Integrating Gradient-Based MPC and Control Barrier Functions for Safe and Performant Control

Aditya Singh*, Aastha Mishra*, Manan Tayal, and Shishir Kolathaya

Abstract— **Ensuring both safety and performance is essential for autonomous systems operating in real-world environments. While Control Barrier Function (CBF)-based safety filters enforce safety by modifying nominal controllers, they can become overly conservative when the nominal policy is not safety-aware. On the other hand, solving state-constrained optimal control problems is computationally challenging in high-dimensional settings. In this work, we propose a two-stage framework that combines gradient-based Model Predictive Control (MPC) with CBF-based safety filtering to jointly optimize safety and performance. First, safety constraints are relaxed as penalty terms in the cost function, enabling efficient gradient-based optimization and improved scalability. Second, a CBF-based Quadratic Program (CBF-QP) minimally modifies the resulting controller to enforce hard safety constraints. We validate the proposed framework on two case studies, demonstrating safe, high-performing, and computationally efficient control in complex high-dimensional systems.**

I. INTRODUCTION

Autonomous systems are increasingly being deployed in applications such as autonomous driving, robotic manipulation, aerospace, and industrial automation. A key challenge in controller design is achieving both *performance* and *safety* simultaneously. Performance focuses on accomplishing tasks efficiently under resource limitations such as time, energy, and actuation, while safety requires satisfying state and input constraints to avoid failures. These objectives are often conflicting; for example, aggressive maneuvers in drone delivery may improve efficiency but reduce safety in cluttered environments. As a result, designing controllers that optimize performance while guaranteeing safety remains a fundamental open problem.

One common approach to this problem is *safety filtering* [1], [2], where real-time corrections are introduced to maintain constraint satisfaction. Techniques such as Control Barrier Function (CBF)-based Quadratic Programs (QP) [3] and Hamilton–Jacobi (HJ) reachability filters [4], [5] enforce safety by minimally adjusting a nominal, potentially unsafe controller. However, if the nominal controller does not account for safety, these methods can produce overly conservative or short-sighted behavior. Another approach is to formulate the problem as a State-Constrained Optimal Control Problem (SC-OCP) and solve it through dynamic programming [6], [7]. Although effective in principle, this approach becomes computationally prohibitive in high-dimensional systems because of the curse of dimensionality.

To address these challenges, we introduce a two-stage framework that connects safety filtering with SC-OCP formulations. In the first stage, the SC-OCP is approximated as an online nonlinear Model Predictive Control (MPC) problem [8], where safety constraints are incorporated into the objective through penalty terms [9]. By avoiding hard nonlinear constraints, this formulation improves feasibility and enables scalable gradient-based optimization, offering advantages over sampling-based approaches such as MPPI [10]–[14] in high-dimensional settings.

In the second stage, the MPC-generated policy serves as the reference input for a CBF-based QP. Since safety considerations are already encoded in the reference policy, the filtering stage avoids unnecessary conservatism while still providing formal safety guarantees. The resulting framework therefore achieves both high performance and provable safety while remaining computationally scalable.

To summarize, the main contributions of this work are:

- We develop a two-stage control framework that combines gradient-based MPC with CBF-based safety filtering to jointly address safety and performance objectives.
- By employing gradient-based optimization instead of sampling-based methods, the proposed approach improves computational efficiency and scalability, especially in high-dimensional systems.
- We integrate a CBF-based Quadratic Program (CBF-QP) to impose safety as a hard constraint, ensuring formal guarantees while minimally altering the nominal controller.
- We demonstrate the effectiveness of the framework on unicycle and quadrotor navigation tasks, showing reliable safety, improved performance, and scalability to more complex systems.

II. PROBLEM FORMULATION

Consider a control-affine system with state $x \in \mathcal{X} \subseteq \mathbb{R}^n$ and control input $u \in \mathcal{U} \subseteq \mathbb{R}^m$, described by the dynamics $\dot{x} = f(x) + g(x)u$, where $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ are continuously differentiable functions. Let $\mathcal{F} \subseteq \mathcal{X}$ represent the failure set corresponding to unsafe states. System performance is quantified through the cost functional:

$$J(t, x(t), \mathbf{u}) = \int_{s=t}^{T} r(x(s))\, ds + \phi(x(T)), \tag{1}$$

where $r : \mathcal{X} \to \mathbb{R}_{\geq 0}$ and $\phi : \mathcal{X} \to \mathbb{R}_{\geq 0}$ are non-negative and Lipschitz continuous functions denoting the running

---

This work was supported by the AI & Robotics Technology Park (ARTPARK) at IISc.

All the authors belong to Cyber-Physical Systems, Indian Institute of Science (IISc), Bengaluru. {adityasingh, aasthamishra, manantayal, shishirk}@iisc.ac.in

* denotes equal contribution.

and terminal costs, respectively. The control trajectory $\mathbf{u} : [t, T) \to \mathcal{U}$ governs the system evolution. The goal is to determine an optimal policy $\pi^* : [t, T) \times \mathcal{X} \to \mathcal{U}$ that minimizes J while ensuring that the system state remains outside the failure set $\mathcal{F}$ over the interval $[t, T]$. To impose the safety requirement, we define a function $l : \mathbb{R}^n \to \mathbb{R}$ such that $\mathcal{F} := \{x \in \mathcal{X} \mid l(x) \leq 0\}$. Using this definition, the control synthesis problem can be formulated as the following State-Constrained Optimal Control Problem (SC-OCP):

Problem 1 (State-Constrained Optimal Control Problem)**.**

$$\begin{aligned} \inf_{\mathbf{u}} J(t, x(t), \mathbf{u}) = \int_t^T r(x(s))\, ds + \phi(x(T)) \\ \textit{s.t.} \quad \dot{x} = f(x) + g(x)u, \\ l(x(s)) > 0 \quad \forall s \in [t, T] \end{aligned} \tag{2}$$

The SC-OCP improves system performance by minimizing the objective function while simultaneously enforcing safety through the state constraint $l(x) > 0$, thereby ensuring avoidance of the failure set $\mathcal{F}$. Consequently, the optimal policy π^* obtained from this SC-OCP jointly optimizes performance and safety.

A. MPC Reformulation

To obtain an efficient and robust solution to the SC-OCP in Problem 1, we employ a Model Predictive Control (MPC) framework. Rather than solving the optimization problem over the full horizon, the problem is repeatedly solved over a shorter horizon h. At each step, the first control action is applied and the optimization is resolved using updated state information. This receding-horizon approach enhances robustness against modeling errors and external disturbances. We first discretize both the system dynamics and the SC-OCP. Let $\mathbf{x}$ and $\mathbf{u}$ denote the discrete-time state and control sequences, respectively. The resulting discrete-time formulation is:

Problem 2 (Discrete-Time SC-OCP)**.**

$$\begin{aligned} \min_{\mathbf{u}} \quad & J(x, \mathbf{u}) = \sum_{k=1}^{K} r(\mathbf{x}(k), \mathbf{u}(k)) + \phi(\mathbf{x}(K)) \\ \textit{s.t.} \quad & \mathbf{x}(k+1) = f_d(\mathbf{x}(k)) + g_d(\mathbf{x}(k))\mathbf{u}(k), \\ & l(\mathbf{x}(k)) > 0, \quad \forall k \in \{1, \dots, K\}. \end{aligned} \tag{3}$$

Here, f_d and g_d denote the differentiable discrete-time system dynamics. At each time step j, we solve a finite-horizon approximation of Problem 2 over a prediction horizon h, resulting in the following MPC formulation:

Problem 3 (MPC Reformulation of the SC-OCP)**.**

$$\begin{aligned} \min_{\mathbf{u}} \quad & \sum_{k=j}^{j+h} r(\mathbf{x}(k), \mathbf{u}(k)) + \phi(\mathbf{x}(j+h)) \\ \textit{s.t.} \quad & \mathbf{x}(k+1) = f_d(\mathbf{x}(k)) + g_d(\mathbf{x}(k))\mathbf{u}(k), \\ & l(\mathbf{x}(k)) > 0, \quad \forall k \in \{j, \dots, j+h\}. \end{aligned} \tag{4}$$

Based on this formulation, we define the primary objective of this work:

Objective 1. Our objective is to design a scalable and computationally efficient framework for solving Problem 3, which can subsequently be used to synthesize an optimal policy $\pi^* : [t, T) \times \mathcal{X} \to \mathcal{U}$ that minimizes the cost J while guaranteeing that the system state remains outside the failure set $\mathcal{F}$ throughout the horizon $[t, T]$.

III. METHODOLOGY

As discussed in Section II, the SC-OCP formulation involves solving a nonlinear MPC problem (Problem 3) at every time step over a finite horizon. This is difficult because nonlinear system dynamics and safety constraints must both be handled efficiently for real-time implementation. A critical design decision therefore concerns the choice of planning framework. Existing methods can broadly be categorized into sampling-based approaches (e.g., MPPI) and gradient-based approaches. Sampling-based methods evaluate randomly sampled control trajectories using scalar cost feedback, which can result in limited scalability in high-dimensional problems. By contrast, gradient-based methods exploit differentiable dynamics to compute full cost gradients, leading to more informative updates, faster convergence, and improved scalability with respect to action dimensionality. Motivated by these advantages, we employ a gradient-based planning strategy to solve the MPC problem in Problem 3.

A. Gradient-Based Planning

Although gradient-based approaches provide computational advantages, directly solving Problem 3 remains challenging due to the existence of hard safety constraints. To overcome this issue, we reformulate the problem by relaxing the safety constraints and incorporating them into the objective function as soft penalties. This preserves the structure of the original problem while producing a smoother and more tractable optimization landscape:

Problem 4 (MPC Formulation with Soft Safety Constraint)**.**

$$\begin{aligned} \min_{\mathbf{u}} \quad & C(x, \mathbf{u}) = \sum_{k=j}^{j+h} \Big(r(\mathbf{x}(k), \mathbf{u}(k)) + \phi(\mathbf{x}(j+h)) \\ & \qquad + \lambda \max\{0, -l(\mathbf{x}(k)) + \delta\} \Big) \\ \textit{s.t.} \quad & \mathbf{x}(k+1) = f_d(\mathbf{x}(k)) + g_d(\mathbf{x}(k))\mathbf{u}(k), \ \forall k \in \mathbb{N}, \end{aligned} \tag{5}$$

In this formulation, the parameter $\lambda > 0$ controls the trade-off between safety and performance, while $\delta > 0$ introduces a conservative margin to improve robustness against constraint violations. Larger values of λ place greater emphasis on safety, whereas smaller values favor performance. We solve Problem 4 using the L-BFGS algorithm [18], a quasi-Newton optimization method that efficiently approximates second-order information without explicitly computing the Hessian. The control sequence is iteratively updated as:

$$\mathbf{u}(k)_{i+1} = \mathbf{u}(k)_i - \alpha_i H_i \nabla C(\mathbf{x}(k), \mathbf{u}(k)_i), \tag{6}$$

Algorithm 1 Gradient-based Planning with CBF-QP Safety Filter

```
Input: Initial state $x_0$, goal position $x_{goal}$, horizon $H$, number of steps $T$
Parameters: Dynamics $f$ and $g$, cost $C$
for $t = 0$ to $T-1$ do
    Solve MPC:
    $\mathbf{x_0} \leftarrow x_t$
    Initialize $u_{0:H-1}$
    Optimize $u_{0:H-1}$ using L-BFGS on cost:
        $J = \sum_{k=0}^{H-1} C(\mathbf{x_k}, u_k)$,
        $\mathbf{x_{k+1}} = \mathbf{x_k} + (f(\mathbf{x_k}) + g(\mathbf{x_k})u_k)\Delta t$
    $u_{\text{nom}} \leftarrow u_0$
    Apply CBF-QP Safety Filter:
    $u^* \leftarrow$ Solve CBF-QP from (8)
    $x_{t+1} \leftarrow x_t + (f(x_t) + g(x_t)u^*)\Delta t$
end for
```

where α_i is determined through line search and H_i denotes the inverse Hessian approximation. This facilitates efficient optimization with low memory requirements, making the method suitable for high-dimensional and real-time applications. Gradients are obtained through automatic differentiation, which naturally handles non-differentiability in the objective function. The resulting relaxation produces a nominal controller that balances safety and performance. However, because safety is enforced only softly, constraint violations may still occur. This motivates the introduction of a safety filtering mechanism.

B. Safe and Performant Controller Synthesis using CBFs

To provide formal safety guarantees, we augment the nominal controller with a CBF-based safety filter [19]. The main idea is to minimally adjust the control input generated by the planner such that safety constraints are strictly satisfied. We first define a continuously differentiable function $h : \mathcal{X} \to \mathbb{R}$ whose super-level set defines the safe region $\mathcal{C} = \{x \in \mathcal{X} : h(x) \geq 0\}$. For the control-affine system $\dot{x} = f(x) + g(x)u$, the function h is considered a valid CBF [3], [20] if there exists an extended class-$\mathcal{K}$ function κ such that:

$$\sup_{u \in \mathcal{U}} \left[\mathcal{L}_f h(x) + \mathcal{L}_g h(x)u + \kappa(h(x))\right] \geq 0, \tag{7}$$

where, $\mathcal{L}_f h = \frac{\partial h}{\partial x} f$ and $\mathcal{L}_g h = \frac{\partial h}{\partial x} g$. This condition guarantees forward invariance of the safe set. Using this property, we formulate a quadratic program that minimally deviates from the nominal control input u_{mpc}, following the standard CBF-QP formulation [19]:

$$\begin{aligned} \mathbf{u}^*(x) &= \arg\min_{u \in \mathcal{U}} \; \|u - u_{\text{mpc}}\|^2 \\ \text{s.t.} \quad & \mathcal{L}_f h(x) + \mathcal{L}_g h(x)u + \kappa(h(x)) \geq 0. \end{aligned} \tag{8}$$

This formulation ensures that the applied control input remains as close as possible to the performance-oriented MPC solution while strictly enforcing safety constraints. The overall approach is summarized in Algorithm 1.

IV. RESULTS AND DISCUSSION

We evaluate the proposed framework using three metrics that capture both performance and safety. (1) **Cumulative Cost**: $\sum_{k=1}^{K} r(\mathbf{x}(k), \mathbf{u}(k)) + \phi(\mathbf{x}(K))$, which measures task performance along safe trajectories. (2) **Safety Rate**: the percentage of trajectories that remain inside the safe set, i.e., trajectories that avoid the failure set $\mathcal{F}$. (3) **Computation Time**: the average time required to solve the MPC optimization at each step. **Baselines and Implementation Details**: To isolate the role of gradient-based planning, we compare against MPPI-CBF [21], where MPPI provides the reference control for the CBF filter. We also include MPPI and gradient-based MPC (GMPC) without CBF filtering as ablations, directly solving Problem 4. All experiments were conducted on an 11th Gen Intel Core i9-11900K CPU (16 cores) with 128GB RAM.

A. Unicycle Robot Navigation with Collision Avoidance

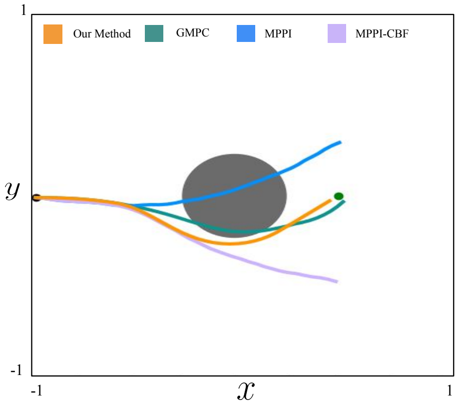

Fig. 1: Trajectories from a common initial condition are shown, where the dark grey circle denotes the obstacle and the green dot represents the goal at $[0.5, 0]^T$. The proposed method is the **only approach** that **successfully reaches the goal** while **satisfying safety constraints**.

In this experiment, we consider a **3D Dubins Car** with dynamics: $\dot{x} = v\cos\theta, \quad \dot{y} = v\sin\theta, \quad \dot{\theta} = u_1$, where (x, y, θ) represent the position and heading, v is a constant velocity, and u_1 denotes the control input. The objective is to navigate the system from a random initial state to a target location in the xy-plane while avoiding a circular obstacle over a 2-second horizon. The MPC formulation uses a time step of 0.05 seconds and a planning horizon of 20 steps. The running cost in Problem 3 is defined as $\left\|[x, y]^T - [x_g, y_g]^T\right\|$, without a terminal cost term. Safety is imposed through the constraint $l = \left\|[x, y]^T - [x_o, y_o]^T\right\| - r$, where $[x_o, y_o]^T$ and r correspond to the obstacle center and radius. Although l may be used as a CBF candidate, it does not explicitly depend on the control input. To address this limitation, we employ a Higher-Order Control Barrier Function (HOCBF) [22], [23]:

$$h(x, t) = 2xv\cos\theta + 2yv\sin\theta + \alpha(x^2 + y^2 - r^2), \tag{9}$$

where $\alpha > 0$ determines the convergence rate toward the safe set. As illustrated in Figure 1, the proposed method reaches

	Unicycle Navigation			Quadrotor Navigation		
Method	**Cost**	**Safety Rate (%)**	**Computation Time (s)**	**Cost**	**Safety Rate (%)**	**Computation Time (s)**
GMPC	1.417	85	0.14	1.036	88	0.3
MPPI	2.085	85	0.24	2.366	87	0.7
MPPI-CBF	2.071	100	0.20	2.302	100	0.7
GMPC-CBF (Ours)	1.403	100	0.14	1.093	100	0.3

TABLE I: Comparative analysis of all evaluated methods based on the defined evaluation metrics

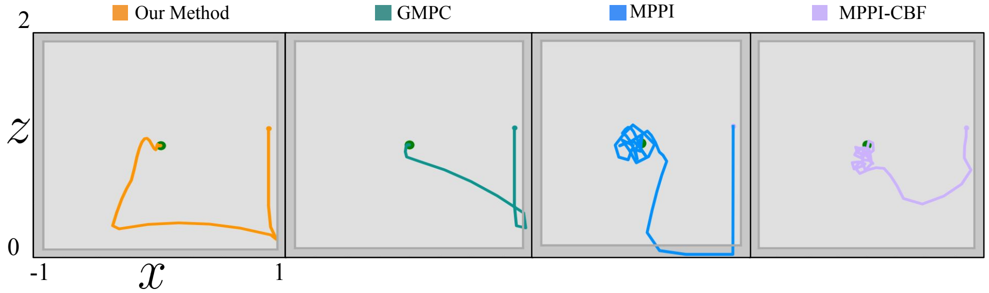

Fig. 2: Trajectories initialized from a common starting state are shown, where the dark grey region indicates the walls, floor, and ceiling boundaries. The proposed framework successfully **reaches the goal at $[\mathbf{0}, \mathbf{0}]^{\mathbf{T}}$** (green) while **maintaining collision avoidance**. In comparison, MPPI and GMPC violate safety constraints, whereas MPPI-CBF remains safe but fails to converge to the goal, instead oscillating near it and incurring a higher performance cost.

the goal while preserving safety, even under challenging initial conditions. In contrast, MPPI and GMPC violate safety constraints, whereas MPPI-CBF, although safe, is overly conservative and fails to reach the goal. This behavior highlights a limitation of sampling-based approaches, which depend on locally optimal solutions obtained from finite sampling and degrade over longer horizons. In comparison, gradient-based planning provides more informative updates and produces improved reference controls for the CBF-QP filter. Table I further validates these observations. MPPI-CBF incurs a 47.6% higher mean cost together with greater computational overhead, while GMPC and MPPI achieve lower safety rates. Although GMPC is computationally efficient, the absence of formal safety guarantees limits its practical applicability. In contrast, the proposed framework achieves a balanced trade-off between safety and performance.

B. *Quadrotor Navigation within a Closed Room Setup*

We next consider a **6-dimensional planar quadrotor** system [24] with state $[x, z, \theta, \dot{x}, \dot{z}, \dot{\theta}]$ governed by the dynamics: $\dot{x} = \dot{x}, \quad \dot{z} = \dot{z}, \quad \dot{\theta} = \dot{\theta}, \quad m\ddot{x} = F\cos\theta, \quad m\ddot{z} = F\sin\theta - mg, \quad J\ddot{\theta} = -M$, where (x, z, θ) denote the position and pitch angle, and F, M represent the thrust and torque inputs. The task is to drive the system to a target position in the x–z plane while avoiding collisions with the walls, ceiling, and floor over a 3-second horizon. The MPC implementation uses a time step of 0.05 seconds and a planning horizon of 20 steps. The running cost is given by $r(x) = \|[x, z]^T - [x_g, z_g]^T\|$, while safety is encoded using $\min(1.9 - z,\ z - 0.1,\ 0.9 - x,\ 0.9 + x)$, which represents distances to the environment boundaries. To enforce these constraints, we formulate four prioritized CBF-QPs, each associated with a particular boundary. At every time step, the QPs are solved sequentially according to proximity, with the solution of one QP used as the reference input for the next, thereby producing a safety-prioritized filtering mechanism. The corresponding HOCBF conditions are:

$$\begin{aligned}
C_1 &= -F\omega\cos\theta - 3\alpha(g - F\sin\theta) - 3\alpha^2\dot{z} + \alpha^3(1.9 - z),\\
C_2 &= F\omega\cos\theta + 3\alpha(F\sin\theta - g) + 3\alpha^2\dot{z} + \alpha^3(z - 0.1),\\
C_3 &= F\omega\sin\theta - 3\alpha F\cos\theta - 3\alpha^2\dot{x} + \alpha^3(0.9 - x),\\
C_4 &= -F\omega\sin\theta + 3\alpha F\cos\theta - 3\alpha^2\dot{x} + \alpha^3(0.9 + x).
\end{aligned} \tag{10}$$

As shown in Figure 2, the proposed method successfully reaches the goal while strictly satisfying safety constraints, even from challenging initial conditions. By comparison, MPPI and GMPC violate constraints, while MPPI-CBF, although safe, fails to reach the goal and incurs a higher cost. Table I further supports these findings: MPPI-CBF achieves safety at the expense of a 110.6% higher cost and increased computation time, whereas GMPC and MPPI exhibit lower safety rates. Notably, the gap between sampling-based and gradient-based methods becomes more pronounced in this higher-dimensional setting, since the computational complexity of sampling-based approaches scales superlinearly. Overall, the proposed framework achieves an effective balance between safety, performance, and scalability.

V. CONCLUSION

We presented a two-stage control framework that integrates gradient-based MPC with CBF-based safety filtering to jointly achieve safety and performance. Experimental results demonstrate that proposed controllers are safe, high-performing, and computationally efficient across multiple tasks. However, the current framework lacks global convergence guarantees and requires manual tuning of CBF parameters. Future work will focus on leveraging learning-based methods for verified CBF construction, and extending the framework to more complex high-dimensional systems.

REFERENCES

[1] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, no. Volume 7, 2024, pp. 47–72, 2024. [Online]. Available: https://www.annualreviews.org/content/journals/10.1146/annurev-control-071723-102940

[2] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems Magazine*, vol. 43, no. 5, pp. 137–177, 2023.

[3] 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.

[4] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017.

[5] J. Borquez, K. Chakraborty, H. Wang, and S. Bansal, "On safety and liveness filtering using hamilton–jacobi reachability analysis," *IEEE Transactions on Robotics*, vol. 40, pp. 4235–4251, 2024.

[6] A. Altarovici, O. Bokanowski, and H. Zidani, "A general hamilton-jacobi framework for non-linear state-constrained control problems," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 19, no. 2, pp. 337–357, 2013.

[7] H. Wang, A. Dhande, and S. Bansal, "Cooptimizing safety and performance with a control-constrained formulation," *IEEE Control Systems Letters*, vol. 8, pp. 2739–2744, 2024.

[8] H. Wang, A. Jordana, L. Righetti, and S. Bansal, "Safe and performant deployment of autonomous systems via model predictive control and hamilton-jacobi reachability analysis," 2025. [Online]. Available: https://arxiv.org/abs/2506.23346

[9] K. Srinivasan, B. Eysenbach, S. Ha, J. Tan, and C. Finn, "Learning to be safe: Deep rl with a safety critic," 2020. [Online]. Available: https://arxiv.org/abs/2010.14603

[10] H. Bharadhwaj, K. Xie, and F. Shkurti, "Model-predictive control via cross-entropy and gradient-based optimization," in *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, ser. Proceedings of Machine Learning Research, A. M. Bayen, A. Jadbabaie, G. Pappas, P. A. Parrilo, B. Recht, C. Tomlin, and M. Zeilinger, Eds., vol. 120. PMLR, 10–11 Jun 2020, pp. 277–286. [Online]. Available: https://proceedings.mlr.press/v120/bharadhwaj20a.html

[11] B. Käpernick and K. Graichen, "The gradient-based nonlinear model predictive control software grampc," in *2014 European Control Conference (ECC)*, 2014, pp. 1170–1175.

[12] S. Hall, L. Ortmann, M. Picallo, and F. Dörfler, "Real-time projected gradient-based nonlinear model predictive control with an application to anesthesia control," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 6193–6198.

[13] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.

[14] L. Streichenberg, E. Trevisan, J. J. Chung, R. Siegwart, and J. Alonso-Mora, "Multi-agent path integral control for interaction-aware motion planning in urban canals," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1379–1385.

[15] M. Tayal, A. Singh, P. Jagtap, and S. Kolathaya, "Semi-supervised safe visuomotor policy synthesis using barrier certificates," *2025 IEEE 64th Conference on Decision and Control (CDC)*, 2025.

[16] A. Singh, Z. Feng, and S. Bansal, "Exact imposition of safety boundary conditions in neural reachable tubes," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025. [Online]. Available: https://arxiv.org/abs/2404.00814

[17] M. Tayal, A. Singh, S. Kolathaya, and S. Bansal, "A physics-informed machine learning framework for safe and optimal control of autonomous systems," in *Forty-second International Conference on Machine Learning*, 2025. [Online]. Available: https://openreview.net/forum?id=SrfwiloGQF

[18] J. Nocedal, "Updating quasi-newton matrices with limited storage," *Mathematics of Computation*, vol. 35, no. 151, pp. 773–782, 1980. [Online]. Available: http://www.jstor.org/stable/2006193

[19] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 6271–6278.

[20] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *18th European control conference (ECC)*. IEEE, 2019, pp. 3420–3431.

[21] P. Rabiee and J. B. Hoagg, "Guaranteed-safe mppi through composite control barrier functions for efficient sampling in multi-constrained robotic systems," 2024. [Online]. Available: https://arxiv.org/abs/2410.02154

[22] W. Xiao and C. Belta, "High-order control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3655–3662, 2022.

[23] ——, "Control barrier functions for systems with high relative degree," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 474–479.

[24] G. Wu and K. Sreenath, "Safety-critical control of a planar quadrotor," in *2016 American Control Conference (ACC)*, 2016, pp. 2252–2258.