for every new marine environment, we apply classifier guidance to dynamically steer the reverse diffusion process toward probabilistic safety. We approximate the intractable safety-conditioned posterior $p(\tilde{\tau}^{i-1} \,|\, \tilde{\tau}^i, i, y, \mathcal{C})$ by shifting the Gaussian mean of the learned diffusion model using the gradients of our CBF and CLF conditions with respect to the control input $\mathbf{u}$. The guided denoising step becomes:

$$p_\theta(\tilde{\tau}^{i-1} \mid \tilde{\tau}^i, y) = \mathcal{N}(\mu_\theta + \lambda_c \nabla_{\mathbf{u}} \psi_h - \lambda_l \nabla_{\mathbf{u}} \Phi_V, \beta_t \mathbf{I}), \quad (1)$$

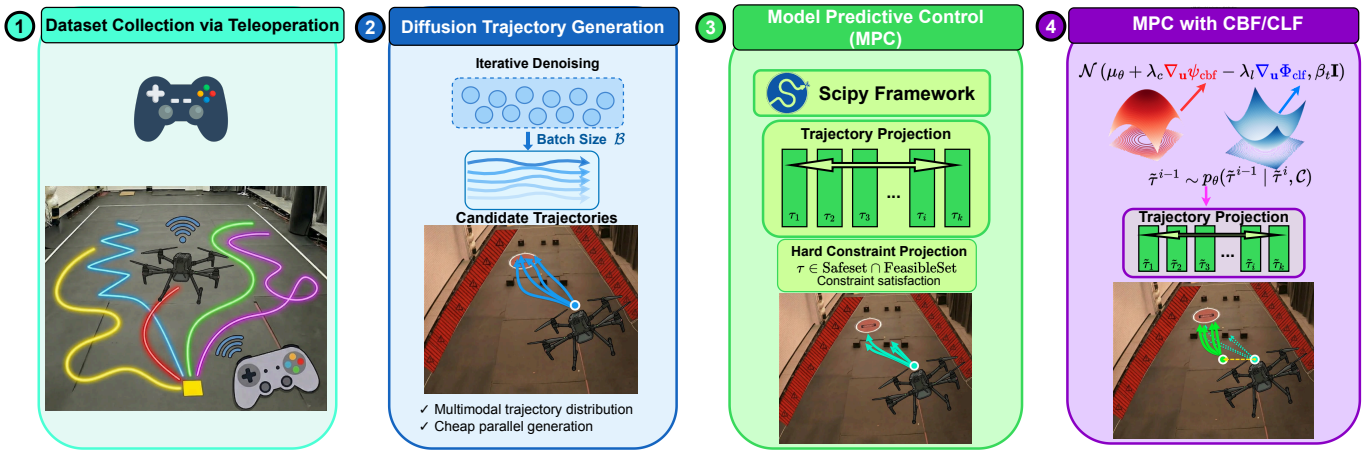

Fig. 1. Overview of D-SafeMPC. The four-stage pipeline. **(1) Dataset Collection:** Expert trajectories are collected by teleoperating a drone via joystick and used to train the conditional diffusion model offline. **(2) Diffusion Inference:** Given the current drone state and goal, the trained diffusion model generates a batch of $\mathcal{B}$ candidate trajectories over an H-step horizon (blue arrows). **(3) MPC without CBF/CLF Guidance:** Each candidate trajectory is used to initialize an MPC (scipy-based). When the diffusion model's initial trajectory lies close to obstacles, the MPC may fail to converge to a feasible solution within its computational budget, resulting in trajectories that stagnate before obstacles (solid blue arrows). **(4) MPC with CBF/CLF-Guided Diffusion:** Discrete-time CBFs and CLFs steer the reverse diffusion process, shifting the initial states away from obstacles (green circle) and providing the MPC with well-initialized, probabilistically safe trajectories. With this improved warm start, the MPC converges efficiently to collision-free, dynamically feasible solutions (green arrows). For comparison, dashed blue arrows depict the unguided MPC trajectories that remain stuck before obstacles, whereas the guided green trajectories successfully navigate around them.

TABLE I. Planning performance across static and dynamic environments. Mean ± std of *Safety-Compliant Goal Reached Rate* (SCGR) and *Goal Reached Rate* (GR). **Best** and 2nd best per column are highlighted. † Tightened constraint variant. ↑ higher is better.

	Static Env		Dynamic Env 1		Dynamic Env 2		Dynamic Env 3	
Planner	**SCGR ↑**	**GR ↑**	**SCGR ↑**	**GR ↑**	**SCGR ↑**	**GR ↑**	**SCGR ↑**	**GR ↑**
D-SafeMPC (ours)	**0.80** ± .33	**0.80** ± .33	**1.00** ± .00	**1.00** ± .00	**0.88** ± .27	**0.88** ± .27	**0.93** ± .18	**0.93** ± .18
DPCC-C	0.90 ± .30	0.90 ± .30	0.75 ± .34	0.85 ± .23	0.45 ± .42	0.60 ± .34	0.43 ± .33	0.43 ± .33
DPCC-C Tight.†	0.70 ± .40	0.70 ± .40	0.70 ± .33	0.75 ± .25	0.58 ± .40	0.58 ± .40	0.65 ± .39	0.65 ± .39
DPCC-R	0.35 ± .39	0.50 ± .39	0.45 ± .35	0.65 ± .39	0.28 ± .30	0.33 ± .33	0.68 ± .29	0.83 ± .24
DPCC-R Tight.†	0.70 ± .33	0.70 ± .33	0.80 ± .33	0.90 ± .30	0.60 ± .34	0.60 ± .34	0.75 ± .25	0.75 ± .25
DPCC-T	0.80 ± .33	0.95 ± .15	0.35 ± .32	0.50 ± .32	0.05 ± .15	0.05 ± .15	0.28 ± .33	0.30 ± .37
DPCC-T Tight.†	0.85 ± .32	0.85 ± .32	0.45 ± .35	0.45 ± .35	0.13 ± .22	0.13 ± .22	0.35 ± .42	0.35 ± .42
CoBL	0.05 ± .15	**1.00** ± .00	0.10 ± .30	0.85 ± .32	0.78 ± .37	0.93 ± .24	0.80 ± .33	0.85 ± .32
Diffuser	0.05 ± .15	0.65 ± .45	0.20 ± .33	0.40 ± .44	0.18 ± .33	0.33 ± .43	0.33 ± .43	0.45 ± .42
Guidance	0.40 ± .49	0.40 ± .49	0.40 ± .49	0.40 ± .49	0.53 ± .49	0.53 ± .49	0.53 ± .49	0.53 ± .49
Guidance Tight.†	0.40 ± .49	0.40 ± .49	0.40 ± .49	0.40 ± .49	0.53 ± .49	0.53 ± .49	0.53 ± .49	0.53 ± .49
Model-Free	0.00 ± .00	0.55 ± .42	0.35 ± .45	0.45 ± .47	0.23 ± .37	0.30 ± .40	0.35 ± .42	0.45 ± .42
Model-Free Tight.†	0.00 ± .00	0.60 ± .37	0.35 ± .45	0.50 ± .45	0.28 ± .40	0.43 ± .40	0.38 ± .42	0.43 ± .43
Post-Proc.	0.55 ± .35	0.75 ± .34	0.55 ± .35	0.60 ± .30	0.40 ± .37	0.43 ± .40	0.55 ± .35	0.73 ± .30
Post-Proc. Tight.†	0.70 ± .33	0.70 ± .33	0.75 ± .40	0.75 ± .40	0.60 ± .30	0.60 ± .30	0.70 ± .29	0.70 ± .29

where $\lambda_c \geq 0$ (adaptive) and $\lambda_l \geq 0$ (fixed) weigh safety and goal-reaching, while $\nabla_{\mathbf{u}} \psi_h$ and $\nabla_{\mathbf{u}} \Phi_V$ steer the trajectory generation.

B. D-SafeMPC: MPC Projection

While CBF/CLF guidance structurally steers the diffusion output $\tilde{\tau}^{i-1}$ away from obvious hazards, it remains a *soft* constraint. To guarantee hard physical safety without causing solver divergence, the guided $\tilde{\tau}^{i-1}$ serves as a highly feasible warm-start for a Model Predictive Control (MPC) projection step:

$$\begin{aligned} \min_{\mathbf{u}_{0:H}, \bar{\mathbf{x}}_{0:H}} \quad & \|\tau^i\|_Q^2 - (\tilde{\tau}^{i-1})^\top Q(\tau^i) \\ \text{s.t} \quad & h(\bar{\mathbf{x}}_k) \geq 0, \quad \bar{\mathbf{x}}_{k+1} = \bar{\mathbf{F}}(\bar{\mathbf{x}}_k, \mathbf{u}_k) \end{aligned} \tag{2}$$

By tightly integrating (1) and (2), the MPC actively projects the probabilistically safe sample onto a strictly feasible set at each denoising iteration i. This refined trajectory τ^i is then fed back into the diffusion model for the next step, ensuring structural

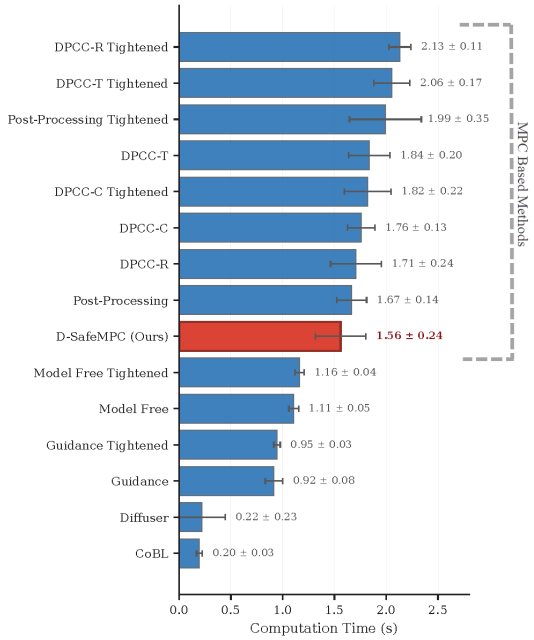

Fig. 2. A comparison of the average time (in seconds) each planner requires to generate a control action in the static obstacle environment.

safety throughout the generation process (as illustrated in Step 4 of Fig. 1).

III. Experiments

We evaluate D-SafeMPC in simulated drone inspection scenarios with both static and dynamic obstacles. Static obstacles remain at fixed positions, while dynamic obstacles follow circular orbits with randomized radii, angular velocities, and phase offsets, tracked via PD control. All planners are assessed across 10 random seeds using two metrics (Table I): the *Goal Reached Rate* (GR), measuring the proportion of collision-free goal-reaching trials, and the *Safety-Compliant Goal Reached Rate* (SCGR), which additionally requires all safety constraints to be satisfied throughout the trajectory.

D-SafeMPC demonstrates consistently superior performance across both static and dynamic environments. While CoBL achieves marginally higher success rates and DPCC-C attains better safe-success rates in the static scenario, our method remains highly competitive at 80% success and safe-success. Notably, D-SafeMPC excels in dynamic obstacle scenarios, achieving the highest or near-highest rates across all metrics, underscoring its robustness in enforcing safety constraints under challenging conditions. We also evaluated computational efficiency (Fig. 2). D-SafeMPC achieves a favorable balance between high safe-success rates and computational overhead, outperforming other MPC-based methods such as DPCC variants. We attribute this improvement to the CBF and CLF guidance, which provides the MPC with better-initialized trajectories, leading to faster convergence.

IV. Conclusion

In this work, we introduce D-SafeMPC, a novel framework that addresses the critical challenge of enforcing safety and dynamic feasibility in diffusion-based trajectory planning. By guiding the reverse diffusion process with discrete-time CBF and CLF, our approach steers trajectory generation toward safe, goal-directed regions. This guidance provides high-quality warm starts for a subsequent MPC. Our framework's effectiveness is demonstrated through an iterative-projection scheme where the MPC refines the trajectory at each denoising step, providing robust refinement and enforcing strict safety and dynamic constraints. Through extensive simulations in environments with both static and dynamic obstacles, alongside a sim-to-real experiment on a physical quadrotor drone, we have shown that D-SafeMPC significantly improves safety, task success rates, and planning efficiency compared to state-of-the-art baselines.

References

[1] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851.

[2] M. Janner, Y. Du, J. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *ICML*, 2022.

[3] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, 2025.

[4] J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters, "Motion planning diffusion: Learning and planning of robot motions with diffusion models," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1916–1923.

[5] J. Carvalho, A. T. Le, P. Kicki, D. Koert, and J. Peters, "Motion planning diffusion: Learning and adapting robot motion planning with diffusion models," *IEEE Transactions on Robotics*, 2025.

[6] W. Xiao, T.-H. Wang, C. Gan, R. Hasani, M. Lechner, and D. Rus, "Safediffuser: Safe planning with diffusion probabilistic models," in *The Thirteenth International Conference on Learning Representations*, 2023.

[7] V. Kurtz and J. W. Burdick, "Equality constrained diffusion for direct trajectory optimization," in *American Control Conference (ACC)*. IEEE, 2025, pp. 535–540.

[8] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[9] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019.

[10] K. Mizuta and K. Leung, "CoBL-Diffusion: Diffusion-based conditional robot planning in dynamic environments using control barrier and Lyapunov functions," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 13 801–13 808.

[11] U. Eren, A. Prach, B. B. Koçer, S. V. Raković, E. Kayacan, and B. Açıkmeşe, "Model predictive control in aerospace systems: Current state and opportunities," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 7, pp. 1541–1566, 2017.

[12] A. Amer, M. Mehndiratta, J. le Fevre Sejersen, H. X. Pham, and E. Kayacan, "Visual tracking nonlinear model predictive control method for autonomous wind turbine inspection," in *2023 21st International Conference on Advanced Robotics (ICAR)*, 2023, pp. 431–438.

[13] A. Amer, M. Mehndiratta, Y. Brodskiy, and E. Kayacan, "Empowering autonomous underwater vehicles using learning-based model predictive control with dynamic forgetting gaussian processes," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 5, pp. 1913–1920, 2025.

[14] R. Römer, A. von Rohr, and A. Schoellig, "Diffusion predictive control with constraints," in *L4DC*, 2025.

[15] G. Zhou, S. Swaminathan, R. V. Raju, J. S. Guntupalli, W. Lehrach, J. Ortiz, A. Dedieu, M. Lázaro-Gredilla, and K. Murphy, "Diffusion model predictive control," *arXiv preprint arXiv:2410.05364*, 2024.