A Comprehensive Review on Control Barrier Functions: Uncertainty Handling, Design Optimization, and Feasibility Analysis

Shanshan Zhao, Zheng Yan, Tingwen Huang, *Fellow, IEEE*, and Shiping Wen, *Senior Member, IEEE*

***Abstract*—Control Barrier Functions (CBFs) provide a rigorous framework for enforcing safety in control-affine systems by ensuring system states remain within predefined safe sets. However, practical deployment faces fundamental challenges that limit real-world applicability. This review analyzes recent progress in CBF methodologies across three interconnected domains: uncertainty handling, structural optimization, and feasibility assurance. For uncertainty, we distinguish strategies tailored to unknown dynamics, modeling discrepancies, and dynamic environments, spanning robust theoretical methods and learning-based approaches. For structural design, we examine class-$\mathcal{K}$ function selection, parameter tuning, and advanced modifications that jointly address conservatism and feasibility. For feasibility, we identify root causes of CBF-QP infeasibility and survey solution strategies including constraint relaxation, structural redesign, and mathematical guarantees. By synthesizing these directions into a unified framework, this review highlights key interdependencies and outlines future research opportunities for advancing CBF-based safety-critical control in robotics, autonomous systems, and beyond.**

***Index Terms*—Control Barrier Functions, Safety-Critical Control, Uncertainty Handling, Structural Optimization, Feasibility Analysis**

I. Introduction

Safety assurance is a fundamental requirement in modern control systems, particularly in domains such as autonomous driving [1, 2], robotics [3], and multi-agent coordination [4]. These systems must operate in dynamic and uncertain environments where failures may cause catastrophic consequences. Control Barrier Functions (CBFs) have emerged as a mathematically rigorous tool for safety-critical control, enabling real-time constraint enforcement while maintaining control flexibility [5]. Through optimization-based formulations, CBFs define admissible control sets that guarantee forward invariance of safe states, establishing themselves as a cornerstone methodology for safety-critical applications. However, despite their theoretical appeal, practical deployment of CBFs faces significant limitations, leaving a noticeable gap between theoretical advances and real-world implementation.

This gap primarily arises from three interconnected challenges: (i) *uncertainty handling*, since real systems inevitably deviate from idealized models; (ii) *structural optimization*, as CBF design choices strongly influence both conservatism and feasibility; and (iii) *feasibility assurance*, because CBF-QP formulations may fail to admit feasible solutions under physical constraints. These challenges are not isolated but deeply coupled—uncertainty complicates structural design, conservative structures hinder feasibility, and infeasibility limits safety enforcement. Addressing them coherently is essential to advance CBF theory toward reliable real-world deployment.

These challenges span multiple research directions. Uncertainty handling addresses real-world dynamics affected by unknown components, modeling errors, and time-varying disturbances [6, 7], with approaches divided between robust theoretical methods providing guarantees under bounded uncertainties [8] and learning-based strategies leveraging data for adaptation [9, 10]. Structural optimization tackles the inherent coupling between conservatism and feasibility, evolving from fixed class-$\mathcal{K}$ functions toward systematic parameter optimization [11] and structural redesigns such as adaptive or auxiliary-variable formulations [12, 13]. Feasibility assurance addresses infeasibility in CBF-QPs when constraints and input bounds yield empty feasible sets [14, 15], with solutions spanning constraint relaxation, backup mechanisms, and rigorous feasibility guarantees [16], though scalability and real-time solvability remain critical bottlenecks.

This review systematically examines these three domains and their interconnections. Unlike prior surveys that treat them independently, we develop a unified framework linking uncertainty characteristics, structural design choices, and feasibility constraints. Section II analyzes uncertainty handling approaches, distinguishing

The statements made herein are solely the responsibility of the authors. *(Corresponding author: Shiping Wen.)*

S. Zhao, Z. Yan, and S. Wen are with the Australian AI Institute, Faculty of Engineering and Information Technology, The University of Technology Sydney, Sydney, NSW 2007, Australia (e-mail: shanshan.zhao@student.uts.edu.au; yan.z01@foxmail.com; shiping.wen@uts.edu.au). T. Huang is with Faculty of Computer Science and Control Engineering, Shenzhen University of Advanced Technology, Shenzhen 518055, China (huangtw2024@163.com).

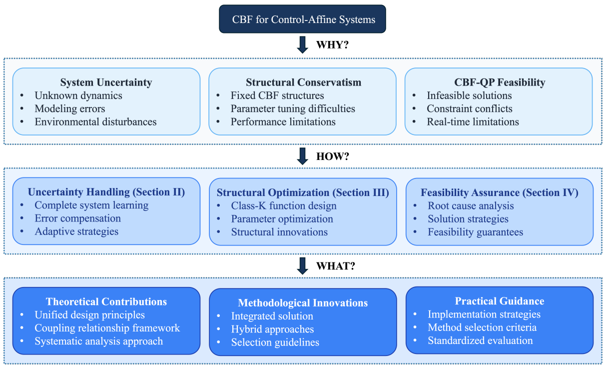

Figure 1: Unified view linking uncertainty, design, and feasibility in CBF-based control.

robust and learning-based methods. Section III explores structural optimization, covering class-$\mathcal{K}$ functions, parameter tuning, and structural redesign. Section IV investigates feasibility in CBF-QPs, examining root causes and solution strategies. Section V concludes with key findings, research gaps, and future directions. Figure 1 illustrates this unified framework, which highlights why bridging these challenges is essential for realizing the full potential of CBF-based safety-critical control.

II. Uncertainty Handling in CBFs

Control-affine systems, widely used in robotics, autonomous driving, and multi-agent coordination, are typically represented as

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

where $x \in \mathbb{R}^n$ is the state and $u \in \mathbb{R}^m$ the control input. CBFs enforce safety by maintaining system trajectories within a safe set $\mathcal{C} = \{x \mid h(x) \geq 0\}$ [5], provided the following condition holds: $\frac{\partial h}{\partial x}(f(x) + g(x)u) \geq -\alpha(h(x))$, with $\alpha(\cdot)$ a class-$\mathcal{K}$ function. Since CBF constraints explicitly depend on the system dynamics $f(x)$ and $g(x)$, their reliability is strongly affected by modeling uncertainty. Inaccuracies from unmodeled dynamics, parameter variations, sensor noise, or external disturbances can compromise safety guarantees, leading to conservatism or even violations. To address this, uncertainty handling in CBF-based control can be grouped into three categories: (i) Complete system learning, when both $f(x)$ and $g(x)$ are unknown; (ii) Error compensation, when nominal models exist but contain discrepancies; (iii) Adaptive policy optimization, when operating in dynamic, time-varying environments.

A. *Complete System Learning for Unknown Dynamics*

Historically, early CBF-based control relied on precisely known analytical models to construct explicit safety constraints. As robotic and autonomous systems became increasingly complex, such model dependence limited applicability in real-world scenarios where analytical modeling is infeasible. To overcome this, research has progressively expanded from data-intensive supervised learning to include probabilistic and model-free frameworks that learn dynamics and safety structures directly from experience.

Early data-driven approaches leveraged supervised learning paradigms through two complementary directions: expert demonstrations and neural networks. Expert demonstrations extract CBF structures from human-provided safe trajectories [17], offering strong initial safety guarantees but remaining constrained by data availability and quality. Neural networks (NNs) learn complex dynamics and safety structures directly from data, with representative studies including [18] for formal safety verification through convex relaxation and [19] for joint learning of dynamics and CBFs via Neural ODEs. While both methods provide pathways to unknown system learning, expert demonstrations face limited generalization beyond training scenarios, and neural methods require substantial training data and face challenges in formal safety verification.

To address data scarcity and provide principled uncertainty quantification, Gaussian Process Regression (GPR) offers a probabilistic framework that balances flexibility with formal safety guarantees. When both $f(x)$ and $g(x)$ are unknown, they can be modeled as Gaussian processes with posterior uncertainties incorporated into modified CBF constraints [20]. This enables explicit tuning of conservatism–feasibility trade-offs through confidence bounds, where higher uncertainty leads to more conservative safety margins. Recent extensions include matrix-variate GPs for structured learning [21], event-triggered updates for computational efficiency [22], and applications in robotic motion planning [23]. GPR provides formal safety guarantees while adapting to modeling requirements, though scalability in high-dimensional systems remains challenging.

Moving toward model-free paradigms, Reinforcement Learning (RL) represents an alternative that bypasses explicit dynamics estimation, instead refining control policies through trial-and-error interaction [24, 25]. Safety constraints can be embedded into RL objectives through penalty terms on barrier function violations, constrained optimization formulations, or safety shielding mechanisms. Foundational works established learning frameworks for barrier functions [26, 27], while recent studies proposed unified CLF–CBF–RL architectures [28]. While RL provides the greatest adaptability to unknown and changing environments, maintaining safety during exploration remains a fundamental challenge, as trial-and-error learning may violate constraints before policy convergence.

Overall, these approaches form a historical and methodological continuum from model-guided to model-free learning. Expert demonstrations require substantial prior knowledge but offer initial safety, NNs provide expressive capacity for complex dynamics, GPR balances flexibility with uncertainty quantification, and RL achieves maximum adaptability through direct interaction. This progression reflects a fundamental shift from data-intensive supervised learning toward adaptive frameworks that trade deterministic guarantees for flexibility in unknown environments. The central trade-off spans rigorous safety guarantees versus adaptability to uncertain dynamics, with future research focusing on hybrid frameworks that unify these advantages for real-time safety-critical applications.

B. Error Compensation for Modeling Discrepancies

Nominal models derived from physics or system identification inevitably contain inaccuracies due to unmodeled effects, parameter variations, noise, or external disturbances. If left unaddressed, such discrepancies can degrade CBF performance, causing conservatism or even safety violations. Historically, the study of uncertainty compensation within CBF frameworks has evolved through three main stages: early robust formulations based on worst-case guarantees, observer-based compensation that reconstructs disturbances in real time, and probabilistic or learning-based approaches that adapt to uncertainty statistically. Each stage reflects a distinct trade-off between conservatism, adaptability, and formal safety guarantees.

Early developments treated uncertainty as bounded disturbances and employed robust CBF formulations to guarantee safety under worst-case conditions [29]. These methods established rigorous safety certificates by ensuring system states remain within a subset of the pre-defined safe set, even under maximum anticipated disturbances. A key conceptual advance was the introduction of Input-to-State Safety (ISSf) [30], which introduced projection-to-state safety: rather than bounding high-dimensional disturbances explicitly, ISSf projects disturbance effects onto the CBF condition itself, enabling safety guarantees through scalar bounds without requiring full disturbance characterization. While these methods provided deterministic safety guarantees without requiring data, they often produced overly conservative control due to loose or uncertain disturbance bounds.

To overcome the conservatism of worst-case robust designs, observer-based methods emerged as an intermediate paradigm that combines robustness with adaptivity. Disturbance observer-based control integrates estimation and compensation mechanisms directly into the CBF framework, enabling real-time uncertainty reconstruction rather than static bounding. Early works introduced lumped disturbance estimates into CBF dynamics [31], demonstrating enhanced robustness through online compensation. This direction evolved toward tunable observer-based CBFs [32], which quantified estimation error impacts via ISSf principles, and nonlinear observers were extended to sampled-data systems [33]. Measurement error compensation [34], extended state observers [35], and environmental CBFs [36] further broadened applicability by embedding observers into noisy or partially observed systems. Together, these efforts established the foundation for real-time adaptive CBF control but were still largely restricted to matched disturbances.

Building upon these observer-based foundations, recent research has turned toward handling unmatched disturbances, where uncertainties affect safety constraints through channels different from the control input. When input relative degree differs from disturbance relative degree, disturbances and their derivatives can adversely impact both CBF dynamics and control performance through multiple channels. Advanced disturbance observers [37] and higher-order formulations such as the disturbance-rejection CBF [38] have demonstrated robust performance recovery while maintaining safety guarantees. This evolution integrates disturbance estimation with higher-order CBF structures, ensuring safety even under dynamically coupled uncertainty effects. Related works such as estimation error-based CBFs [39] and invariant-tube formulations [40, 41] further extend these ideas to more general nonlinear systems.

In parallel with observer-based developments, probabilistic and learning-based methods have been introduced to achieve uncertainty compensation through statistical adaptation rather than deterministic estimation. GPR models residual errors $\Delta(x, u)$ as Gaussian processes [42, 43], enabling adaptive confidence-based safety margins that balance conservatism and feasibility. Recent work demonstrates improved data efficiency [44] and event-triggered learning [45]. Neural network approaches [46–48] further generalize this trend by learning high-dimensional error patterns and even co-adapting CBF structures with system dynamics. In contrast to observer-based methods that rely on deterministic estimates, these learning-based strategies provide probabilistic safety adaptation but typically trade off formal guarantees for empirical flexibility.

Overall, the development of uncertainty compensation in CBFs demonstrates a clear historical trajectory: from worst-case robustness ensuring strict but conservative safety, through adaptive observer-based estimation achieving real-time compensation, to probabilistic and neural learning methods introducing flexibility and scalability. This evolution highlights an ongoing balance between formal safety guarantees and adaptive performance—a core theme in advancing safe control under modeling discrepancies.

C. Adaptive Policy Optimization for Dynamic Environments

Complementing the offline robustness and error compensation strategies discussed previously, a third research direction emerged—adaptive policy optimization—aimed at maintaining safety in dynamic and non-stationary environments where uncertainties evolve over time. Safety-critical systems often face time-varying uncertainties caused by changing external conditions or interactive agents, such as autonomous vehicles in varying traffic or multi-agent teams with dynamic coordination. Unlike static modeling errors, these environments require real-time adaptation to preserve safety under distributional shifts.

Early studies focused on risk-aware and probabilistic formulations, extending robust CBFs to stochastic settings. Probabilistic CBFs introduced safety guarantees under random disturbances [7, 8], marking a shift from deterministic guarantees to probabilistic risk bounds. They offer flexibility when uncertainty distributions are known but degrade in poorly characterized or fast-changing environments.

Building on these foundations, research expanded to multi-agent coordination, where uncertainty stems from inter-agent coupling rather than noise. Distributed CBF frameworks [49, 50] enforced cooperative safety through local constraints and communication-based optimization, representing a step from single-agent robustness to distributed adaptability, though scalability and real-time performance remain open challenges.

The latest stage introduced RL as a model-free paradigm for adaptive policy optimization in unknown or time-varying environments. RL-based safe control integrates CBF constraints into learning via penalties, constrained optimization, or shielding [26–28]. Recent progress includes event-triggered adaptation [51], distributed coordination [52], and scalable neural implementations [53]. While RL offers strong adaptability, it reintroduces the safety–exploration dilemma, as policies may violate constraints during training.

Together, these developments trace a historical trajectory from probabilistic risk-aware control to distributed multi-agent coordination and finally to RL-based adaptive learning. Each stage enhances adaptability but relaxes formal guarantees. Risk-aware methods suit stochastic yet stationary settings, multi-agent CBFs handle structured interactions, and RL adapts to unknown or rapidly changing dynamics but demands careful supervision.

Overall, uncertainty handling in CBFs has evolved from deterministic robustness to adaptive learning. When dynamics are unknown, learning-based approaches such as expert demonstrations, neural networks, Gaussian processes, and RL jointly infer dynamics and safety structures. When nominal models contain discrepancies, robust and observer-based methods compensate modeling errors through estimation and correction. In dynamic environments, adaptive policy optimization extends these principles to real time, enabling safe control under evolving uncertainties. These three stages—system learning, error compensation, and adaptive policy optimization—form the conceptual backbone of uncertainty-aware CBFs, highlighting the trade-off between formal guarantees and adaptive flexibility. Future progress will likely arise from hybrid frameworks that unify theoretical rigor with data-driven adaptability for complex, dynamic systems.

III. Design and Structural Optimization of CBFs

CBFs provide a formal mechanism for enforcing safety constraints in dynamic systems by defining admissible sets of control inputs through barrier-based inequality constraints. Over time, the design philosophy of CBFs has evolved through three main stages, each addressing the limitations of its predecessor. The first stage (function-level design) established an analytical foundation based on fixed class-$\mathcal{K}$ functions, which were tractable but often conservative or infeasible. The second stage (parameter optimization) focused on systematically tuning these functions through analytical or data-driven methods, yet remained limited by the fixed structure of CBF formulations. The third stage (structural redesign) introduced architectural innovations—such as auxiliary variables, backup safety layers, and compositional frameworks—to enhance feasibility and adaptability through structural flexibility rather than parameter adjustment. This section reviews these developments in chronological order, emphasizing how each evolution mitigated the conservatism–feasibility coupling in CBF-based control.

A. Role of the Class-$\mathcal{K}$ Function in CBF Constraints

The class-$\mathcal{K}$ function $\alpha(\cdot)$ in the CBF condition

$$\dot{h}(x,u) = \frac{\partial h}{\partial x}(f(x) + g(x)u) \geq -\alpha(h(x))$$

defines how strictly the system must approach or maintain the safe set, controlling the balance between conservatism and aggressiveness. Early formulations used a simple linear function $\alpha(h) = \gamma h$, which ensured analytical simplicity but imposed a rigid proportional relationship between $h(x)$ and its derivative. Large γ values enforced fast convergence but could make the QP infeasible when control limits or multiple constraints conflicted, while small γ values weakened safety enforcement. To address this rigidity, nonlinear forms—such as power, logarithmic, and exponential functions—were introduced, allowing the constraint strength to vary with the state and producing smoother transitions [54–56]. Later, adaptive or time-varying formulations [11, 57] adjusted γ dynamically based on environmental uncertainty or risk. For systems where the control input affects the barrier condition only after several derivatives, high-order CBFs (HOCBFs) [58] generalized this idea by differentiating the barrier constraint until the control appeared explicitly, extending applicability to higher-order systems. Together, these developments show how CBF design evolved from fixed linear constraints to nonlinear, adaptive, and high-order formulations that better balance safety and feasibility.

B. *Optimization Strategies for CBF Constraints*

While class-$\mathcal{K}$ design improved flexibility, it raised the question of how to tune parameters systematically. Manual selection was often problem-specific and brittle, motivating optimization frameworks that could calibrate parameters automatically from models or data. Analytical optimization methods, such as convex programming, Lyapunov-based tuning, and sum-of-squares formulations [59, 60], provided theoretical guarantees but depended heavily on accurate dynamics. To reduce model dependence, probabilistic and data-driven approaches were introduced, including Bayesian optimization and GPR [21, 61], which offered uncertainty-aware tuning, and RL frameworks [55, 62, 63], which adapted CBF parameters through reward feedback and safe exploration. These methods improved adaptability but also revealed a deeper limitation: parameter-level tuning cannot fully resolve infeasibility when multiple CBFs or actuator constraints conflict. This realization motivated the shift from tuning within fixed structures to redesigning the CBF architecture itself.

C. *Structural Optimization of CBFs*

Structural optimization extends CBF design beyond parameter tuning by modifying the internal formulation of the barrier condition. When inputs saturate or safety constraints interact nonlinearly, even optimally tuned CBFs can become infeasible. Adaptive and auxiliary-variable CBFs [12, 13] introduce auxiliary states that track constraint violations with bounded recovery guarantees. They transform hard safety conditions into soft formulations by embedding penalty dynamics or slack variables, maintaining QP feasibility while ensuring eventual return to the safe set. Backup CBFs [64] further extend this idea by establishing hierarchical safety layers. When the primary CBF becomes infeasible, a certified backup controller takes over and drives the system into a smaller safe subset, guaranteeing re-entry to the primary safe region within finite time [40, 65]. Learning-based frameworks such as imitation-driven barrier shaping [17] and safe exploration policies [66] adapt the CBF structure from experience, while compositional and switching designs [15, 67] combine multiple barriers through logical or mode-dependent coordination. These approaches enhance scalability in hybrid and multi-objective systems and demonstrate that structural flexibility, rather than numerical tuning, offers a principled way to address feasibility while maintaining safety guarantees.

The evolution of CBF research follows a clear trajectory from function design to parameter optimization and finally to structural innovation. Function-level shaping established tractable analytical foundations but was limited by conservatism. Parameter optimization improved adaptability yet could not eliminate structural infeasibility. Structural redesign incorporated flexibility directly into the architecture, enabling real-time feasibility and robustness. This progression reflects how the field has matured from theoretical safety formulations toward adaptive, architecture-aware safety control. The next chapter builds on this foundation by analyzing the feasibility of CBF-QP formulations and strategies to ensure reliable solvability in real-time safety-critical applications.

IV. Ensuring Feasibility in CBF-QP

The study of CBF–QP feasibility has evolved alongside the broader development of safety–critical control. Early works established the theoretical foundation for constraint enforcement [2, 68], but practical deployment quickly revealed feasibility failures under input saturation and model uncertainty. Subsequent studies recognized infeasibility not as an isolated numerical issue, but as a structural property of the CBF formulation itself. This realization motivated a research trajectory from constraint relaxation and auxiliary variables to redundancy-based and robustness-oriented designs such as Backup CBFs and Robust Policy CBFs.

CBF-QPs have become a widely adopted framework for safety-critical control, offering a principled way to enforce safety constraints while optimizing performance objectives. Unlike heuristic safety filters, CBF-QPs synthesize control inputs by solving the constrained optimization problem:

$$u^* = \arg\min_u \frac{1}{2} u^T H u + f^T u \tag{2a}$$

$$\text{s.t.} \quad L_f h(x) + L_g h(x) u \geq -\alpha(h(x)), \quad u \in \mathcal{U}, \tag{2b}$$

where H and f define the performance objective, (2b) enforces the CBF condition, and $\mathcal{U}$ denotes admissible control inputs constrained by actuator limits. Early CBF-QP implementations [2, 68] established theoretical foundations but revealed feasibility challenges when actuator limits conflicted with safety requirements. Systematic studies [14, 64] identified feasibility failures as a fundamental bottleneck, motivating structured approaches including constraint relaxation [29], learning-based adaptations [17, 69], mathematical feasibility guarantees [16, 70], and backup mechanisms [15, 64].

Despite its theoretical elegance, CBF-QP feasibility remains a central challenge in real-world deployment. The problem (2) may admit no solution when the feasible control set becomes empty due to conflicts between CBF constraints and bounded control inputs. This section examines the root causes of infeasibility, surveys solution strategies, and highlights implementation considerations for making CBF-QPs reliably solvable in real time.

A. Root Causes of CBF-QP Infeasibility

The infeasibility of problem (2) arises when the intersection of all constraints becomes empty, leaving no admissible control input. The origins of infeasibility can be grouped into three categories: (i) physical limitations that fundamentally restrict control authority, (ii) inappropriate CBF design that mismatches geometric safety definitions with dynamic feasibility, and (iii) improper parameter selection that renders otherwise sound formulations unsolvable.

The most basic cause stems from actuator and hardware limitations. Real systems operate under strict input bounds imposed by actuator saturation, power constraints, or mechanical safety limits. When safety constraints require control efforts beyond these limits, the feasible set collapses. [15] showed that multi-CBF compositions may demand control authority exceeding actuator capabilities, while [14] demonstrated that disturbances can entirely eliminate the feasible region in multi-agent coordination. This type of infeasibility reflects physical impossibility: no control input exists that can simultaneously satisfy all requirements. Resolution often requires relaxing safety constraints, redesigning tasks, or improving hardware capacity.

Poor CBF design can introduce infeasibility even when sufficient control authority exists. A common pitfall is constructing barrier functions directly from geometric safety boundaries without ensuring compatibility with system dynamics. In double-integrator systems, selecting high-relative-degree functions leads to derivative conditions that do not involve the control input, making the QP unsolvable. [64] noted that such design flaws push trajectories toward regions where no admissible control can restore safety. Learning-based CBF construction suffers similar issues: [17] highlighted that learned barriers may fit training data but fail to enforce derivative conditions in practice. The essence is a mismatch between geometric safety definitions and dynamic feasibility.

The third source arises when theoretically sound CBF candidates are paired with poorly chosen parameters. Tuning class-$\mathcal{K}$ functions, coordinating multiple CBFs, or updating adaptive parameters can introduce infeasibility if handled incorrectly. [71] showed that adaptive CBFs may fail when update rates are inappropriate, while [15] demonstrated that naive composition of individually feasible CBFs can yield collective infeasibility when input bounds are ignored. [69] emphasized that many methods verify CBF conditions but neglect to learn the feasibility domain, resulting in constraints that cannot be satisfied near boundaries. The root problem is numerical sensitivity: poor parameterization transforms otherwise valid CBFs into practically infeasible formulations.

These three categories are tightly coupled rather than independent. Physical limitations shrink the feasible design space, inappropriate CBF structures exacerbate parameter sensitivity, and poor parameter choices can mask structural flaws. Recognizing their interdependence is crucial for developing systematic solutions.

B. Solution Strategies for Feasibility

To mitigate the infeasibility of CBF-QPs, researchers have developed strategies targeting different root causes. These methods fall into three complementary categories: (i) constraint relaxation and learning approaches that modify or adapt safety constraints, (ii) structural design and redundancy mechanisms that improve CBF formulations, and (iii) mathematical feasibility guarantees that provide rigorous solvability conditions.

Constraint relaxation offers flexibility when strict formulations cannot be satisfied. Early unified CLF–CBF frameworks [2, 68] introduced slack variables to prioritize safety (CBF) over performance (CLF) when conflicts arise, establishing a foundation for handling infeasibility through constraint softening. Robust CBFs [29] extended this idea by balancing safety–performance trade-offs under bounded disturbances. Sum-of-squares programming [72] and projected dynamical systems [73] enable feasibility analysis across multiple constraints, though these approaches remain limited when the feasible set is fundamentally empty. Learning-based methods extend this work by refining safety constraints from data. Examples include GPR for uncertainty modeling [21, 43], reinforcement- and imitation-learning–based CBF construction [17, 27], and feasibility-aware neural networks [69]. While these methods enhance adaptability, they depend on data coverage and provide only probabilistic guarantees.

Structural modifications improve feasibility by refining the CBF formulation itself. Adaptive CBFs [12] dynamically tune parameters, zeroing CBFs [74] enable multi-robot collision avoidance, and disturbance-robust variants [38] integrate observers for enhanced robustness. Within this line, Backup Control Barrier Functions (bCBFs) represent a particularly influential family of approaches that directly target CBF-QP feasibility. The work in [64] presents the foundational formulation of the bCBF framework, which embeds an implicitly defined control-invariant set via a fixed backup policy and forward integration of system dynamics. This structure guarantees the solvability of CBF-QPs even under limited control authority and provides a rigorous foundation for feasibility assurance in safety-critical control. The study in [75] extends this idea to robotic systems with multiple constraints, introducing a rectified Backup CBF-QP that maintains feasibility under time-varying, time-invariant, and input constraints through predictive backup trajectories and greedy switching policies.

Building upon these foundations, subsequent works extend the bCBF framework to address the robustness–feasibility nexus under model uncertainty. The formulations in [40, 41] construct invariant tubes around nominal trajectories and incorporate online disturbance estimation, ensuring that control-invariant sets remain valid—and thus CBF-QPs remain feasible—even when system dynamics deviate from nominal assumptions. A closed-form bounded-input controller [76] removes the runtime quadratic program while analytically guaranteeing safety and feasibility. The formulation in [77] further generalizes the backup concept into a policy-based structure, where Robust Policy CBFs adaptively learn safe policies online and ensure feasible safety filtering under uncertain, input-constrained, and high–relative-degree systems. Together, these studies form a coherent progression from structural redundancy to robustness and computational tractability, bridging the gap between design-oriented and robustness-oriented feasibility strategies.

Rigorous approaches, in contrast, aim to establish mathematical conditions that guarantee solvability. Pointwise feasibility analysis [70] provides necessary and sufficient conditions using Gaussian–Process–based second-order cone programs, while sufficient-condition methods [78] dynamically reshape the feasible set through auxiliary constraints. Feasible-set–reshaping frameworks [14] ensure solvability in multi-agent systems via small-gain theory, and inverse-optimal safety filters [16] employ Hamilton–Jacobi–Isaacs equations to secure infinite-horizon feasibility. Although these methods provide strong theoretical guarantees, their computational costs limit real-time deployment.

Overall, the three categories complement one another: constraint relaxation emphasizes adaptability, structural design—particularly through backup and robustness-oriented formulations—improves practical feasibility, and mathematical guarantees offer strict solvability conditions. Integrating these perspectives yields a comprehensive toolkit for deploying feasible and reliable CBF-QPs in safety-critical robotic systems.

C. Implementation Considerations and Future Directions

While solution strategies demonstrate promising ways to address CBF-QP infeasibility, practical implementation faces major challenges. At the core lies a fundamental tension between theoretical guarantees and the practical requirements of real-time, large-scale deployment. The main obstacles are threefold. First, computational complexity remains a barrier: rigorous methods such as Hamilton–Jacobi–Isaacs formulations [16] and small-gain analysis frameworks [14] offer strong guarantees but incur prohibitive costs in embedded systems. Second, real-time implementation is demanding in high-speed applications, where learning-based adaptation and dynamic switching across CBF formulations introduce latency that may invalidate feasibility. Third, scalability issues arise in high-dimensional and multi-agent systems, where multi-CBF compositions [15] create combinatorial growth in constraints, burdening computation and increasing infeasibility risks.

Bridging the gap between theoretical rigor and practical deployment requires several advances. Structurally feasible CBF formulations that guarantee solvability by construction should integrate actuator limitations and safety margins directly into the design stage. To reduce computational burden, convex formulations and approximate solvers merit further study, with hierarchical approaches and hardware acceleration enabling efficient deployment. For real-time feasibility, predictive and adaptive schemes can anticipate infeasible states and adjust constraints proactively. Hybrid frameworks that combine structural feasibility-aware designs with data-driven adaptation may achieve both provable guarantees and real-time adaptability. Finally, scalability demands decentralized multi-agent frameworks, distributed optimization, and communication-efficient coordination protocols, supported by benchmarking tools to systematically evaluate trade-offs. Overall, the evolution of feasibility research—from recognizing infeasibility as a practical failure to formulating it as a structural and robustness problem—illustrates the ongoing convergence between theoretical rigor and real-time applicability in CBF-QP control.

V. Conclusion and Future Work

This review has presented a comprehensive analysis of CBF methodologies, focusing on three interconnected challenges that limit their practical deployment: uncertainty handling, structural design optimization, and feasibility in CBF-QPs. Unlike existing surveys that treat these aspects independently, our study emphasizes their interdependencies and develops a unified framework for advancing CBF-based safety-critical control.

A. Key Contributions and Insights

The contributions of this review can be summarized in three insights. First, uncertainty handling should be classified according to system unknowns rather than algorithmic categories, clarifying when robust control suffices and when learning-based methods are required: learning strategies for complete system unknowns, compensation methods for modeling discrepancies, and adaptive policies for dynamic environments. Second, structural optimization reveals that conservatism reduction and feasibility preservation are inherently coupled; while parameter tuning cannot resolve this coupling, structural redesign through auxiliary variables, backup mechanisms, and compositional frameworks enables dynamic architectures that jointly optimize safety and feasibility. Third, feasibility in CBF-QPs emerges from the interaction of physical limitations, CBF design, and parameter selection, requiring integrated solutions that combine constraint relaxation for adaptability, structural methods for robustness, and mathematical guarantees for rigorous solvability. Together, these insights establish unified design principles: (i) embed uncertainty considerations at the design stage, (ii) optimize conservatism and feasibility jointly, and (iii) balance theoretical rigor with computational scalability.

B. Research Challenges and Future Directions

Despite advances, several challenges remain in both design and implementation. For uncertainty handling, high-dimensional and multi-agent systems face computational bottlenecks, motivating distributed frameworks and hybrid robust–learning methods that dynamically adapt to changing conditions. Structural optimization still relies heavily on expert knowledge, suggesting future work on automated CBF synthesis using machine learning, particularly for distributed designs with communication-efficient and fault-tolerant coordination in multi-agent settings. Real-time feasibility remains a pressing issue, as mathematical guarantees such as feasible-set reshaping or inverse optimal safety filters often exceed computational limits; research should therefore emphasize predictive feasibility management combined with scalable solvers and hardware acceleration. Finally, standardized evaluation frameworks are urgently needed, as current studies rely on isolated case analyses, hindering systematic comparison. Developing benchmarks across uncertainty types, system scales, and application domains will accelerate progress.

In summary, advancing CBF methodologies requires integrated approaches that jointly address uncertainty, structural adaptability, and feasibility. The unified framework proposed here provides a foundation for such integration, guiding future research toward scalable, robust, and practically deployable safety-critical control systems.

References

[1] 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.

[2] 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.

[3] F. Ferraguti, C. T. Landi, A. Singletary, H.-C. Lin, A. Ames, C. Secchi, and M. Bonfè, "Safety and efficiency in robotics: The control barrier functions approach," *IEEE Robotics & Automation Magazine*, vol. 29, no. 3, pp. 139–151, 2022.

[4] D. Panagou, D. M. Stipanović, and P. G. Voulgaris, "Distributed coordination control for multi-robot networks using lyapunov-like barrier functions," *IEEE Transactions on Automatic Control*, vol. 61, no. 3, pp. 617–632, 2015.

[5] 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.

[6] X. Xu, J. W. Grizzle, P. Tabuada, and A. D. Ames, "Correctness guarantees for the composition of lane keeping and adaptive cruise control," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 1216–1229, 2017.

[7] A. Singletary, M. Ahmadi, and A. D. Ames, "Safe control for nonlinear systems with stochastic uncertainty via risk control barrier functions," *IEEE Control Systems Letters*, vol. 7, pp. 349–354, 2022.

[8] M. Black, G. Fainekos, B. Hoxha, D. Prokhorov, and D. Panagou, "Safety under uncertainty: Tight bounds with risk-aware control barrier functions," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 686–12 692.

[9] 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.

[10] M. Cohen and C. Belta, *Adaptive and Learning-Based Control of Safety-Critical Systems*. Springer, 2023.

[11] X. Wang, "Ensuring safety of learning-based motion planners using control barrier functions," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4773–4780, 2022.

[12] W. Xiao, C. Belta, and C. G. Cassandras, "Adaptive control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2267–2281, 2022.

[13] S. Liu, W. Xiao, and C. A. Belta, "Auxiliary-variable adaptive control barrier functions for safety critical systems," in *Proceedings of the 62nd IEEE Conference on Decision and Control (CDC)*, 2023, pp. 8602–8607.

[14] S. Wu, T. Liu, M. Egerstedt, and Z.-P. Jiang, "Quadratic programming for continuous control of safety-critical multiagent systems under uncertainty," *IEEE Transactions on Automatic Control*, vol. 68, no. 11, pp. 6664–6679, 2023.

[15] J. Breeden and D. Panagou, "Compositions of multiple control barrier functions under input constraints," in *2023 American Control Conference (ACC)*. IEEE, 2023, pp. 3688–3695.

[16] M. Krstic, "Inverse optimal safety filters," *IEEE Transactions on Automatic Control*, vol. 69, no. 1, pp. 16–31, 2023.

[17] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3717–3724.
[18] A. Mazouz, S. Dutta, and M. Althoff, "Safety guarantees for neural network dynamic systems via stochastic barrier functions," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10 367–10 373.
[19] V. Zinage and E. Bakolas, "Neural differentiable integral control barrier functions for unknown nonlinear systems with input constraints," in *2023 American Control Conference (ACC)*. IEEE, 2023, pp. 4680–4687.
[20] P. Jagtap, G. J. Pappas, and M. Zamani, "Control barrier functions for unknown nonlinear systems using gaussian processes," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3699–3704.
[21] V. Dhiman, M. J. Khojasteh, M. Franceschetti, and N. Atanasov, "Control barriers in bayesian learning of system dynamics," *IEEE Transactions on Automatic Control*, vol. 68, no. 1, pp. 214–229, 2021.
[22] A. Lederer, A. Begzadi'c, S. Hirche, J. Cort'es, and S. Herbert, "Safe barrier-constrained control of uncertain systems via event-triggered learning," *ArXiv*, vol. abs/2408.16144, 2024.
[23] J. Zhang, Y. Liu, X. Wu, and Q. Fan, "Safe control of robotic systems under nonparametric uncertainty using gaussian process regression," in *2024 International Conference on Advanced Robotics and Mechatronics (ICARM)*, 2024, pp. 777–782.
[24] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 411–444, 2022.
[25] D. Xu and M. Zhang, "Reinforcement-learning-based control of convectively unstable flows," *Journal of Fluid Mechanics*, vol. 954, p. A37, 2023.
[26] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, "Episodic learning with control lyapunov functions for uncertain robotic systems," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6878–6884.
[27] A. Taylor, A. Singletary, Y. Yue, and A. Ames, "Learning for safety-critical control with control barrier functions," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 708–717.
[28] J. Choi, F. Castaneda, C. J. Tomlin, and K. Sreenath, "Reinforcement learning for safety-critical control under model uncertainty, using control lyapunov functions and control barrier functions," in *Proceedings of Robotics: Science and Systems*, 2020.
[29] M. Jankovic, "Robust control barrier functions for constrained stabilization of nonlinear systems," *Automatica*, vol. 96, pp. 359–367, 2018.
[30] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 108–113, 2019.
[31] E. Daş and R. M. Murray, "Robust safe control synthesis with disturbance observer-based control barrier functions," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 5566–5573.
[32] A. Alan, T. G. Molnar, E. Daş, A. D. Ames, and G. Orosz, "Disturbance observers for robust safety-critical control with control barrier functions," *IEEE control systems letters*, vol. 7, pp. 1123–1128, 2022.
[33] J. Sun, J. Yang, and Z. Zeng, "Safety-critical control with control barrier function based on disturbance observer," *IEEE Transactions on Automatic Control*, vol. 69, no. 7, pp. 4750–4756, 2024.
[34] R. K. Cosner, A. W. Singletary, S. Coogan, and A. D. Ames, "Measurement-robust control barrier functions: Certainty in safety with uncertainty in state," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6786–6793.
[35] J. Chen, Z. Gao, and Q. Lin, "Robust control barrier functions for safe control under uncertainty using extended state observer and output measurement," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 8477–8482.
[36] V. Hamdipoor, N. Meskin, and C. G. Cassandras, "Safe control synthesis using environmentally robust control barrier functions," *European Journal of Control*, vol. 74, p. 100840, 2023.
[37] E. Daş and J. W. Burdick, "Robust control barrier functions using uncertainty estimation with application to mobile robots," *IEEE Transactions on Automatic Control*, 2025.
[38] X. Wang, J. Yang, C. Liu, Y. Yan, and S. Li, "Safety-critical disturbance rejection control of nonlinear systems with unmatched disturbances," *IEEE Transactions on Automatic Control*, 2024.
[39] Y. Wang and X. Xu, "Disturbance observer-based robust control barrier functions," in *2023 American Control Conference (ACC)*, 2023, pp. 3681–3687.
[40] D. E. Van Wijk, S. Coogan, T. G. Molnar, M. Majji, and K. L. Hobbs, "Disturbance-robust backup control barrier functions: Safety under uncertain dynamics," *IEEE Control Systems Letters*, 2024.
[41] D. E. van Wijk, E. Das, A. Alan, S. Coogan, T. G. Molnar, J. W. Burdick, M. Majji, and K. L. Hobbs, "Disturbance observers for robust backup control barrier functions," *arXiv preprint arXiv:2503.15734*, 2025.
[42] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*. Springer, 2003, pp. 63–71.
[43] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, "Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects and dynamics," in *2021 American Control Conference (ACC)*, 2021, pp. 3683–3690.
[44] J. J. Choi, F. Castaneda, W. Jung, B. Zhang, C. J. Tomlin, and K. Sreenath, "Constraint-guided online data selection for scalable data-driven safety filters in uncertain robotic systems," *IEEE Robotics and Automation Letters*, vol. 8, pp. 4567–4575, 2023.
[45] F. Castaneda, J. J. Choi, W. Jung, B. Zhang, C. J. Tomlin, and K. Sreenath, "Probabilistic safe online learning with control barrier functions," *arXiv preprint arXiv:2208.10733*, 2022.
[46] L. Wang, A. D. Ames, and M. Egerstedt, "Learning control barrier functions with high relative degree for safety-critical control," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 562–569.
[47] C. Dawson, S. Gao, and C. Fan, "Safe nonlinear control using robust neural lyapunov-barrier functions," in *2021 Conference on Robot Learning*. PMLR, 2021, pp. 1724–1735.
[48] T. Dai, D. Romeres, R. Dai, and A. D. Ames, "Learning a better control barrier function under uncertain dynamics,"

in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 3464–3471.

[49] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "Safe multi-agent interaction through robust control barrier functions with learned uncertainties," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 2052–2057.

[50] S. Yaghoubi and G. Fainekos, "Training neural network controllers using control barrier functions in the presence of disturbances," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3657–3662.

[51] B. Yan, Y. Sun, P. Shi, and C.-C. Lim, "Event and learning-based resilient formation control for multiagent systems under dos attacks," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 54, no. 8, pp. 4876–4886, 2024.

[52] B. Yan, P. Shi, C. P. Lim, Y. Sun, and R. K. Agarwal, "Security and safety-critical learning-based collaborative control for multiagent systems," *IEEE Transactions on Neural Networks and Learning Systems*, 2024.

[53] W. Bai, P. X. Liu, and H. Wang, "Neural-network-based adaptive fixed-time control for nonlinear multiagent non-affine systems," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 1, pp. 570–583, 2024.

[54] W. Xiao, C. A. Belta, and C. G. Cassandras, "Feasibility-guided learning for constrained optimal control problems," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 1896–1901.

[55] Z. Marvi and B. Kiumarsi, "Safe reinforcement learning: A control barrier function optimization approach," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 6, pp. 1923–1940, 2021.

[56] A. Agrawal and D. Panagou, "Safe control synthesis via input constrained control barrier functions," *IEEE Transactions on Automatic Control*, vol. 66, no. 7, pp. 3311–3318, 2021.

[57] H. Parwana, A. Mustafa, and D. Panagou, "Trust-based rate-tunable control barrier functions for non-cooperative multi-agent systems," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2222–2229.

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

[59] A. J. Taylor and A. D. Ames, "Adaptive safety with control barrier functions," in *2020 American Control Conference (ACC)*, 2020, pp. 1399–1405.

[60] A. Isaly, M. Ghanbarpour, R. G. Sanfelice, and W. E. Dixon, "On the feasibility and continuity of feedback controllers defined by multiple control barrier functions for constrained differential inclusions," in *2022 American Control Conference (ACC)*, 2022, pp. 5160–5165.

[61] F. Vinter-Hviid, C. Sloth, T. R. Savarimuthu, and I. Iturrate, "Safe contact-based robot active search using bayesian optimization and control barrier functions," *Frontiers in Robotics and AI*, vol. 11, p. 1344367, 2024.

[62] L. Lindemann, A. Robey, L. Jiang, S. Das, S. Tu, and N. Matni, "Learning robust output control barrier functions from safe expert demonstrations," *IEEE Open Journal of Control Systems*, 2024.

[63] M. Pereira, Z. Wang, I. Exarchos, and E. Theodorou, "Safe optimal control using stochastic barrier functions and deep forward-backward sdes," in *Conference on Robot Learning*. PMLR, 2021, pp. 1783–1801.

[64] Y. Chen, M. Jankovic, M. Santillo, and A. D. Ames, "Backup control barrier functions: Formulation and comparative study," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 6835–6841.

[65] P. M. Rivera, "Forward and control invariance analysis of backup control barrier function induced safe sets for online safety of nonlinear systems," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*. IEEE, 2024, pp. 8150–8157.

[66] P. Rabiee and A. Safari, "Safe exploration in reinforcement learning: training backup control barrier functions with zero training time safety violations," *arXiv preprint arXiv:2312.07828*, 2023.

[67] A. J. Taylor and A. D. Ames, "Adaptive safety with control barrier functions," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 1399–1405.

[68] 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*, 2014, pp. 6271–6278.

[69] W. Xiao, C. G. Cassandras, and C. A. Belta, "Learning feasibility constraints for control barrier functions," in *2023 European Control Conference (ECC)*. IEEE, 2023, pp. 1–6.

[70] F. Castaneda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, "Pointwise feasibility of gaussian process-based safety-critical control under model uncertainty," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 6762–6769.

[71] S. Liu, W. Xiao, and C. Belta, "Feasibility-guaranteed safety-critical control with applications to heterogeneous platoons," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*. IEEE, 2024, pp. 8066–8073.

[72] A. Isaly, M. Ghanbarpour, R. G. Sanfelice, and W. E. Dixon, "On the feasibility and continuity of feedback controllers defined by multiple control barrier functions for constrained differential inclusions," in *2022 American Control Conference (ACC)*, 2022, pp. 5160–5165.

[73] G. Delimpaltadakis, J. Cortés, and W. P. M. H. Heemels, "Continuous approximations of projected dynamical systems via control barrier functions," *IEEE Transactions on Automatic Control*, vol. 70, no. 1, pp. 681–688, 2025.

[74] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.

[75] D. Ko and W. K. Chung, "A backup control barrier function approach for safety-critical control of mechanical systems under multiple constraints," *IEEE/ASME Transactions on Mechatronics*, 2024.

[76] D. E. van Wijk, E. Das, T. G. Molnar, A. D. Ames, and J. W. Burdick, "Safety-critical control with bounded inputs: A closed-form solution for backup control barrier functions," *arXiv preprint arXiv:2510.05436*, 2025.

[77] L. Knoedler, O. So, J. Yin, M. Black, Z. Serlin, P. Tsiotras, J. Alonso-Mora, and C. Fan, "Safety on the fly: Constructing robust safety filters via policy control barrier functions at runtime," *IEEE Robotics and Automation Letters*, 2025.

[78] W. Xiao, C. A. Belta, and C. G. Cassandras, "Sufficient conditions for feasibility of optimal control problems using control barrier functions," *Automatica*, vol. 135, p. 109960, 2022.