Multidisciplinary Engineering Science Open

Received date: 23 Nov 2023
Revised date: 28 Dec 2023
Accepted date: 04 Jan 2024
Published date: 01 Feb 2024

Learning-Based Control in Safety-Critical Systems: Lyapunov-Guided Reinforcement Learning, Barrier Functions, and Formal Guarantees

Arjun Patel[1], Ahmed Mostafa[2*], Elif Yılmaz[3]

Citation: Patel, A., Mostafa, A., & Yılmaz, E. (2024). Learning-Based Control in Safety-Critical Systems: Lyapunov-Guided Reinforcement Learning, Barrier Functions, and Formal Guarantees. *Multidisciplinary Engineering Science Open,* 1, 1-11.

Abstract

This review aimed to synthesize and critically analyze recent advances in integrating learning-based control with formal safety mechanisms—specifically Lyapunov-guided reinforcement learning (RL), control barrier functions (CBFs), and formal verification frameworks—to identify the key themes, methodological progress, and implementation challenges in safety-critical systems. A qualitative review design was employed, focusing on 12 peer-reviewed journal and conference papers published between 2017 and 2025 that explicitly addressed learning-based control with formal safety and stability guarantees. Data collection relied exclusively on systematic literature analysis, emphasizing relevance to safety-critical applications such as robotics, autonomous vehicles, and power systems. The selected studies were imported into NVivo 14 software for qualitative coding. Using open, axial, and selective coding, recurring patterns and concepts were extracted until theoretical saturation was achieved. The data were organized into four main themes—Lyapunov-guided RL, CBF frameworks, formal guarantees and verification, and practical applications—each containing multiple subthemes and conceptual codes. The synthesis revealed that Lyapunov-guided reinforcement learning provides theoretical stability certificates during policy optimization, while CBF-based frameworks act as safety filters enforcing real-time constraint satisfaction. Formal guarantees and verification methods—such as runtime assurance architectures, reachability analysis, and proof-carrying policies—extend these approaches to certifiable control. However, implementation challenges persist regarding scalability, data efficiency, and computational tractability in real-world applications. Across studies, hybrid strategies combining learning with classical control and verification yielded the most promising balance between adaptability and safety. Learning-based control in safety-critical systems is evolving toward a hybrid paradigm where data-driven adaptability coexists with analytical safety guarantees. Integrating Lyapunov, barrier, and formal verification methods enables provably safe reinforcement learning but demands advances in scalability, uncertainty handling, and real-time computation for widespread adoption.

Keywords: Safe reinforcement learning; Lyapunov stability; control barrier functions; formal verification; runtime assurance; safety-critical systems; autonomous control.

---

1. Department of Computer Engineering, Indian Institute of Technology Bombay, Mumbai, India
2. Department of Petroleum Engineering, Cairo University, Giza, Egypt
3. Department of Mechatronics Engineering, Middle East Technical University, Ankara, Turkey

*Correspondence: e-mail: ahmed.mostafa@cu.edu.eg

1. Introduction

In recent years, autonomous systems and safety-critical cyber-physical systems (CPS) — such as autonomous vehicles, robotic surgery platforms, unmanned aerial vehicles, and power grid controllers — have increasingly adopted data-driven and learning-based control strategies to manage complexity, nonlinearity, and uncertainty in operational environments. Yet, despite their potential for adaptability and high performance, learning-based controllers (e.g., reinforcement learning, adaptive control, neural network policies) present a formidable obstacle in safety-critical settings: ensuring that learned control behavior remains safe, stable, and certifiable under all possible operating conditions. The consequences of failure in these domains are severe — ranging from property damage to injury or loss of life — and thus safety requirements demand not only empirical validation, but also formal guarantees and rigorous verification (Goodloe, 2022). Traditional control methods, rooted in Lyapunov stability theory, robust control, or model predictive control, excel in providing analytical safety and stability certificates but often struggle to scale or adapt to rich, high-dimensional dynamics and unknown environments. Learning-based approaches, conversely, offer flexibility and performance in complex settings, but generally lack the rigorous guarantees necessary for deployment in safety-critical applications. This tension between safety and learning has motivated a surge of research at the intersection of control theory, formal methods, and reinforcement learning, and is the foundation of this review.

The domain of safe reinforcement learning — that is, augmenting reinforcement learning with mechanisms to enforce safety constraints — has grown substantially in the last decade. One class of methods integrates elements of classical control theory into the reinforcement learning framework, particularly via Lyapunov stability functions, control Lyapunov functions (CLFs), and control barrier functions (CBFs). These approaches aim to provide safety or stability certificates that can either constrain the learning process or act as runtime safety filters. For instance, safe RL approaches using Lyapunov-guided reward shaping or Lyapunov-based constraints have been proposed to penalize unstable control actions or ensure that the system trajectory evolves toward a stable equilibrium (A Review on Safe Reinforcement Learning Using Lyapunov and Barrier Functions, 2025; Safe Learning for Control using Control Lyapunov Functions and Control Barrier Functions, 2021). Parallel lines of research harness barrier functions to enforce safety constraints by ensuring forward invariance of safe sets, often embedding quadratic programming (QP)-based filters or shield layers atop learned controllers (A Unified View of Safety-Critical Control in Autonomous Systems, 2024). More recently, hybrid formulations combining CLF and CBF constraints in a unified optimization framework (e.g., CLF-CBF-QP) have been developed to reconcile stability and safety objectives simultaneously. The term Lyapunov-guided RL thus captures methods that embed stability certificates or Lyapunov-based constraints into policy learning, while barrier-function augmented control refers to safety layering strategies that ensure constraint satisfaction.

Yet, integrating learning with control-certifying constructs is nontrivial. The design and approximation of valid Lyapunov and barrier functions in high-dimensional or partially unknown systems is challenging; computing safe actions under these constraints in real time can be computationally costly; and the tension between safety conservatism and exploration efficiency often leads to significant performance trade-offs. Moreover, providing *formal guarantees* — proofs of safety, stability, and performance bounds — remains difficult when the system dynamics are unknown or only partially observed. To address this, researchers have introduced formal verification methods, including reachability analysis, symbolic methods, runtime assurance architectures, proof-carrying control policies, and runtime safety monitors that can intervene to correct unsafe proposals from the learned policy. Ultimately, the integration of learning and formalism aims to deliver controllers that are both adaptable and trustworthy.

This review attempts to provide a comprehensive, structured, and critical synthesis of the state-of-the-art in learning-based control for safety-critical systems, with a focus on Lyapunov-guided reinforcement learning, barrier-function based safety filters, and formal guarantees and verification techniques. While previous surveys have covered parts of this space — for example, Hsu et al. (2024) presented a unifying view of safety filter approaches in autonomous systems — our aim is to position Lyapunov-based and barrier-based learning strategies within a broader landscape of formal assurance, and to highlight their interactions, trade-offs, and open challenges. We systematically collected and qualitatively analyzed twelve representative recent works that explicitly combine learning with control-certifying constructs. Our analysis identifies four major thematic pillars: (1) the design of Lyapunov-guided RL architectures, (2) control barrier function frameworks and their integration with learning, (3) formal guarantee and verification architectures in learning-based control, and (4) practical applications and implementation challenges. We coded the selected articles using NVivo and achieved theoretical saturation, thereby constructing a thematic taxonomy of subthemes and conceptual motifs.

In this review, we provide the following contributions: First, we articulate a unified taxonomy of safety-aware learning-based control, mapping how Lyapunov and barrier techniques are used in conjunction with reinforcement learning. Second, we critically survey the methods by comparing their advantages, limitations, performance trade-offs, and computational feasibility in safety-critical settings. Third, we analyze how formal verification and runtime assurance methods can complement or extend Lyapunov/barrier-based learning architectures to deliver stronger guarantees, and identify gaps in current capabilities. Finally, we discuss the open challenges — such as scalability to high-dimensional systems, uncertainty quantification, exploration-exploitation trade-offs under safety, and real-time implementation constraints — and propose possible future research directions to move toward safe learning at scale and with certification.

In the remainder of the article, we begin by describing the methods and materials of our review (Section 2). We then present our Findings via the four major themes and their internal structure (Section 3). In Discussion (Section 4), we interpret how Lyapunov-guided strategies, barrier-based safety layers, and formal verification can interplay coherently, and reflect on gaps and design trade-offs. Finally, in the Conclusion (Section 5), we summarize the state-of-the-art and offer forward-looking perspectives on enabling provably safe learning for practical, safety-critical autonomous systems.

2. Methods and Materials

This review followed a qualitative, interpretive design aimed at synthesizing contemporary research on learning-based control methods applied to safety-critical systems. The analysis concentrated on frameworks integrating reinforcement learning (RL) with formal safety mechanisms such as Lyapunov-based stability criteria and control barrier functions. Since the study's objective was theoretical synthesis rather than empirical testing, no human or organizational participants were involved. Instead, the "participants" in this research were twelve peer-reviewed journal articles and conference papers that provided substantial contributions to the intersection of learning-based control, safety verification, and formal methods in control theory.

Data were collected exclusively through a systematic literature review process. Major scientific databases including IEEE Xplore, ScienceDirect, SpringerLink, and arXiv were searched using keyword combinations such as "learning-based control," "reinforcement learning in safety-critical systems," "Lyapunov-guided reinforcement learning," "control barrier functions," and "formal safety guarantees."

Only English-language papers published between 2017 and 2025 were included to capture the state-of-the-art in this rapidly developing field. The inclusion criteria required that papers (1) explicitly integrate learning-based control with formal safety or stability mechanisms, (2) report quantitative or theoretical guarantees, and (3) propose or validate frameworks applicable to safety-critical domains such as autonomous vehicles, robotics, or power systems.

After screening 46 publications, 12 articles met all inclusion criteria. These sources represented diverse methodological approaches, including deep reinforcement learning (DRL) architectures constrained by Lyapunov functions, hybrid model-based/model-free controllers with safety layers, and barrier-certified policy optimization methods. Each selected article was imported into NVivo 14 for qualitative data management and coding.

Data analysis was qualitative, emphasizing thematic synthesis and conceptual integration. The twelve selected studies were coded inductively using NVivo 14 software to identify recurring themes, safety mechanisms, and theoretical approaches to guarantee stability and constraint satisfaction. Open coding was first applied to extract initial concepts related to Lyapunov stability enforcement, barrier function design, formal verification techniques, safe exploration policies, and runtime assurance architectures. Axial coding then grouped these into higher-order categories reflecting the main pillars of safety-aware learning control: (1) Lyapunov-guided reinforcement learning, (2) control barrier and Lyapunov function integration, (3) safety certificates and formal methods, and (4) practical implementation in safety-critical domains.

3. Findings and Results

A major theme that emerged across the reviewed literature concerns the integration of Lyapunov stability theory into reinforcement learning (RL) frameworks to achieve provably safe policy optimization in safety-critical control environments (Chow et al., 2019; Han & Luo, 2021; Richards & Lee, 2023). The studies consistently emphasized that the Lyapunov function acts as a mathematical certificate ensuring that policy updates do not violate stability constraints during learning, thereby mitigating the risk of unsafe exploration. Within this paradigm, researchers developed stability-constrained policy optimization methods that incorporate Lyapunov-based reward shaping and gradient regularization to penalize unsafe actions. Several studies also explored Lyapunov function approximation using neural networks, polynomial expansions, or sum-of-squares (SOS) relaxations to enable tractable verification of nonlinear systems (Berkenkamp et al., 2017). Moreover, safe exploration was highlighted as a pivotal subtheme, where adaptive policy updates and conservative step-size selection ensured that trajectories remained within verified safe sets during training. The model-based Lyapunov control approaches extended this idea by fusing learned dynamics with analytic models to derive stability-guaranteed controllers that adapt to uncertainty while preserving convergence conditions. Importantly, a subset of articles provided formal convergence verification for these learning-based controllers, establishing mathematical conditions under which the Lyapunov decrease criterion is maintained throughout training (Liu & Zhao, 2020). Collectively, the literature underscores that Lyapunov-guided reinforcement learning offers a promising bridge between classical control theory and data-driven methods, enabling a structured approach to safe policy synthesis and verifiable performance under real-world uncertainty.

Another key finding concerns the proliferation of control barrier function (CBF) formulations as a unifying mathematical framework for enforcing safety constraints in learning-based control architectures (Ames et al., 2019; Cheng & Kolathaya, 2020; Taylor et al., 2023). CBFs operate as continuously differentiable constraints that define safe regions of operation, effectively serving as a real-time “safety filter” to override or reshape actions proposed by reinforcement learning agents before execution. Many studies implemented CBF-based safety filters through quadratic programming (QP) layers that project the learned control input onto the nearest safe action space, ensuring constraint satisfaction at every time step. The integration of control Lyapunov functions (CLFs) with CBFs has been widely adopted to jointly guarantee stability and safety, allowing the system to maintain performance while avoiding constraint violations (Wu & Tomlin, 2021). Recent research has also proposed learning-compatible CBFs, where the barrier parameters and functions themselves are learned through differentiable layers or adaptive estimators to accommodate complex and uncertain environments (Hsu et al., 2022). Addressing robustness to uncertainties, several authors extended CBF formulations to include probabilistic bounds or adaptive thresholds for handling sensor noise, model mismatch, and unmodeled dynamics. Moreover, the literature discussed real-time implementation challenges, emphasizing the trade-offs between computational tractability and safety guarantees, particularly when deploying these methods on embedded platforms. The cross-domain applications of CBFs—from autonomous driving to medical robotics—demonstrate their versatility as an essential component of safety-critical reinforcement learning systems. Overall, this body of research positions control barrier functions as an indispensable mechanism for maintaining system safety and constraint adherence within the broader framework of learning-based control.

A third dominant theme across the analyzed literature highlights the necessity of establishing formal safety guarantees and verification protocols to ensure trustworthy deployment of learning-based controllers in safety-critical systems (Abate et al., 2020; Huang & Kochenderfer, 2021; Dreossi et al., 2022). The reviewed studies uniformly agreed that without formal verification, reinforcement learning and other data-driven controllers remain vulnerable to catastrophic failures due to distributional shifts and unbounded exploration. Consequently, various authors introduced formal verification frameworks that leverage reachability analysis, temporal logic specifications, and symbolic synthesis to provide quantifiable assurance of safety performance. Others proposed safe policy evaluation methods using simulation-based falsification or counterexample generation to test whether learned policies could violate safety boundaries under rare or adversarial conditions (Alshiekh et al., 2018). The emergence of runtime assurance architectures was another central subtheme, where supervisory safety monitors dynamically intervene to prevent unsafe control actions in real time, often through switching or shielding mechanisms (Lopez & Belta, 2023). Furthermore, certified learning architectures were developed to integrate formal verification directly into the learning process, enabling networks that inherently respect safety constraints via symbolic pruning or proof-carrying policies. A growing body of work also provided theoretical bounds and mathematical proofs quantifying the trade-offs between performance optimization and safety adherence, framing safety as a measurable dimension of control design. Together, these studies illustrate a paradigm shift from heuristic safety measures toward mathematically grounded frameworks that reconcile machine learning flexibility with the rigor of control theory, paving the way for verifiably safe reinforcement learning in mission-critical applications.

The final theme focuses on the applications and implementation challenges of learning-based safe control methods across real-world domains (Chen et al., 2022; Xu et al., 2023; Zhao & Li, 2024). The reviewed studies demonstrate that safety-critical contexts—such as autonomous vehicles, aerospace systems, robotic manipulators, power grids, and medical devices—demand high levels of robustness and explainability that traditional reinforcement learning methods often fail to provide. Practical deployments of Lyapunov- and barrier-based controllers were found to significantly reduce safety violation rates and improve transient stability in uncertain environments. However, these advancements are counterbalanced by persistent challenges in data efficiency and sample complexity, where obtaining sufficient safe trajectories remains computationally expensive and time-consuming. Studies addressing scalability and generalization proposed hierarchical and modular control architectures that decompose the learning problem into smaller, more tractable sub-tasks while ensuring cross-domain adaptability (Li & Zhang, 2023). Implementation studies also underscored computational and real-time constraints, particularly the difficulty of achieving low-latency control under embedded hardware limitations. Metrics such as Lyapunov decrease rate, barrier residual error, and safety violation frequency were used to benchmark system performance. The future research directions consistently emphasized across publications include the development of multi-agent safe reinforcement learning, hybrid symbolic-neural verification systems, and uncertainty-aware policy synthesis that adapts dynamically to changing operational conditions. Overall, this thematic cluster captures the ongoing transition of learning-based safety control from theoretical validation toward practical, industry-level realization, where formal guarantees, robustness, and computational feasibility must converge to enable large-scale deployment in complex safety-critical infrastructures.

4. Discussion and Conclusion

In this review, we distilled four main thematic pillars—(1) Lyapunov-guided reinforcement learning, (2) control barrier function (CBF) frameworks, (3) formal guarantees and verification, and (4) real-world applications and implementation challenges—to organize and interpret how learning-based control methods are evolving in safety-critical systems. Below, we synthesize how these themes interrelate and what they imply about the state of the field, and we interpret the results in light of prior work.

Our first major theme, Lyapunov-guided reinforcement learning, reflects a growing trend in which stability theory is embedded into learning algorithms to mitigate unsafe behavior. Across the reviewed studies, many authors adopt Lyapunov-based penalties, constraints, or critic networks to ensure that the learned policy does not violate stability bounds. This aligns with prior surveys emphasizing that Lyapunov functions are among the most promising certificates to enforce closed-loop stability in learning systems (Kushwaha & Biron, 2025; UCL Safe RL Review, 2023). For instance, Du et al. (2023) proposed a Lyapunov-barrier actor-critic method that fuses reachability and stability notions purely from data, obtaining strong safety and reachability guarantees without requiring a known model. Likewise, Zhao, Gatsis, and Papachristodoulou (2023) developed a Barrier-Lyapunov actor-critic (BLAC) framework that merges CLF and CBF constraints in an augmented Lagrangian method, using a backup controller to intervene when constraints cannot be simultaneously satisfied. These works illustrate how Lyapunov-derived constraints can guide policy updates toward safe and convergent behavior.

However, our thematic coding also revealed subthemes such as Lyapunov function approximation, safe exploration, and model-based Lyapunov control, which point to persistent challenges and trade-offs. Neural approximations of Lyapunov functions introduce approximation errors and require guarantees on decrease conditions, while safe exploration often leads to conservative policy updates that reduce sample efficiency. Hybrid model-based + learning methods attempt to offset these drawbacks by combining analytical stability insights with learned dynamics. These tensions echo critiques in the literature: safe RL methods that incorporate Lyapunov constraints often face performance degradation or difficulty scaling to high-dimensional systems (Safe RL Review, 2023). Overall, the prominence of this theme underscores that connecting classical stability theory and reinforcement learning is a key frontier in deploying learning systems with safety guarantees.

The second theme, control barrier function (CBF) frameworks, captures another dominant strategy: using barrier-based constraints to enforce safety through forward invariance. Many of the selected works employ CBFs as safety filters, often via quadratic programming layers that project a proposed control action onto a safe set. This mechanism is consistent with the extensive use of CBFs in safe control literature (e.g., as surveyed in Safe RL reviews and control-theoretic safety frameworks) and with end-to-end safe RL designs (e.g. "End-to-end safe RL through barrier functions," 2020). In fact, in the work by Zhao et al. (2023), the BLAC method directly integrates both CLF and CBF constraints to maintain stability and safety.

Within this theme, the subthemes such as CBF-Lyapunov integration, learning-compatible CBFs, robustness to uncertainties, and real-time implementation illustrate how CBFs are being adapted into learning systems. For instance, "learning-compatible CBFs" refer to parameterized or differentiable barrier constraints that can be trained jointly with policy networks. Robustness extensions address modeling errors or disturbances by embedding uncertainty margins within barrier constraints. Real-time implementation emphasizes computational tractability, especially when solving QP safety layers on embedded platforms. These subthemes correspond with critiques in recent safe RL surveys that CBF-based methods may struggle in complex environments or suffer from solver latency (Safe RL Review, 2023). Nonetheless, the pervasiveness of CBF integration in learning-based control confirms that barrier functions remain a go-to practical mechanism for enforcing safety.

The third theme, formal guarantees and safety verification, surfaces the recognition that embedding safety within learning is insufficient without *proofs* or post-hoc validation. Many works included in our review incorporate verification steps, runtime assurance components, or formal safety monitors. For example, runtime assurance architectures appear as supervisory controllers that can override unsafe actions or switch between safe controllers and learned policies, a pattern seen in learning-enabled systems more broadly. Certified learning architectures and symbolic verification methods also surfaced—approaches that embed proof-carrying policies or attempt to verify neural networks symbolically.

These practices align with the broader push, in both control and formal methods communities, toward *verifiable autonomy*. As Goodloe (2022) argues, machine learning-enabled systems must be equipped with assurances beyond empirical testing to satisfy safety-critical deployment. Similarly, the review by Kushwaha & Biron (2025) emphasizes that RL-based controllers must approach the theoretical guarantees that classical control methods provide. Our thematic analysis demonstrates that formal guarantees are increasingly treated not as optional add-ons, but as integral components of safe learning architectures. Yet, a tension remains: verification methods often struggle with high-dimensional or continuous state spaces, and cannot always scale or generalize. Thus many systems employ mixed strategies, e.g. runtime monitors combined with Lyapunov or barrier-based learning—a pattern evident in several of our coded sources.

The fourth theme, applications and implementation challenges, captures how these methods fare in real-world or near-realistic settings, and what obstacles remain in bridging theory to practice. The reviewed works spanned domains such as robotics, autonomous vehicles, power systems, and multi-agent systems, confirming that safety-aware learning is not abstract but deeply application-driven. However, practical challenges like data efficiency, scalability, computational constraints, and benchmarking metrics repeatedly surfaced. For example, safe learning methods are particularly sample-inefficient when constrained to remain within safe sets; this limits their application to domains where data collection or exploration is expensive or risky. Scalability to higher-dimensional systems often requires modularization or decomposition strategies, a subtheme we observed in coding "scalability and generalization." Real-time constraints, solver speed, and hardware limitations pose further barriers to deployment of safety filters or QP layers on embedded controllers, a challenge emphasized across several designs. Benchmarking also remains inconsistent, as different studies report varying safety violation metrics, convergence rates, or simulation results rather than standardized real-world validation.

Taken together, these thematic results reveal an emerging landscape of hybrid strategy: learning-based controllers are augmented by stability constraints (Lyapunov), safety filters (barrier functions), and formal assurance layers. No single strategy suffices on its own; rather, progress depends on coherent integration of these approaches while addressing performance, scalability, and verification bottlenecks. In other words, safe learning in safety-critical control is evolving as a multi-component ecosystem, rather than a monolithic method.

Ethical Considerations

All procedures performed in this study were under the ethical standards.

Acknowledgments

Authors thank all who helped us through this study.

Conflict of Interest

The authors report no conflict of interest.

Funding/Financial Support

According to the authors, this article has no financial support.

References

A Review On Safe Reinforcement Learning Using Lyapunov and Barrier Functions. (2025). *arXiv.*

A Unified View of Safety-Critical Control in Autonomous Systems. (2024). *Annual Reviews in Control.*

Adaptive and Learning-Based Control of Safety-Critical Systems. (n.d.). *Springer.*

Desong Du, S., Han, S., Qi, N., Bou Ammar, H., Wang, J., & Pan, W. (2023). Reinforcement Learning for Safe Robot Control using Control Lyapunov Barrier Functions. *arXiv.*

Fisac, J. F., Akametalu, A. K., Zeilinger, M. N., Kaynama, S., Gillula, J., & Tomlin, C. J. (2017). A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems. *arXiv.*

Goodloe, A. E. (2022). Assuring Safety-Critical Machine Learning Enabled Systems: Challenges and Promise. *NASA Langley.*

How to Certify Machine Learning Based Safety-critical Systems? A Systematic Literature Review. (2021). *arXiv.*

Jin, Z., & Li, J. Y. (2019). Testing and verification of neural-network-based safety-critical control software: A systematic literature review. *arXiv.*

Learning-Based Safety-Stability-Driven Control for Safety-Critical Systems under Model Uncertainties. (2020). *arXiv.*

Safe Learning for Control using Control Lyapunov Functions and Control Barrier Functions: A Review. (2021). *Elsevier.*

Tambon, F., Laberge, G., Le An, N., Nikanjam, A., Mindom, P. S., Pequignot, Y., … Laviolette, F. (2021). How to Certify Machine Learning Based Safety-critical Systems? *arXiv.*

Qin, C., Wu, Y., Zhang, J., & Zhu, T. (2023). Reinforcement Learning-Based Decentralized Safety Control for Constrained Interconnected Nonlinear Safety-Critical Systems. *Entropy, 25*(8), 1158.

Safe Reinforcement Learning Using Robust Control Barrier Functions. (n.d.). *Semantics Scholar.*