Velocity Obstacle-Based Control Barrier Function for Safety-Critical Predictive Control

Francesco Trotti[1,*] and Daniele Meli[1]

***Abstract*—This paper introduces the Velocity Obstacle-based Control Barrier Function (VO-CBF), a novel safety framework bridging geometric collision prediction with dynamic control. While classical Velocity Obstacles ignore system dynamics, the VO-CBF encodes collision cones directly into the dynamic state space. By natively coupling position and velocity, this formulation mathematically reduces the constraint's relative degree from two to one for force-controlled systems, fundamentally eliminating the solver conservatism and input saturation conflicts associated with standard high-order predictive barriers. We propose a hybrid Nonlinear Model Predictive Control (NMPC) architecture that enforces strict discrete-time invariance for the immediate step and affine approximations across the prediction horizon, balancing formal safety guarantees with computational tractability. Comparative simulations and hardware experiments demonstrate that the VO-CBF effectively throttles velocity to prevent collisions in high-inertia scenarios where standard kinematic baselines fail.**

***Index Terms*—Optimization and Optimal Control, Integrated Planning and Control, Constrained Motion Planning, Autonomous Agents, Collision Avoidance.**

I. INTRODUCTION

AUTONOMOUS systems are increasingly required to operate safely in dynamic and uncertain environments. Ensuring collision avoidance and the forward invariance of safe sets remains a primary challenge for real-world deployment. Because safety cannot be treated merely as a high-level planning constraint [1], [2], [3], [4], it must be enforced in real-time within the low-level control loop.

Control Barrier Functions (CBFs) [5], [6] effectively enforce set invariance in nonlinear control-affine systems, often combined with Control Lyapunov Functions (CLFs) to balance safety and stability [7], [8]. However, traditional velocity-dependent CBFs rely on High-Order CBFs (HO-CBF) [9], which react to instantaneous state derivatives rather than anticipating the geometric trajectories of moving obstacles. This lack of foresight yields complex constraints that easily conflict with input saturation [10]. Furthermore, integrating Discrete-Time CBFs (DTCBF) into Model Predictive Control (MPC) frameworks [11], [12], [13] typically relies on simple position-based distance bubbles. Lacking predictive velocity alignment, these formulations fail to anticipate future interactions, resulting in highly conservative evasive maneuvers or freezing-robot behaviors when facing fast-moving obstacles. Alternatively, the Velocity Obstacle (VO) framework [14], [15], [16] provides rigorous geometric collision avoidance in velocity space. While computationally efficient and widely used in multi-agent navigation, standard VOs operate purely reactively and kinematically [17]. They assume instantaneous velocity tracking, ignoring inertia and actuation limits. Therefore, they cannot formally guarantee constraint feasibility under strict dynamic limits or optimize maneuvers over a predictive horizon.

A critical gap remains: unifying the proactive geometric foresight of VO with the optimal predictive horizon of NMPC without sacrificing strict mathematical safety guarantees under actuator limits.

This paper proposes the Velocity Obstacle-based Control Barrier Function (VO-CBF) to bridge this gap. Unlike prior velocity-dependent CBFs, the VO-CBF mathematically maps the infinite-horizon predictive geometry of relative kinematics directly into the optimization landscape. It encodes both positional separation and relative velocity alignment, providing a dynamic safety constraint that naturally accounts for the system's momentum and the obstacle's predicted motion. We integrate this formulation within an NMPC architecture to ensure real-time safety at the control input level.

The main contributions of this paper are:

- *Analytic Regularization & Relative Degree Reduction:* We propose the VO-CBF, transforming the non-convex Velocity Obstacle cone into a continuously differentiable barrier. Crucially, by natively coupling position and velocity, this formulation structurally reduces the constraint's relative degree from two to one for force-controlled systems. This eliminates the need for complex high-order backstepping, encoding predictive collision risks without inducing artificial input saturation conflicts.
- *Hybrid Discrete-Continuous Safety Layer:* We introduce a decoupled hybrid constraint architecture for NMPC that resolves the discretization error inherent in standard CBF implementations. By enforcing strict discrete-time invariance (forward difference) for the immediate control action and affine approximations (Lie derivatives) for the prediction horizon, we achieve rigorous safety guarantees without sacrificing solver tractability.
- *Recursive Feasibility Analysis:* We provide a formal proof of safety preservation under the proposed hybrid scheme, demonstrating that the discrete-time evolution of the closed-loop system remains forward invariant with respect to the VO-CBF set while avoiding solver failure under strict actuation limits.

Manuscript received: January, 28, 2026; Revised March, 19, 2026; Accepted May 19, 2026.

This paper was recommended for publication by Editor Lucia Pallottino upon evaluation of the Associate Editor and Reviewers' comments. This project has been funded by the grant "Monitoraggio automatico dell'impatto ambientale dei processi agricoli sulle acque fluviali", funded by Bando Ricerca e Sviluppo 2025 by Cassa di Risparmio di Trento e Rovereto CARITRO (code 2024.0571)

*Corresponding author.

[1]Francesco Trotti and Daniele Meli are with the Department of Computer Science, University of Verona, 37134, Italy. francesco.trotti@univr.it, daniele.meli@univr.it

The remainder of this paper is organized as follows: Section II reviews the necessary background. Section III details the VO-CBF design. Section IV integrates the VO-CBF into the NMPC and analyzes the discretization strategy. Section V presents simulated and real-world experiments. Section VI draws conclusions.

II. BACKGROUND

In this section, we cover the main foundations related to the research topics considered in the paper.

A. Control Barrier Function

Consider a nonlinear control-affine system described by the differential equation:

$$\dot{x} = f(x) + g(x)u, \tag{1}$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ denotes the system state, and $u \in \mathcal{U} \subset \mathbb{R}^m$ represents the control input. The vector fields $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ are assumed to be locally Lipschitz continuous, ensuring the existence and uniqueness of solutions to (1). In this context, the notion of a Control Barrier Function is introduced to enforce the forward invariance of a designated safe set.

Definition 1 (Control Barrier Function [5]): Let $h : \mathbb{R}^n \to \mathbb{R}$ be a continuously differentiable function, and define the associated safe set $\mathcal{C} \subset \mathbb{R}^n$ as:

$$\begin{aligned} \mathcal{C} &= \{x \in \mathbb{R}^n : h(x) \geq 0\}, \\ \partial\mathcal{C} &= \{x \in \mathbb{R}^n : h(x) = 0\}, \\ \text{Int}(\mathcal{C}) &= \{x \in \mathbb{R}^n : h(x) > 0\}. \end{aligned} \tag{2}$$

The function h is said to be a Control Barrier Function if $\nabla h(x) \neq 0$ for all $x \in \partial\mathcal{C}$, and there exists an extended class-$\mathcal{K}$ function α such that for all $x \in \mathcal{C}$, there exists a control input $u \in \mathcal{U}$ satisfying:

$$L_f h(x) + L_g h(x)u \geq -\alpha(h(x)), \tag{3}$$

where the Lie derivatives are defined as:

$$L_f h(x) := \nabla h(x)^T f(x), \quad L_g h(x) := \nabla h(x)^T g(x). \tag{4}$$

To illustrate the concept, consider the scalar single integrator system $\dot{x} = u, \quad x \in \mathbb{R}, \quad u \in \mathbb{R}$. Define the safe set $\mathcal{C} = \{x \in \mathbb{R} : h(x) = 1 - x^2 \geq 0\}$, which corresponds to the interval $[-1, 1]$. Let $h(x) = 1 - x^2$, so that $\nabla h(x) = -2x, \quad f(x) = 0, \quad g(x) = 1$. The Lie derivatives become $L_f h(x) = 0, \quad L_g h(x) = -2x$. The CBF condition requires the existence of a control input u such that $-2xu \geq -\alpha(1 - x^2)$. Choosing $\alpha(h) = h$, the condition simplifies to $u \leq \frac{1-x^2}{2x}$ for $x > 0$, $u \geq \frac{1-x^2}{2x}$ for $x < 0$. Thus, the control input must satisfy the above inequality to ensure forward invariance of the safe set $\mathcal{C}$.

B. Velocity Obstacle

Let an agent A be represented as a circle of radius r_A, located at position $\mathbf{p}_A \in \mathbb{R}^2$, with velocity $\mathbf{v}_A \in \mathbb{R}^2$. Similarly, let a moving obstacle B be a circle of radius r_B, located at $\mathbf{p}_B$, with velocity $\mathbf{v}_B$. The Velocity Obstacle for agent A induced by B, denoted $\text{VO}_{A|B}$, is defined as a set of all velocities $\mathbf{v}_A$ such that the relative velocity $\mathbf{v}_A - \mathbf{v}_B$ will result in a collision within a finite time horizon τ. Formally,

$$\begin{aligned} \text{VO}^{\tau}_{A|B} = \big\{ \mathbf{v}_A \in \mathbb{R}^2 \,\big|\, \exists t \in [0, \tau] \\ : \|\mathbf{p}_A + t\mathbf{v}_A - (\mathbf{p}_B + t\mathbf{v}_B)\| < r_A + r_B \big\}. \end{aligned} \tag{5}$$

To avoid collision, agent A must select a velocity $\mathbf{v}_A \notin \text{VO}^{\tau}_{A|B}$.

III. VELOCITY OBSTACLE-BASED CONTROL BARRIER FUNCTION

In this section, we formalize the Velocity Obstacle concept as a differentiable Control Barrier Function. Unlike classical VO formulations which operate on kinematic assumptions, we embed the collision cone constraint directly into the dynamic state space.

A. Geometric Definitions

Let the robot and obstacle positions be $\mathbf{p}_r, \mathbf{p}_o \in \mathbb{R}^n$ and their velocities be $\mathbf{v}_r, \mathbf{v}_o \in \mathbb{R}^n$. We define the relative state vectors as:

$$\mathbf{p}_{rel} \coloneqq \mathbf{p}_o - \mathbf{p}_r, \quad \mathbf{v}_{rel} \coloneqq \mathbf{v}_r - \mathbf{v}_o. \tag{6}$$

Here, $\mathbf{p}_{rel}$ represents the vector pointing from the robot to the obstacle. The system is on a collision course if the relative velocity $\mathbf{v}_{rel}$ lies within the VO cone defined by the combined radius $R = r_r + r_o$.

We define the regularized distance d and relative speed ν as:

$$d = \sqrt{\mathbf{p}_{rel}^{\top}\mathbf{p}_{rel} + \varepsilon_d}, \quad \nu = \sqrt{\mathbf{v}_{rel}^{\top}\mathbf{v}_{rel} + \varepsilon_v}, \tag{7}$$

To guarantee Lipschitz continuity and strict numerical conditioning for the optimization solvers, small regularization constants ($\epsilon_d, \epsilon_v \approx 10^{-8}$) are introduced into the distance and velocity norms. In standard Velocity Obstacle formulations, exact zero relative velocity or zero relative distance induces singularities (division by zero) in the gradient calculations, which natively crash gradient-based SQP and NLP solvers. Because these constants are on the scale of numerical precision, they successfully bound the Jacobians to finite values without inducing any measurable geometric distortion to the physical safety boundary or compromising the theoretical forward invariance of the set.

The VO cone half-angle α satisfies:

$$\cos\alpha = \sqrt{1 - \left(\frac{R}{d}\right)^2}. \tag{8}$$

We define the alignment metric Φ, which represents the cosine of the angle between the relative velocity and the line-of-sight vector:

$$\Phi = \frac{\mathbf{p}_{rel}^{\top}\mathbf{v}_{rel}}{d\nu}. \tag{9}$$

Geometrically, if $\Phi \geq \cos\alpha$, the relative velocity vector lies inside the collision cone.

B. Barrier Function Design

To ensure safety, we require the system to maintain a configuration outside the collision cone and maintain a safe distance from the obstacle. We propose the following candidate CBF:

$$h(x) = \gamma_d(d - R) + \gamma_a(\cos\alpha - \Phi). \tag{10}$$

The parameters $\gamma_d > 0$ and $\gamma_a > 0$ serve as weighting coefficients that balance spatial separation against proactive velocity alignment. To ensure safe and non-conservative behavior, we establish a strict tuning hierarchy where the distance penalty heavily dominates the velocity alignment ($\gamma_d \gg \gamma_a$). This rationale ensures that the physical distance boundary, the absolute condition for collision, strictly dominates the constraint space as the robot approaches the safety limit. The velocity alignment weight, γ_a, acts as a predictive shaping term; tuning it too high results in overly conservative, premature evasive maneuvers, while tuning it too low reduces the constraint to a purely reactive position bubble. The safe set is defined as $\mathcal{C} = \{x \in \mathbb{R}^n \mid h(x) \geq 0\}$. The term $\gamma_d(d-R)$ enforces spatial separation, while $\gamma_a(\cos\alpha - \Phi)$ penalizes velocity alignment with the obstacle.

Conversely, the proposed VO-CBF $h(p, v)$ natively embeds the velocity state. Evaluating its first time derivative yields:

$$\dot{h}(p, v) = \frac{\partial h}{\partial p}\dot{p} + \frac{\partial h}{\partial v}\dot{v} = L_f h(x) + L_g h(x)u, \tag{11}$$

because the velocity alignment terms ensure that $\frac{\partial h}{\partial v} \neq 0$, the Lie derivative $L_g h(x)$ is non-zero. The control input u explicitly appears in the first derivative, strictly reducing the constraint to relative degree one. This structural reduction eliminates the need for high-order barrier backstepping, directly mitigating the input saturation conflicts and solver conservatism typically associated with second-order predictive safety constraints.

C. System Dynamics and Time Derivatives

We assume the robot dynamics are control-affine:

$$\dot{x} = f(x) + g(x)u \implies \mathbf{a}_r = f_a(x) + g_a(x)u. \tag{12}$$

Assumption 1: We assume the obstacle's state $\boldsymbol{x}_o = [p_o, v_o, a_o]^\top$ is available to the robot. While p_o and v_o are standard in VO frameworks, the proposed dynamic formulation explicitly utilizes the obstacle acceleration a_o to anticipate the rotation of the collision cone. In practice, a_o can be estimated via high-gain observers [18] or Extended Kalman Filters [19]; here, we treat a_o as a known bounded disturbance to isolate the safety guarantees of the controller.

To enforce the CBF condition $\dot{h}(x, u) \geq -\kappa h(x)$, we derive the time derivative of (10). Note that $\dot{\mathbf{p}}_{rel} = -\mathbf{v}_{rel}$ and $\dot{\mathbf{v}}_{rel} = \mathbf{a}_r - \mathbf{a}_o$.

The time derivative is given by:

$$\dot{h}(x, u) = \gamma_d\dot{d} + \gamma_a(\dot{\cos}\,\alpha - \dot{\Phi}). \tag{13}$$

We expand the terms individually:

$$\dot{d} = \frac{\mathbf{p}_{rel}^\top(-\mathbf{v}_{rel})}{d}, \tag{14}$$

$$\frac{d}{dt}(\cos\alpha) = \frac{R^2}{d^3\cos\alpha}\dot{d} = -\frac{R^2(\mathbf{p}_{rel}^\top\mathbf{v}_{rel})}{d^4\cos\alpha}. \tag{15}$$

The derivative of the alignment metric Φ is:

$$\dot{\Phi} = \frac{\dot{N}D - N\dot{D}}{D^2}, \text{ with } N = \mathbf{p}_{rel}^\top\mathbf{v}_{rel},\ D = d\nu. \tag{16}$$

Expanding $\dot{N}$ and $\dot{D}$:

$$\dot{N} = -\mathbf{v}_{rel}^\top\mathbf{v}_{rel} + \mathbf{p}_{rel}^\top(\mathbf{a}_r - \mathbf{a}_o), \tag{17}$$

$$\dot{D} = \dot{d}\nu + d\frac{\mathbf{v}_{rel}^\top(\mathbf{a}_r - \mathbf{a}_o)}{\nu}. \tag{18}$$

By substituting $\mathbf{a}_r = f_a(x) + g_a(x)u$ into these expressions, we can isolate the control input u.

D. Affine Formulation

We rearrange $\dot{h}(x, u)$ into the standard affine form:

$$\dot{h}(x, u) = L_f h(x) + L_g h(x)u. \tag{19}$$

The control-dependent Lie derivative $L_g h(x)$ is derived from the $\mathbf{a}_r$ terms in $\dot{\Phi}$:

$$L_g h(x) = -\gamma_a\left[\frac{\mathbf{p}_{rel}^\top}{d\nu} - \frac{(\mathbf{p}_{rel}^\top\mathbf{v}_{rel})\mathbf{v}_{rel}^\top}{d\nu^3}\right]g_a(x). \tag{20}$$

All remaining terms independent of u (including obstacle acceleration $\mathbf{a}_o$) are collected into $L_f h(x)$. This analytic derivation allows the VO-CBF to be used directly in quadratic programs or nonlinear MPC.

Remark 1 (Relative Degree Reduction): Standard position-based CBFs, e.g., $h_{pos}(p) = \|p_{rel}\| - R$, have a relative degree of 2 for force-controlled systems ($\ddot{p} = u$). Because the first time derivative depends solely on velocity, the Lie derivative along the control vector field is zero ($L_g h_{pos} \equiv 0$). Thus, the control input does not appear until the second derivative, requiring complex High-Order CBF techniques.

In contrast, the proposed VO-CBF (10) explicitly includes the velocity terms $\mathbf{v}_{rel}$. Evaluating its first time derivative (11), the velocity alignment components ensure that $\frac{\partial h}{\partial v} \neq 0$, the Lie derivative evaluates to $L_g h(x) \neq 0$. Consequently, the control input u explicitly appears in the first time derivative, strictly reducing the constraint to relative degree 1. This formulation naturally simplifies the safety constraint design by avoiding the need for nested barrier functions or backstepping, while natively retaining the predictive lookahead of the collision cone.

IV. TIME-VARYING VO-CBF FOR CONSTRAINED OPTIMAL CONTROL

Transitioning a continuous-time safety barrier into a discrete-time digital controller inherently introduces structural complexity. To ensure a clear exposition, this section proceeds in three logical steps. First, we define the standard discrete-time NMPC mapping. Second, we introduce our decoupled Hybrid Constraint Architecture, which explicitly splits the barrier into a strict formulation for the immediate control step and an affine formulation for the prediction horizon. Finally, we provide the formal theoretical analysis bridging the continuous-time guarantees with this hybrid discrete-time implementation.

A. Discrete-Time NMPC Formulation

Let $\Delta t > 0$ be the sampling time. We denote the discrete time steps as $t_k = k\Delta t$ and the current measured state as $x_k := x(t_k)$. The control input is held constant between samples (Zero-Order Hold): $u(t) = u_k$ for $t \in [t_k, t_{k+1})$. At each time step k, the NMPC solves an optimization problem over a prediction horizon N. Let $x^{(i)}$ and $u^{(i)}$ denote the predicted state and control at step i relative to the current time k, where $x^{(0)} = x_k$. The discrete evolution is governed by the state transition map $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$, defined as the integration of the continuous dynamics (1) over the sampling interval:

$$\begin{aligned} x^{(i+1)} &= F(x^{(i)}, u^{(i)}) \\ &:= x^{(i)} + \int_0^{\Delta t} \Big(f(x(\tau)) + g(x(\tau))u^{(i)} \Big)\, d\tau, \end{aligned} \tag{21}$$

where $x(\tau)$ is the solution to (1) starting from $x(0) = x^{(i)}$ with constant input $u^{(i)}$. In implementation, F is typically approximated using numerical integration schemes such as Runge-Kutta (RK4).

B. Hybrid Constraint Structure

Existing discrete-time CBF-MPC approaches typically employ a homogeneous constraint architecture, applying either a strict nonlinear difference equation or a linearized approximation uniformly across the entire prediction horizon. However, homogeneous application forces a critical compromise between theoretical rigor and computational tractability. Applying a strict non-convex discrete formulation across the full horizon severely compromises real-time solver tractability, while applying an affine approximation uniformly introduces linearization errors at the initial time step, weakening the formal safety guarantee of the immediate control action executed by the robot. To resolve this fundamental trade-off, we propose a decoupled, hybrid constraint structure. By enforcing the exact nonlinear discrete invariance condition strictly at the immediate control step ($k = 0$) and utilizing the convex affine approximations solely for the prediction tail ($k > 0$), the proposed architecture guarantees mathematically exact, error-free safety for the physical control input while preserving the high-speed convex tractability of the predictive solver.

1) Strict Formulation (Discrete Invariance): For the immediate next step ($i = 0$), we enforce a discrete-time decay condition based on the finite difference of the barrier function. This captures the integrated effect of the control input over Δt:

$$h(x^{(1)}) \geq (1 - \kappa\Delta t)h(x^{(0)}) - \delta_{safe}, \tag{22}$$

where $\kappa > 0$ determines the convergence rate to the safe set boundary, and $\delta_{safe} \geq 0$ is a slack variable heavily penalized in the cost function to ensure recursive feasibility. This constraint is nonlinear in u but crucial for robust safety.

2) Affine Formulation (Tractability): For the remainder of the horizon ($i = 1, \ldots, N-1$), we enforce the standard continuous-time CBF condition approximated at the grid points:

$$L_f h(x^{(i)}) + L_g h(x^{(i)})u^{(i)} \geq -\kappa h(x^{(i)}) - \delta^{(i)}. \tag{23}$$

This constraint is linear in the decision variable u, preserving the convex substructure of the optimization problem (beneficial for SQP solvers) while guiding the planner toward safe regions.

A critical limitation of the Strict VO-CBF formulation (CBF-Strict) is its inherent geometric non-convexity. The boundary of a Velocity Obstacle collision cone naturally forms a non-convex feasible space. Utilizing this strict formulation across the entire NMPC prediction horizon forces the underlying Nonlinear Programming (NLP) solver to navigate this non-convexity at every predictive step. This severely increases computational latency and introduces a high risk of the solver converging to unsafe local minima during extreme dynamic interactions. This mathematical bottleneck is the primary motivation for our proposed Affine formulation (CBF-Affine). By linearizing the barrier via Lie derivatives at each predictive step, the non-convex collision cone is mathematically relaxed into a sequence of convex half-spaces. This structural transformation allows the NMPC to be solved utilizing Sequential Quadratic Programming (SQP) techniques, fundamentally eliminating local minima risks within the prediction horizon and guaranteeing robust, globally optimal convergence for the QP sub-problems.

C. Theoretical Analysis

We now provide the formal justification for the proposed scheme. We first introduce a Lemma relating the continuous derivative condition to the discrete evolution.

Lemma 1 (Sampled-Data Safety): Consider the system (1) under a zero-order hold control $u(t) = u_k$ for $t \in [t_k, t_{k+1}]$. Assume the barrier function $h(x)$ is globally Lipschitz with constant L_h, its time derivative $\dot{h}(x, u)$ is Lipschitz in x with constant $L_{\dot{h}}$, and the system dynamics are bounded such that $\|\dot{x}\| \leq M$ for all $x \in \mathcal{C}, u \in \mathcal{U}$. If the control input u_k satisfies the condition at the sampling instant:

$$\dot{h}(x_k, u_k) \geq -\kappa h(x_k) + \epsilon \tag{24}$$

with a robustness margin chosen as $\epsilon \geq (\kappa L_h + L_{\dot{h}})M\Delta t$, then the continuous-time differential inequality $\dot{h}(x(t), u_k) \geq -\kappa h(x(t))$ holds for all $t \in [t_k, t_{k+1}]$, guaranteeing the discrete decay:

$$h(x_{k+1}) \geq e^{-\kappa\Delta t} h(x_k). \tag{25}$$

Proof 1: Over the interval $t \in [t_k, t_{k+1}]$, the control input is held constant at u_k. The deviation of the state from the sampling instant is bounded by $\|x(t) - x_k\| \leq M(t - t_k) \leq M\Delta t$. Due to the Lipschitz continuity of h and $\dot{h}$, we have:

$$\begin{aligned} h(x_k) &\geq h(x(t)) - L_h M\Delta t \\ \dot{h}(x(t), u_k) &\geq \dot{h}(x_k, u_k) - L_{\dot{h}} M\Delta t. \end{aligned}$$

Substituting the sampling-instant condition $\dot{h}(x_k, u_k) \geq -\kappa h(x_k) + \epsilon$ into the derivative bound yields:

$$\dot{h}(x(t), u_k) \geq -\kappa h(x_k) + \epsilon - L_{\dot{h}} M\Delta t.$$

Replacing $h(x_k)$ with its lower bound relative to $h(x(t))$ gives:

$$\dot{h}(x(t), u_k) \geq -\kappa h(x(t)) + \kappa L_h M\Delta t + \epsilon - L_{\dot{h}} M\Delta t.$$

To ensure $\dot{h}(x(t), u_k) \geq -\kappa h(x(t))$ for all $t \in [t_k, t_{k+1}]$, we require the residual terms to be non-negative:

$$\epsilon - (\kappa L_h + L_{\dot{h}})M\Delta t \geq 0 \implies \epsilon \geq (\kappa L_h + L_{\dot{h}})M\Delta t.$$

If this condition is met, the comparison lemma implies $h(x(t)) \geq h(x_k)e^{-\kappa(t-t_k)}$. Evaluating at $t = t_{k+1}$ yields $h(x_{k+1}) \geq e^{-\kappa\Delta t}h(x_k)$.

Theorem 1 (Recursive Safety of Hybrid NMPC): Let $x_0 \in \mathcal{C}$ be the initial safe state ($h(x_0) \geq 0$). Assume the Optimization Problem (27) is recursively feasible. If the exact penalty slack variable $\delta_{safe} = 0$ at every time step k, then the closed-loop system state $x(t)$ remains in the safe set $\mathcal{C}$ for all $t > 0$.

Proof 2: The proof proceeds by induction. Base Case: At $k = 0$, $h(x_0) \geq 0$ by assumption. Inductive Step: Assume $h(x_k) \geq 0$. The NMPC applies the optimal control $u_k^* = u^{(0)}$. Since $\delta_{safe} = 0$, the Strict formulation (22) is strictly satisfied:

$$h(x^{(1)}) \geq (1 - \kappa\Delta t)h(x^{(0)}). \tag{26}$$

For a sufficiently small Δt such that $\kappa\Delta t \leq 1$, $h(x^{(0)}) \geq 0 \implies h(x^{(1)}) \geq 0$. By Lemma 1, the discrete satisfaction of this constraint ensures that the underlying continuous-time trajectory does not cross the zero-level set between samples, implying $h(x(t)) \geq 0$ for all $t \in [t_k, t_{k+1}]$. Thus, the system remains safe.

Remark 2 (Strict Safety vs. Recursive Feasibility): Theorem 1 establishes that strict forward invariance is guaranteed when $\delta_{safe} = 0$. However, in physical systems with strict actuation limits ($u_{min} \leq u \leq u_{max}$), ensuring the existence of a control input that strictly satisfies the barrier constraint at all times is challenging and can lead to solver failure. By introducing $\delta_{safe} \geq 0$ as a heavily penalized slack variable in the cost function, we relax the strict safety requirement into an exact penalty method [20]. This guarantees the recursive feasibility of the NMPC. During extreme dynamic transients where avoiding the obstacle violates input constraints, the solver will allow $\delta_{safe} > 0$, degrading gracefully to provide the "least unsafe" control action rather than aborting.

Remark 3 (Sampling Time Bound): The formal validity of the discrete-time safety guarantee inherently relies on a sufficiently small sampling time Δt. Analytically bounding Δt to guarantee strict forward invariance requires satisfying an inequality dominated by the Lipschitz constants of both the closed-loop system dynamics and the barrier gradient. Because computing exact online Lipschitz bounds for highly nonlinear vehicle dynamics and moving obstacles is computationally prohibitive, deriving a non-conservative analytic bound remains an open challenge for real-time discrete CBF implementations.

D. Optimization Problem Summary

We define the decision variables as the sequence of control inputs $\boldsymbol{u} \coloneqq [u^{(0)\top}, \dots, u^{(N-1)\top}]^\top$ and the sequence of slack variables for the affine constraints $\boldsymbol{\delta} \coloneqq [\delta^{(1)}, \dots, \delta^{(N-1)}]^\top$. Combining these elements, the NMPC solves the following problem at each step k:

$$\begin{aligned}
\min_{\boldsymbol{u}, \delta_{safe}, \boldsymbol{\delta}} \quad & \sum_{i=0}^{N-1} \left(\|x^{(i)} - x_{ref}\|_Q^2 + \|u^{(i)}\|_R^2 \right) \\
& + \|x^{(N)} - x_{ref}\|_P^2 + \rho_s \delta_{safe}^2 + \sum_{i=1}^{N-1} \rho \|\delta^{(i)}\|^2 \\
\text{s.t.} \quad & x^{(i+1)} = F(x^{(i)}, u^{(i)}), \quad x^{(0)} = x_k, \\
& u_{min} \leq u^{(i)} \leq u_{max}, \\
& h(x^{(1)}) \geq (1 - \kappa\Delta t)h(x^{(0)}) - \delta_{safe}, \qquad \text{(27a)} \\
& A_{cbf}^{(i)} u^{(i)} \geq b_{cbf}^{(i)} - \delta^{(i)}, \quad i = 1, \dots, N-1, \qquad \text{(27b)} \\
& \delta_{safe} \geq 0, \quad \delta^{(i)} \geq 0.
\end{aligned}$$

Constraint (27a) (Strict formulation) ensures robust one-step safety. Constraint (27b) (Affine formulation) enforces the linearized CBF condition derived in Section III, where the terms are defined as:

$$A_{cbf}^{(i)} \coloneqq L_g h(x^{(i)}), \quad b_{cbf}^{(i)} \coloneqq -\kappa h(x^{(i)}) - L_f h(x^{(i)}). \tag{28}$$

These definitions map the standard affine inequality $L_f h + L_g h u \geq -\kappa h$ directly into the linear constraint form required by the QP solver [21].

V. EXPERIMENTAL EVALUATION

In this section, we validate the proposed VO-CBF framework through comparative simulations and real robot experiments. The experiments are designed to highlight two distinct sources of safety violations: (1) unmodeled inertia in kinematic structures, and (2) unmodeled obstacle acceleration and estimation noise.

A. Baselines and Setup

To rigorously evaluate the proposed architecture, we benchmark the Hybrid VO-CBF against four alternative formulations across three distinct obstacle interaction scenarios (Linear, Sine, and Aggressive):

- *Baseline (Kinematic VO):* A standard Velocity Obstacle planner that projects desired velocity outside the collision cone, tracked by a low-level PD controller. It assumes instantaneous velocity tracking.
- *Position DTCBF:* A standard Discrete-Time Control Barrier Function MPC relying solely on spatial separation ($h_{pos}(x) = \|p_{rel}\| - R \geq 0$), serving as a strong dynamic baseline.
- *CBF-Strict:* NMPC enforcing the exact, nonlinear discrete-time barrier condition (22) at every step across the horizon.
- *CBF-Affine:* NMPC enforcing the linearized Lie derivative condition (23) at every step across the horizon.
- *CBF-Hyb:* The proposed Hybrid architecture, which enforces the Strict formulation at the first execution step ($k = 0$) and the Affine formulation for the prediction tail ($k > 0$).

For all experiments, we set the obstacle and robot radius to $0.3\,\text{m}$, the NMPC horizon to $N = 10$, the control frequency

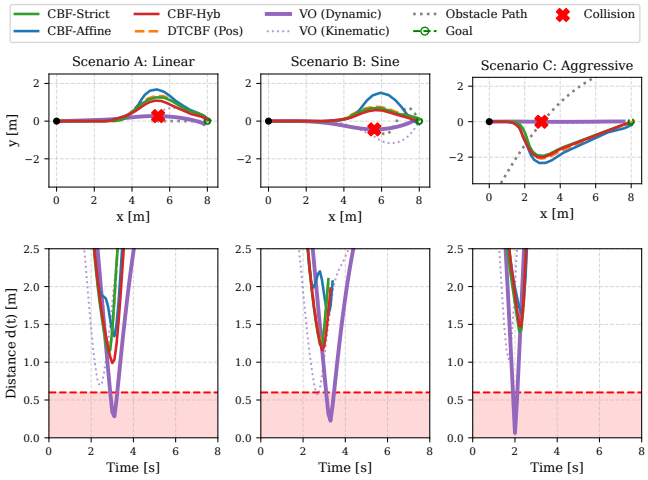

Fig. 1: Comparative evaluation across three dynamic obstacle interactions. Top row: Executed spatial trajectories. Bottom row: Minimum inter-agent distances (red shaded region indicates hard collision threshold at $0.60\,\text{m}$).

to $20\,\text{Hz}$, and the weighting parameters were set to $\gamma_d = 15.0$ and $\gamma_a = 1.5$. Following the theoretical rationale established in Section III, this strict order-of-magnitude difference guarantees that the absolute physical collision boundary dominates predictive velocity alignment during extreme close-proximity evasions.

B. Dynamic Obstacle Avoidance Scenarios

We evaluate the framework across three dynamic interactions (Fig. 1): a Head-On linear approach (Scenario A), a time-varying Sine maneuver (Scenario B), and an Aggressive high-acceleration cut-in (Scenario C).

1) Kinematic Failures vs. Dynamic Safety: Across all scenarios, the inherent flaws of pure kinematic planning are evident. Because the standard Kinematic VO ignores the robot's strict actuation limits ($F \leq F_{max}, \tau \leq \tau_{max}$), it assumes instantaneous velocity tracking. In Scenario A, the robot cannot generate the required centripetal force to track the planned evasive turn, physically drifting into the collision zone. In Scenarios B and C, where the obstacle executes sudden lateral maneuvers, the kinematic planner reacts too late; the resulting sudden steering intervention violates physical bounds, causing severe understeer and collisions (minimum distances dropping well below the $0.60\,\text{m}$ threshold).

Conversely, all CBF-based methods successfully maintain strict safety ($h(x) \geq 0$) across all scenarios. By natively embedding the dynamic actuation limits directly into the NMPC optimization, these formulations gracefully throttle velocity and initiate braking maneuvers well before the geometric collision boundary is breached.

2) Trajectory Efficiency and Conservatism: While all CBF variants ensure safety, their trajectory efficiencies differ drastically. The Position DTCBF relies solely on shrinking geometric distance and is blind to the directional collision cone. This limited anticipatory capability forces it to act highly conservatively; for instance, in Scenario B, it commands a wide, suboptimal detour ($8.22\,\text{m}$ path length). This efficiency penalty also affects the pure Affine formulation (CBF-Affine). While safe, relying uniformly on linear approximations over the prediction horizon induces artificial conservatism during dynamic lateral maneuvers, commanding an even wider detour ($9.17\,\text{m}$ path length).

The proposed Hybrid VO-CBF (CBF-Hyb) cleanly resolves this conservatism. By enforcing the mathematically exact nonlinear discrete decay (Strict formulation) solely at the immediate execution step, and anticipating the obstacle's acceleration ($\mathbf{a}_o$) via the Lie derivatives, the Hybrid architecture executes a much tighter, optimal evasion ($8.03\,\text{m}$ path length in Scenario B). Furthermore, as observed in Scenario C, while the Strict formulation (CBF-Strict) technically allows for a marginally tighter geometric curve than the Hybrid approach, this minimal spatial difference represents the deliberate, necessary trade-off. By accepting the slight geometric relaxation of affine half-spaces in the predictive tail, the Hybrid NMPC maintains convexity, solving in sub-milliseconds without sacrificing the rigorous safety of the immediate control action.

C. Computational Performance and Real-Time Feasibility

To validate the real-time viability of the proposed framework, aggregate computational metrics over 50 randomized trials are reported in Table I. Reporting both the standard deviation and maximum solve times is critical to validating real-time feasibility for the $50\,\text{ms}$ ($20\,\text{Hz}$) control loop. To accurately mirror real-world hardware deployment, the simulations were benchmarked using a Sequential Quadratic Programming (SQP) solver backed by the fast qpOASES engine.

Method	Success	Solve Time ($\mu \pm \sigma$)	Max Time
Kinematic VO	0%	N/A	N/A
Position DTCBF	**100%**	0.3 ± 0.2 ms	1.6 ms
CBF-Strict	**100%**	0.6 ± 0.4 ms	3.5 ms
CBF-Affine	**100%**	0.7 ± 0.4 ms	4.1 ms
CBF-Hyb	**100%**	$\mathbf{1.0 \pm 1.3}$ **ms**	**9.1 ms**

TABLE I: Quantitative Performance (50 randomized trials)

As demonstrated by the empirical data, the SQP architecture natively exploits the affine approximations, allowing the dynamic CBF formulations to execute with average solve times $\leq 1.0\,\text{ms}$. While the Hybrid method (CBF-Hyb) exhibits a microscopic overhead of a few hundred microseconds compared to the homogeneous formulations, this variance is strictly an artifact of constructing a heterogeneous constraint graph (mixing exact bounds at $k = 0$ with affine bounds at $k > 0$) within the high-level CasADi interface. In a compiled C++ deployment, these structural overheads vanish entirely. Most importantly, explicit variance quantification definitively proves that the proposed Hybrid approach comfortably bounds its worst-case solve time ($9.1\,\text{ms}$) well within the strict $50\,\text{ms}$ real-time limit, simultaneously ensuring trajectory optimality and a 100% dynamic safety rate.

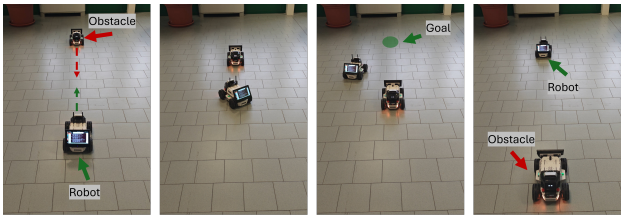

Fig. 2: Time-lapse sequence of the Head-on scenario. **(Left)** The robot (green marker) detects the moving obstacle via shared odometry. **(Center)** The Hybrid VO-CBF controller anticipates the future collision cone and initiates an early braking and steering maneuver ($t = 2.0s$). **(Right)** The robot successfully passes behind the obstacle, maintaining the safe distance R despite the dynamic constraints.

D. Robustness to Estimation Uncertainty

A critical assumption in dynamic collision avoidance is the availability of accurate obstacle state estimates (v_o, a_o). To quantify the impact of real-world estimation noise, we conducted a robustness analysis simulating severe tracking degradation. During the Aggressive cut-in scenario, which already places extreme demands on the dynamic solver, we injected zero-mean Gaussian noise ($\sigma = 1.0$) directly into the obstacle state predictions over the NMPC horizon. This emulates a severe hardware fault where the upstream Extended Kalman Filter (EKF) momentarily diverges during a high-speed interaction.

Method	Safety Rate	Min Distance	Path Length
Position DTCBF	0%	0.31 m (Collision)	8.12 m
CBF-Affine	0%	0.27 m (Collision)	13.08 m
CBF-Hyb	**100%**	**1.33 m**	**10.21 m**

TABLE II: Robustness Analysis with $\sigma = 1.0$

Because the Position DTCBF and the pure CBF-Affine methods rely entirely on numerical predictions over the receding horizon, the severe injected noise fundamentally corrupted their spatial forecasting. As detailed in Table II, the chaotic prediction tail caused both baseline solvers to synthesize highly unsafe control inputs, resulting in physical collisions (minimum distances dropping to 0.31 m and 0.27 m, respectively, failing the 0.60 m hard safety threshold).

In contrast, the proposed Hybrid approach (CBF-Hyb) structurally isolates the immediate physical reality from future uncertainty. By enforcing the mathematically exact nonlinear discrete decay (Strict formulation) at the immediate execution step ($k = 0$), the solver is anchored by a rigid, uncorrupted spatial barrier. Consequently, the Hybrid NMPC effectively absorbed the massive prediction variance in the horizon tail without compromising the immediate control action, maintaining a strict 100% safety rate (minimum distance of 1.33 m).

VI. REAL EXPERIMENTS

The experimental setup consists of two differential-drive Agilex Limo Pro robots operating under the ROS2 framework. One robot acts as the controlled agent, while the second serves as a dynamic obstacle. The Hybrid VO-CBF NMPC runs onboard the robot's embedded NVIDIA Jetson Orin Nano computer at a frequency of 20 Hz. The robot's low-level PID drivers track the optimal force commands.

Concerning the perception system, the robots share their odometry (position $\mathbf{p}$ and velocity $\mathbf{v}$) via the local ROS2 network. Crucially, the obstacle's acceleration $\mathbf{a}_o$, which is required for the VO-CBF Lie derivatives, is not directly measured. Instead, we exploit an onboard discrete-time Kalman Filter to estimate $\mathbf{a}_o$ from the noisy velocity stream. We evaluated the system in two canonical dynamic interaction scenarios: (1) a Head-On scenario where the obstacle moves toward the robot (we test linear and sinusoidal motion types), and (2) a Perpendicular Crossing scenario where the obstacle moves laterally across the robot's path. Consistent with the simulation results, the standard Kinematic VO baseline frequently resulted in collisions or emergency stops. Specifically, both in head-on and crossing scenarios, the kinematic planner failed, generating trajectories that attempted to avoid the obstacle too aggressively, leading to physical drift and collision.

In contrast, the Hybrid VO-CBF successfully navigated both scenarios in all trials. As illustrated in the time-lapse sequence in Figure 2 (head-on scenario), the integration of the estimated acceleration $\mathbf{a}_o$ allowed the controller to correctly anticipate the future rotation of the collision cone. Consequently, the robot initiated avoidance maneuvers earlier than the baseline, maintaining the safe distance R throughout the interaction. Notably, the total control cycle execution time, including ROS2 network latency, Kalman Filter state estimation, and the NMPC solver,remained within the $18-22$ ms range on the embedded hardware. This confirms that the Hybrid architecture effectively reduces the computational burden of the strict safety formulation, easily guaranteeing high-frequency real-time deployment (well within the 50 ms limit) on resource-constrained mobile robots.

VII. DISCUSSION

A key theoretical advantage of the VO-CBF is the reduction of the relative degree. Unlike position-based barriers, which typically have a relative degree of 2 for force-controlled systems, the VO-CBF formulation explicitly depends on relative velocity, resulting in a relative degree of 1. This allows the controller to actuate immediately against collision risks, preventing the inertial drift observed in the Kinematic Baseline.

To implement this practically, the proposed Hybrid architecture resolves the conflict between safety rigor and solver speed. As shown in the results section, the Hybrid approach enforces strict invariance only for the immediate step ($k = 0$) while using efficient affine approximations for the prediction tail ($k > 0$). This strategy maintains the total control cycle well within the $18 - 22$ ms range, successfully decoupling the immediate safety certificate from the computational burden of the long-horizon planning problem.

A. Limitations and Scalability

Despite these advantages, three key limitations exist. First, the predictive tail relies on the observability of the obstacle's

acceleration $\mathbf{a}_o$. While Section V demonstrated the Hybrid formulation's strong resilience to severe momentary estimation noise, sustained or unbounded tracking divergence could eventually degrade the predictive validity of the barrier.

Second, the strict formulation of the VO-CBF introduces a geometrically non-convex constraint. While our Hybrid architecture successfully mitigates local minima risks by utilizing convex affine relaxations for the prediction horizon ($k > 0$), the strict non-convexity at the immediate execution step ($k = 0$) could still challenge the NLP solver in highly cluttered environments.

Finally, while the proposed Hybrid VO-CBF framework successfully ensures real-time tractability and dynamic safety, its current validation is limited to single-obstacle interactions. Theoretically, the architecture scales to multi-agent scenarios by appending additional barrier constraints for each dynamic obstacle. However, synthesizing multiple simultaneous CBFs introduces the challenge of ensuring recursive feasibility, specifically, guaranteeing that the intersection of all dynamic safe sets remains strictly non-empty at all times. Proving the existence of a non-empty intersection in dense, uncoordinated multi-agent environments remains an open challenge.

To practically mitigate these scalability bottlenecks, concrete future strategies will include implementing distance-based active-set filtering to cap the number of simultaneously evaluated constraints, and exploring Decentralized MPC (DMPC) schemes where agents share planned trajectories to guarantee non-conflicting safe sets.

VIII. CONCLUSION

This work presented the VO-CBF, a control framework that embeds geometric collision cones directly into a dynamic safety barrier. By integrating this formulation into a hybrid NMPC architecture, we achieved a rigorous balance between the formal safety guarantees of discrete-time invariance and the computational tractability required for real-time control. Experimental validation on both simulated and real embedded platforms demonstrated that the VO-CBF effectively handles the inertial drift and dynamic obstacles that cause standard kinematic planners to fail. Furthermore, the Hybrid architecture cleanly resolves the artificial conservatism of standard position-based dynamic baselines and exhibits critical resilience to severe estimation noise. While current results are highly promising for single-agent dynamic avoidance, future efforts will focus on addressing the scalability challenges of dense multi-agent environments via decentralized optimization and active-set filtering.

REFERENCES

[1] F. Trotti, A. Farinelli, and R. Muradore, "An online path planner based on pomdp for uavs," in *2023 European Control Conference (ECC)*. IEEE, 2023, pp. 1–6.

[2] F. Trotti, D. Rigo, and R. Muradore, "Geometric methods for aircraft planning and control," *Robotics and Autonomous Systems*, p. 105181, 2025.

[3] F. Trotti, A. Farinelli, and R. Muradore, "Towards aircraft autonomy using a pomdp-based planner," in *2024 American Control Conference (ACC)*. IEEE, 2024, pp. 2399–2404.

[4] ——, "A markov decision process approach for decentralized uav formation path planning," in *2024 European Control Conference (ECC)*. IEEE, 2024, pp. 436–441.

[5] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, 2016.

[6] 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)*. IEEE, 2019, pp. 3420–3431.

[7] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," *American Control Conference (ACC)*, pp. 322–328, 2016.

[8] L. Wang and A. D. Ames, "Safety-critical model predictive control with control barrier functions," in *American Control Conference (ACC)*, 2017, pp. 5213–5219.

[9] W. Xiao and C. Belta, "Control barrier functions for systems with high relative degree," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4055–4072, 2019.

[10] T. G. Molnar, A. Singletary, G. Orosz, and A. D. Ames, "Safety-critical control with input constraints," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 3, pp. 1196–1211, 2021.

[11] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical hybrid systems," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 4998–5005.

[12] J. Zeng, B. Zhang, and K. Sreenath, "Safety-critical model predictive control with discrete-time control barrier function," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3553–3569, 2021.

[13] K. P. Wabersich and M. N. Zeilinger, "Predictive control barrier functions: Enhanced safety mechanisms for learning-based control," *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2638–2651, 2022.

[14] D. Wilkie, J. van den Berg, and D. Manocha, "Generalized velocity obstacles," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009, pp. 5573–5578.

[15] J. Snape, J. van den Berg, S. J. Guy, and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Trans. Robotics*, vol. 27, no. 4, pp. 696–706, 2011.

[16] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The international journal of robotics research*, vol. 17, no. 7, pp. 760–772, 1998.

[17] F. Vesentini, D. Meli, N. Sansonetto, L. Di Persio, and R. Muradore, "Dynamic movement primitives with control barrier functions for constrained trajectory planning," *IEEE Robotics and Automation Letters*, 2025.

[18] H. K. Khalil and L. Praly, "High-gain observers in nonlinear feedback control," *International Journal of Robust and Nonlinear Control*, vol. 24, no. 6, pp. 993–1015, 2014.

[19] A. Gelb *et al.*, *Applied optimal estimation*. MIT press, 1974.

[20] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, "Multi-layered safety for legged robots via control barrier functions and model predictive control," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8352–8358.

[21] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.