Lyapunov– Barrier Duality: A Structural Perspective

Seigo Sasaki

Department of Electrical and Electronic Engineering
National Defense Academy of Japan, Yokosuka, Japan
seigo@nda.ac.jp

Abstract

This article proposes a unified region-based paradigm that reconnects historically independent developments in nonlinear control, optimal control, and safety-critical systems. We began with Lyapunov's original insight that stability is fundamentally a property of regions and their deformation under flow, we trace a conceptual lineage through Zubov's characterization (1950s), the variable-gradient method (1962), extended quadratic Lyapunov functions (1997-1998), contraction theory (2014-), and recent advances in safe reinforcement learning and Hamilton-Jacobi-Bellman methods (2020-). We argue that contemporary challenges in learning-based and safety-critical control stem from treating stability as a pointwise property rather than a geometric one. A unified region-based view clarifies the complementary roles of Lyapunov functions (inner contraction), barrier functions (outer invariance), and value functions (performance landscapes within feasible domains). It also highlights constructive methods that jointly design gradients and level sets. By reconstructing these connections, we provide a coherent framework that situates modern tools within Lyapunov's original geometric philosophy—not as a new theory, but as a synthesis for interpreting and integrating existing methods.

Keywords: Lyapunov function, barrier function, contraction theory, region-based stability, nonlinear control, Hamilton-Jacobi-Bellman (HJB), reinforcement learning

1 Introduction — Stability as a Geometry of Regions

Since Lyapunov's foundational work, stability theory has been shaped by two parallel traditions: one that emphasizes scalar functions whose sublevel sets contract toward an attractor, and the other that emphasizes gradient and vector fields that guide the flow of trajectories. Decades of advances in nonlinear control, optimal control, and safety-critical systems have tended to adopt one of these perspectives and ignore the other. As a result, the fundamental geometric insight present in Lyapunov's original interpretation, namely that stability is a property of a domain and its deformation under flow, has been partially obscured.

This paper revisits the geometric foundations of Lyapunov stability theory and traces its historical progression from classical domain-based interpretations, Lyapunov's monograph [1] and Zubov's construction [2], to variable gradient methods (1962) [3], extended quadratic Lyapunov functions (1997 – 1998) [4, 5], contraction theory (since 2014) [6, 7], and recent developments in reinforcement learning and safety-critical control (since 2020) [8, 9, 10, 11, 12, 13, 14]. We propose a unified domain-based framework that integrates attraction, safety, and optimality through the collaborative design of level sets and gradients.

Zubov's partial differential equation-based methods, developed in the mid-20th century, provide a rigorous characterization of domains of attraction [2], but do so through an ontological and descriptive lens. Unlike constructive approaches such as variable gradient methods [3] and extended quadratic frameworks [4, 5], Zubov's theory does not offer a design-oriented mechanism for forming level sets or gradients. Therefore, this paper treats Zubov's theory as part of its own lineage and does not extend it. Between the rigor of control theory and the expansiveness of robotics, the geometry of stability has been fragmented, forgotten, and rediscovered. This paper seeks to trace the contours of this tension as a foundation for structural synthesis, rather than to resolve it.

The central argument of this paper is that modern control theory requires a unified domain-based paradigm that integrates contraction [6, 7], attraction [2, 3], safety [15, 16, 17], and optimality [18, 19] into a coherent geometric framework. In particular, we argue that Lyapunov design must be understood as a dual problem: forming an inner contraction domain and an outer invariant boundary, and ensuring that they are compatible with optimality and learning. This perspective not only clarifies connections between seemingly contradictory literature but also highlights the constructive potential of methods for designing both gradient and level set structures, such as the extended quadratic Lyapunov framework [4, 5].

By reconstructing the connections between historical developments and current trends, we aim to provide a more unified and physically meaningful conceptual foundation for nonlinear stability theory. This theory will guide the next generation of stability analysis, controller design, and learning-based methods.

Rather than reproducing the technical details of each method, this paper focuses on the conceptual evolution of Lyapunov-based reasoning across disciplines. Our aim is to reveal how geometric, optimization-based, and data-driven perspectives have reshaped the understanding of stability. The arguments are grounded in cited literature and structured to support a unified geometric paradigm. This is not a conventional control theory paper, but a design-oriented synthesis of ideas that have emerged across theory, robotics, and learning.

Accordingly, this paper is structured as a conceptual synthesis rather than a conventional technical exposition. Our goal is not to present new theorems or algorithms, but to reveal the underlying geometric unity that connects diverse methodological developments. We focus on conceptual evolution and structural relationships, with technical details delegated to cited literature. This design-oriented perspective aims to provide a roadmap for integrating stability analysis, safety constraints, and learning-based methods within a coherent geometric framework.

2 Lyapunov's Original Perspective: Stability as a Property of Regions

The contemporary view of Lyapunov theory often begins with the requirement that a candidate function $V(x)$ be positive definite and strictly decreasing along trajectories. Yet this familiar formulation tends to obscure the central insight contained in Lyapunov's original monograph—The General Problem of the Stability of Motion (1892) [1, 20]. Lyapunov consistently emphasized that stability is not a property of a point, but rather of a region of the state space containing the equilibrium. In his formulation, the level sets

$$\Omega_c = \{x : V(x) \le c\}$$

serve as the structural basis for Lyapunov's reasoning: they define closed domains within which the motion remains confined, and they contract under the system dynamics.

From this viewpoint, the Lyapunov function serves not merely as an algebraic certificate, but as a geometric descriptor of an invariant domain. The decrease condition $\dot{V}(x(t)) < 0$ is meaningful precisely because it implies the nesting of these domains: trajectories flow from larger to smaller level sets. Thus, the function $V(x)$ is a coordinate representation of a more fundamental geometric object—the family of invariant regions surrounding the equilibrium [2].

This region-based interpretation is explicit throughout Lyapunov' s text: the analysis proceeds by constructing neighborhoods whose boundaries are level sets of $V(x)$, and by ensuring that trajectories do not exit these neighborhoods. The commonly taught "pointwise" reading of Lyapunov theory—focusing only on positivity and dissipation—therefore overlooks the original geometric foundation of the method. It is this neglected foundation that later developments, from variable-gradient constructions [3] to modern contraction [6, 7] and barrier methods [15, 17], ultimately rely upon.

3 The 1962 Variable-Gradient Method: Constructing ∇V Without Constructing V

The work of Schultz and Gibson (1962) [3] represents one of the earliest systematic attempts to operationalize Lyapunov' s region-based viewpoint. Rather than searching directly for a function $V(x)$ whose level sets define invariant domains, they proposed to construct its gradient field $g(x) = \nabla V(x)$ as the primary object. This shift of emphasis—from $V(x)$ to its differential structure—was motivated by the observation that the geometric requirement $\dot{V}(x(t)) = g(x)^\top f(x) < 0$ depends only on the gradient, and not on the integral form of $V(x)$ itself.

The variable-gradient method makes this idea explicit. One selects a parametric family of vector fields $g(x)$ and imposes two conditions:

1. . Dissipation:
$$\dot{V}(x(t)) = g(x)^\top f(x) < 0 \quad (x \neq 0)$$
ensuring contraction of level sets.

2. . Integrability: The generalized curl of $g(x)$ must vanish,
$$\frac{\partial g_i(x)}{\partial x_j} = \frac{\partial g_j(x)}{\partial x_i}, \tag{1}$$
so that a scalar potential $V(x)$ exists.

In principle, satisfying these conditions yields both the Lyapunov function and its domain-level geometry. But the practical limitation is apparent: ensuring integrability across nonlinear systems amounts to solving a coupled set of partial differential constraints, often as difficult as constructing $V(x)$ directly. Moreover, because the method centers on the gradient rather than the level sets, the resulting $V(x)$ may satisfy $\dot{V}(x(t)) < 0$ without possessing meaningful geometric properties—such as convexity, radial unboundedness, or domain shape—that govern the invariant regions around the equilibrium [21].

Regarding the construction of domains, Zubov's PDE-based method was established in the 1950s and 1960s [2], but this is a theory that focuses on the descriptive characterization of the domain of attraction rather than designability, and belongs to a different lineage from variable gradient.

Thus, the 1962 method stands as a conceptual bridge. It recognizes, as Lyapunov did, that stability is ultimately about the shape and contraction of certain regions in the state space. Yet its reliance on integrable gradient fields left the construction problem unresolved. The missing piece —a way to shape both the gradient and the level sets simultaneously—would only surface decades later [4, 5, 22, 6].

4 1997: Simultaneous Construction of Gradients and Level Sets (EQLF)

The extended quadratic Lyapunov function (EQLF) framework introduced in 1997 represented a conceptual departure from earlier approaches to Lyapunov construction [4]. Unlike the variable-gradient method of 1962—which sought to synthesize $\nabla V(x)$ while leaving the shape of the level sets to the integrability constraints [3]—the 1997 formulation treated both the gradient and the level sets as primary design objects. It provided a parametrized structure for $V(x)$ whose geometry could be directly shaped, while ensuring that dissipation and integrability were satisfied by construction.

The essential insight of the method is that the Lyapunov function can be expressed in the form

$$V(x) = x^\top P(x)x,$$

where $P(x)$ is defined through a set of carefully coordinated nonlinear relations. These relations link three structural components:

1. the derivatives of $P(x)$, determining the gradient $\nabla V(x)$,
2. the algebraic form of $P(x)$, determining the geometry of the level sets, and
3. the system dynamics, ensuring $\dot{V}(x(t)) < 0$.

The key is to ensure consistency between $V(x)$ and $\dot{V}(x(t))$ through the relation

$$\dot{V}(x(t)) = \dot{x}^\top X(P(x), x)^\top x + x^\top X(P(x), x)\dot{x},$$

where $X(P(x), x)$ connects the algebraic form of $V(x)$ with its time derivative. This leads to the structural constraint

$$X(P(x), x) = P(x) + \frac{1}{2}\begin{bmatrix} \frac{\partial P}{\partial x_1}(x)x & \cdots & \frac{\partial P}{\partial x_n}(x)x \end{bmatrix},$$

which encodes the radial deformation of the level sets via the structure of $P(x)$.

This structural constraint stands in sharp contrast to the 1962 variable-gradient method [3], where the gradient field $g(x) = \nabla V(x)$ is chosen first, and integrability is verified post hoc by checking the symmetry of its Jacobian, as expressed in equation (1). In that approach, the level sets of $V(x)$ emerge only if this condition is satisfied, making the geometry of the resulting function an incidental outcome rather than a design target. In contrast, the EQLF framework embeds this consistency structurally, allowing the designer to shape both the gradient and the level sets simultaneously.

Unlike the 1962 method, which separates gradient synthesis from integrability verification, the 1997 framework enforces structural consistency by construction. This transforms Lyapunov construction from a two-step verification process into a unified design problem. The level sets defined by $V(x)$ are not incidental byproducts, but deliberately shaped regions whose boundaries encode contraction surfaces in the state space. This restores Lyapunov' s original emphasis on invariant regions [1, 2]: the function is not merely decreasing along trajectories but is designed so that its sublevel sets form closed domains that contract under the dynamics.

The method can thus be seen as the first broadly applicable framework to place Lyapunov design —rather than Lyapunov analysis—at its core. In 1998, the method also solved robust output control problems under nonlinear dynamics [5, 23]. It bridges the gap between classical stability theory and geometric intuition by enabling a coherent co-design of the function, its gradient field, and its level-set geometry. In doing so, it anticipates later developments—such as contraction metrics [22, 6, 7] and control barrier functions [17, 15]—which similarly emphasize the geometric structure of the state space, though often without explicit recognition of this earlier unifying perspective.

5 2000– 2010: The Point-Stability Bias and the Loss of Geometric Insight

The first decade of the 2000s witnessed extraordinary advances in convex optimization, LMIs, and linear-system theory [19], all of which were linked in the 1990s [24, 25]. These tools greatly strengthened the analytical tractability and computational feasibility of nonlinear control problems. However, the focus on pointwise algebraic conditions sometimes obscured the geometric meaning of invariant regions that motivated Lyapunov's original framework. Stability became increasingly identified with pointwise conditions on the derivative of a Lyapunov function, evaluated through algebraic inequalities. The geometric meaning of Lyapunov' s original theory—its emphasis on invariant regions and the contraction of level sets—faded from view [1, 2].

Within this LMI-centric paradigm, nonlinear stability analysis was frequently reduced to a search for a scalar function satisfying

$$V(x) > 0, \qquad \dot{V}(x(t)) < 0,$$

with additional convex constraints imposed to enable computation. Because these inequalities are local and algebraic, the resulting $V(x)$ often certifies stability without conveying meaningful information about the domain of attraction, boundary geometry, or invariant regions. In many cases, the very object that motivated Lyapunov' s original framework—the region in which the motion is confined—was treated as an incidental byproduct, if it appeared at all.

This point-stability bias was reinforced by the rise of polynomial Lyapunov functions and Sum-of-squares (SOS) programming [26, 27]. SOS methods enabled the automated verification of positivity and dissipation conditions, but they also encouraged an interpretation of Lyapunov theory in which the function is optimized pointwise, and its level-set geometry is neither shaped nor analyzed. In parallel, approaches grounded in Riccati-type inequalities or dissipativity theory preserved the algebraic aspects of Lyapunov analysis while largely overlooking its geometric core.

While the period from 2000 to 2010 produced powerful computational tools for stability verification, the emphasis shifted from constructing geometrically meaningful domains to solving algebraic feasibility problems. The conceptual threads connecting modern practice to Lyapunov' s original region-based philosophy—and to earlier constructive methods such as variable-gradient [3] and extended-quadratic formulations [4, 5]—became increasingly attenuated. This gap would eventually motivate renewed interest in trajectory-based and geometric viewpoints that emerged in the following decade.

The publication of Boyd and Vandenberghe' s Convex Optimization (2004) [19] marked a conceptual shift: stability and control were no longer confined to geometric intuition, but rephrased as tractable inequality-constrained optimization problems. Younger researchers who absorbed this optimization-centric worldview in the early 2000s would, a decade later, reshape Lyapunov theory into a language of constraints, solvers, and certificates. While the mathematical foundations of optimization have long been established, the past decade has witnessed a structural reconfiguration of control theory through the lens of computation—making classical ideas newly actionable. To earlier generations of control theorists, the computational turn may appear as a reinterpretation rather than a reinvention. Yet for the researchers, it has opened new pathways for implementation, safety, and scalability. Having witnessed both the geometric intuition of Lyapunov theory and the rise of optimization-based control, this survey attempts to bridge perspectives across generations.

6 2014–: Contraction and the Revival of Lyapunov's Motion-Based View

The emergence of contraction analysis and its geometric generalizations in the 2010s marked a turning point in stability theory. In contrast to the point-based interpretations that dominated the preceding decades, contraction theory studies the behavior of pairs of trajectories. The central question is no longer whether a particular equilibrium is attractive, but whether all trajectories converge toward one another, or toward a lower-dimensional flow such as a periodic orbit [22]. This shift returns stability analysis to Lyapunov's original concern: the stability of motion, not merely of points [1].

Forni and Sepulchre (2014) [6] provided a rigorous differential-geometric formulation of this idea. It showed that contraction can be characterized through a Riemannian metric $M(x)$, satisfying a differential Lyapunov inequality

$$\dot{V}(x(t), \delta x(t)) = \delta x^\top \left(\dot{M}(x(t)) + M(x)A + A^\top M(x) \right) \delta x < 0,$$

where A is the Jacobian of the dynamics and δx represents virtual displacements between neighboring trajectories. In this framework, stability is expressed not by a scalar function but by the geometry of infinitesimal distances shrinking along the flow. The metric $M(x)$ induces tubes of contracting level sets, providing a geometric interpretation strikingly consistent with the region-based viewpoint of Lyapunov's monograph [1, 2].

Although contraction theory introduces new mathematical language—Riemannian metrics rather than scalar potentials—it may be better understood as a modern reconstruction of Lyapunov's motion-based stability framework, now expressed through differential-geometric tools. Instead of a scalar $V(x)$ whose level sets describe invariant regions, contraction employs a smoothly varying metric whose geodesic balls shrink under the dynamics. Both approaches share the same geometric essence: stability arises when the flow maps regions into strictly smaller regions. The difference lies in the mathematical language—scalar potentials versus Riemannian metrics—rather than in the underlying philosophy.

Contraction theory thus repaired a conceptual gap that had persisted for decades. By focusing on the deformation of regions in the state space, it restored the link to Lyapunov's foundational insight while offering new tools for systems with no equilibrium, time-varying dynamics, or incremental performance requirements [7]. Yet, despite this revival, contraction has often remained conceptually separate earlier constructive methods such as extended-quadratic Lyapunov functions [4, 5], which achieve a similar geometric coherence within an algebraic framework. Reconnecting these threads remains essential for a unified view of modern nonlinear stability.

7 2017–: Barrier Functions as the One-Sided Face of Lyapunov Theory

The recent rise of barrier functions—and in particular of control barrier functions (CBFs) in robotics and safety-critical systems—represents a pragmatic and influential strand of modern control research [15, 16, 17]. Optimization-based barrier methods appear to offer a fundamentally different approach [19] from Lyapunov theory: while Lyapunov seeks functions whose sublevel sets contract toward an attractor, barrier functions impose invariance of a prescribed safe set by ensuring trajectories do not cross a forbidden boundary. Yet a closer inspection reveals that barrier functions are not a replacement for Lyapunov theory but rather its one-sided dual.

A (control) barrier function $B(x)$ is typically used to enforce forward invariance of a set $\mathcal{S} = \{x : B(x) \geq 0\}$ by imposing a differential inequality of the form

$$\dot{B}(x(t)) \geq -\alpha(B(x))$$

for an extended class-$\mathcal{K}$ function α. In control applications, such conditions are often enforced pointwise via real-time quadratic programs (the CLF – CBF – QP architecture), which select control inputs that simultaneously seek performance (through a control Lyapunov function) and safety (through a barrier constraint) [15, 17]. The practical attractiveness of this approach is clear: it permits a modular, online enforcement of safety constraints and fits naturally into existing control stacks.

Geometrically, barrier functions govern the exterior of safe regions (keeping trajectories outside the forbidden set), while Lyapunov functions govern the interior (pulling trajectories inward toward the attractor). Considered this way, barrier functions solve a dual design problem: instead of asking "Which interior region will shrink?", we ask "Which exterior boundary must never be crossed?" Both problems concern the geometry of level sets, differing only in sign and whether we prescribe a center of attraction (Lyapunov function) or a forbidden exterior (barrier function).

This duality helps explain the recent sociological shift in control design practice. Barrier methods offer a tractable alternative, precisely because they allow the designer to specify the set. Constructing Lyapunov functions that yield large, useful regions of attraction—and thus the formation of those regions—remains challenging [21]. Barrier methods circumvent this difficulty by allowing the user to declare safe sets and enforcing invariance. This engineering convenience is a major reason for CBF's popularity.

This practical convenience comes with certain trade-offs that are important to recognize. First, systems using only barrier constraints may not converge while maintaining safety. Second, combining barrier constraints with performance goals (such as CLF) can be infeasible and complicate guarantees. Furthermore, unlike Lyapunov design, which forms regions of attraction through analysis, barrier design shifts the burden of specification to the designer. Finally, barriers provide no information about the internal dynamics or the size/shape of the basins of attraction.

From a philosophical perspective, the trend toward barrier methods is understandable. The difficulty of designing Lyapunov level sets that are both computationally tractable and domain-proof makes practical alternatives attractive. However, from the conceptual perspective developed in the previous section, treating barrier methods as a replacement for, rather than complement to, Lyapunov-based attraction analysis overlooks their fundamentally dual nature. That is, it conflates structurally dual tools with functionally interchangeable ones

Thus, in a design philosophy that respects Lyapunov domain-based insights, barrier methods should be treated as a complementary means of enforcing constraints, powerful and essential in many applications, and not as a replacement for the fundamentally difficult task of forming regions of attraction. The future of robust, safety-aware control lies in principled ways of constructing both interior contraction regions (Lyapunov or contraction-based) and exterior invariant boundaries (barrier-based), and controlling their interaction with strict guarantees rather than ad-hoc workarounds.

8 2020 – : The Region-Based Paradigm in RL, Safety, and HJB — A Necessary Reinterpretation

Over the past decade, concepts in reinforcement learning (RL), safety-critical robotics, and optimal control have converged on a common conceptual bottleneck: the need for domain-based guarantees.

As systems become more autonomous and operate under more complex uncertainties, concepts of pointwise or local stability are no longer sufficient.

What is needed, as Lyapunov demonstrated over a century ago, are guarantees about the domain, its evolution in flow, and the interaction between the domain and constraints. This shift has prompted the re-emergence of domain-based reasoning, sometimes unwittingly, even in communities far removed from classical control theory [9].

In reinforcement learning, the value function $V^{\pi}(x)$ under a policy π can be viewed as a large-scale potential that shapes an agent's behavior [13]. Recent research on stability-aware reinforcement learning (e.g., Lyapunov reinforcement learning, safe policy iteration) explicitly enforces conditions such as

$$V^{\pi}(f(x,u)) - V^{\pi}(x) \leq 0,$$

where $f(x,u)$ denotes the next state under the system dynamics in a discrete-time setting. This formulation can be interpreted as a hidden version of the classical Lyapunov decay condition [10, 11, 12], where a candidate function is expected to decrease along system trajectories. However, unlike classical Lyapunov design, which constructs control laws from such candidates, reinforcement learning attempts to learn both the policy and the potential function simultaneously.

This amplifies the challenge of maintaining stability under exploration. Without domain-based constraints, exploration can easily lead the system outside the domain where the learned dynamics and value approximations are valid. Concepts such as the learning domain, safe set, and reachable region have reemerged in the 2020s [28, 29]. These are reformulations of Lyapunov' s original ideas —now expressed through the geometry of regions and the structure of domains.

Barrier functions have attracted attention because they provide a way to explicitly enforce safety [15, 16]. However, barriers only cover the "outer bounds." On the other hand, practical safety control requires consistency between the inner convergence region (covered by the Lyapunov contraction) and the outer safety boundary (covered by the barrier) [6, 7].

The problem here is the inherent existence of "domain mismatch." Specifically, the CLF domain often lies outside the barrier-defined safe set, and the two may not intersect—or if they do, the overlap is narrow. Furthermore, online QP is often infeasible [15, 17]. In other words, the challenge facing CBF research in the 2020s is a recurrence of the classic Lyapunov problem known as domain mismatch. The HJB, the fundamental equation for optimal control, is essentially a partial differential equation reflecting a global domain structure. However, numerical solutions are local and only weakly guaranteed outside the discretization [18]. As a result, how to guarantee the validity of approximate solutions becomes a practical challenge, which is directly related to the problem of domain identification.

As of 2020, many studies have developed numerical HJB solutions and learning-based approximate solutions into a three-layer structure [10, 11, 8]. The first layer estimates the validity region, the second layer ensures invariance/degeneracy within that region, and the third layer uses barriers to ensure external safety. This should be understood as a complementary relationship between the Lyapunov model (interior), the barrier model (exterior), and HJB (topography of optimality), revealing a complex problem structure that cannot be solved by a single theory.

In this way, the challenges faced by the fields of reinforcement learning, safety, and HJB individually actually converge into a single structure. Specifically, the Lyapunov model designs and forms an internal "region of attraction," the barrier sets an external "safety boundary" to avoid discontinuities, and HJB/reinforcement learning explores the internal "topography of optimality." These three are not mutually exclusive; rather, they are naturally linked within a unified, domain-based perspective. The concept of "internal domain construction," proposed by EQLF in 1997–1998 [4, 5], is an important component of this unified structure.

Many recent studies have simply rediscovered from different angles the principle proposed by Lyapunov in the first edition: "The stability of motion is the geometry of the domain" [1, 20, 2]. In that sense, the trends of the 2020s can be said to be a return to the era of domains.

9 Toward a Unified Region-Based Theory of Stability

Through more than a century of research, from Lyapunov's geometric insights to modern developments in contraction, safety, and learning, one consistent theme emerges: stability is fundamentally a property of domains.

Trajectories evolve within, across, and between domains, where geometry encodes attraction, safety, performance, and invariance. Contemporary challenges in nonlinear control, safety-critical robotics, and learning-based systems reflect the fragmentation that has resulted from this domain-based structure being forgotten or only partially adopted.

What is needed today is not a new theory, but a reunification of existing ideas under a consistent geometric and constructive framework. Such a framework allows attraction to arise from the construction of a contracting interior domain (Lyapunov, contraction metric) and safety to arise from maintaining an external invariant boundary (barrier function, feasibility theory). Optimality is encoded by the value landscape within a feasible domain (HJB, RL). Therefore, achieving constructivity requires a combination of gradients and level sets (variable gradients in 1962, [3], EQLF in 1997, [4]) rather than using either alone.

These are not competing paradigms, but complementary aspects of a single geometric construct. Historically, the research community has tended to favor either "gradients" or "level sets." Schultz and Gibson's (1962, [3]) variable gradients enhanced the former, while SOS [27] and LMI methods[24, 25] paved the way for the algebraic treatment of the latter. However, methods that simultaneously construct both "gradients and contours" are extremely rare, and EQLF (1997, [4]) is one of the few successful examples. Meanwhile, contraction (2014, [6]) reformulated Lyapunov's "stability of motion" in the language of modern differential geometry. The "domain problem" facing reinforcement learning and safety research in the 2020s naturally complements this trend.

Taking these findings together, one conclusion emerges. Lyapunov's ideas are unified, not incomplete. A unified theory requires, at a minimum, a two-domain geometry that simultaneously considers both the inner domain, dealing with convergence, optimization, and trajectory formation, and the outer domain, dealing with safety, constraints, and non-aggression domains. The consistency of this two-layer structure is essential for stability guarantees. That is, as in the 1997 paper, a framework for consistently designing level sets (domains) and their gradients (directions) is necessary. And, as reduction shows, even in complex nonlinear dynamics, capturing local deformations can reveal global structure.

Furthermore, to provide computationally meaningful guarantees, integrating learning-based and approximation-based methods requires estimation of validity domains, theoretical guarantees both within and outside the validity domains, and addressing errors in data-driven estimation. Lyapunov analysis and controller synthesis should not be treated separately, but as two sides of a domain-based design. This is the very spirit of the EQLF framework [4, 5]. This integration is neither novel nor flashy. Rather, it's a natural progression if we understand the genealogy of Lyapunov's original papers in 1962 [3], 1997 [4], and 2014 [6]. It's a structure implicit in recent developments across disciplines, now ready for explicit articulation.

The purpose of this paper is to present a domain-based stability paradigm that clearly bridges past insights and integrates contemporary research themes (reinforcement learning, safety, and optimal control). In that sense, this framework is not a critique but a starting point for strengthening, interconnecting, and extending existing theory. Lyapunov's journey—from 1962 [3] and 1997 [4] to the learning and safety paradigms of the 2020s—has reached a mature stage of integration, awaiting articulation. Concretely, a unified region-based framework should address:

1. Constructive geometry: Techniques that simultaneously shape both level sets and gradient fields, extending ideas from the 1997 extended quadratic framework to modern settings.

2. Inner-outer consistency: Methods for ensuring that the domain of attraction (Lyapunov) and the safe set (barrier) are compatible, avoiding the infeasibility issues common in CLF-CBF-QP formulations.

3. Contraction and invariance: Integration of differential contraction metrics with set-based invariance constraints, unifying trajectory-level and region-level reasoning.

4. Learning within domains: Principled approaches for estimating the validity region of learned dynamics and value functions, with guarantees both inside and outside this region.

10 Conclusion — Toward a Unified Geometry of Stability

Across more than a century of developments, stability theory has oscillated between local, pointwise criteria and global, geometric descriptions. The resulting landscape is rich but fragmented: contraction theory emphasizes incremental behavior [22, 6, 7], barrier methods address safety [15, 17], Lyapunov functions describe attraction [1, 2], and HJB-based approaches encode optimality [18, 10]. Yet each of these perspectives captures only one facet of a broader geometric reality.

This article has argued that the time is ripe to unify these threads under a region-based stability paradigm. In this view, stability is fundamentally about how regions are deformed under the system dynamics:

- inner regions contract (Lyapunov, contraction) [1, 6],

- outer regions remain invariant (barriers, viability) [15, 16],

- internal landscapes encode performance (HJB, RL) [13, 10],

- and constructive methods shape the geometry itself (1962 gradient design, 1997 extended-quadratic formulations) [3, 4, 5].

The emergence of learning-based control and safety-critical autonomy has amplified the need for such a unification [9, 8]. These domains routinely face the issue that classical guarantees fail outside limited operating regions, and thus require explicit reasoning about the geometry of feasible and safe sets. Reinforcing Lyapunov's original insight, modern challenges reaffirm that regions—not points—are the natural unit of stability analysis [1, 20].

Reconstructing stability theory in this geometric manner does not invalidate existing methods; it clarifies their role within a larger conceptual structure. It also highlights directions that deserve renewed attention, such as the constructive co-design of level sets and gradients [4], and the integration of inner contraction with outer invariance [6, 17]. These ideas, already implicit in foundational works from Lyapunov to 1997's extended-quadratic approach, offer a generative basis for a unified and scalable theory of stability.

Ultimately, a perspective that integrates diverse, intertwining methods into a single geometry, rather than pitting them against each other, is the most appropriate framework for meeting modern requirements that span the fields of control, learning, and safety. The ideas left behind by Lyapunov are not a complete system, but rather profound insights that remain expandable. Thoughtfully reweaving these fragmented threads may well form the foundation.

References

[1] A. M. Lyapunov. *The General Problem of the Stability Of Motion.* Taylor & Francis, 1992.

[2] V. I. Zubov. *Methods Of A. M. Lyapunov And Their Applications.* P. Noordhoff, English edition by Leo F. Boron edition, 1964.

[3] D. G. Schultz and J. E. Gibson. The variable gradient method for generating liapunov functions. *Transactions of the American Institute of Electrical Engineers, Part II: Applications and Industry*, 81(4):203–210, 1962. doi:10.1109/TAI.1962.6371818.

[4] S. Sasaki and K. Uchida. A convex characterization of analysis and synthesis for nonlinear systems via extended quadratic Lyapunov function. In *Proceedings of the 1997 American Control Conference*, volume 1, pages 411–415 vol.1, June 1997. doi:10.1109/ACC.1997.611830.

[5] S. Sasaki and K. Uchida. Nonlinear H^∞ Control System Design via Extended Quadratic Lyapunov Function. *IFAC Proceedings Volumes*, 31(17):161–166, 1998. doi:10.1016/S1474-6670(17)40328-4.

[6] F. Forni and R. Sepulchre. A Differential Lyapunov Framework for Contraction Analysis. *IEEE Transactions on Automatic Control*, 59(3):614–628, 2014. doi:10.1109/TAC.2013.2285771.

[7] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview. *Annual Reviews in Control*, 52:135–169, January 2021. doi:10.1016/j.arcontrol.2021.10.001.

[8] C. Dawson, S. Gao, and C. Fan. Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control. *IEEE Transactions on Robotics*, 39(3):1749–1767, 2023. doi:10.1109/TRO.2022.3232542.

[9] 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*, 5:411–444, may 2022. doi:10.1146/annurev-control-042920-020211.

[10] J. Liu, Y. Meng, M. Fitzsimmons, and R. Zhou. Physics-informed neural network Lyapunov functions: PDE characterization, learning, and verification. *Automatica*, 175:112193, may 2025. doi:10.1016/j.automatica.2025.112193.

[11] Y. Meng and J. Liu. Towards Learning and Verifying Maximal Lyapunov-Barrier Functions with a Zubov PDE Formulation, November 2025. arXiv:2511.09523, doi:10.48550/arXiv.2511.09523.

[12] A. Baheri. Distributionally robust Lyapunov–Barrier Networks for safe and stable control under uncertainty. *Results in Control and Optimization*, 19:100556, June 2025. doi:10.1016/j.rico.2025.100556.

[13] R. S. Sutton and A. G. Barto. *Reinforcement Learning - An Introduction.* The MIT Press, second edition, November 2018.

[14] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal Policy Optimization Algorithms, August 2017. `arXiv:1707.06347`, `doi:10.48550/arXiv.1707.06347`.

[15] 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*, 62(8):3861–3876, 2017. `doi:10.1109/TAC.2016.2638961`.

[16] 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)*, pages 3420–3431, June 2019. `doi:10.23919/ECC.2019.8796030`.

[17] M. Z. Romdlony and B. Jayawardhana. Uniting Control Lyapunov and Control Barrier Functions. In *53rd IEEE Conference on Decision and Control*, pages 2293–2298, February 2014. `doi:10.1109/CDC.2014.7039737`.

[18] M. D. S. Aliyu. An approach for solving the Hamilton-Jacobi-Isaacs equation (HJIE) in nonlinear control. *Automatica*, 39(5):877–884, 2003. `doi:10.1016/S0005-1098(03)00025-6`.

[19] S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge University Press, March 2004. `doi:10.1017/CBO9780511804441`.

[20] J. Mawhin. Alexandr Mikhailovich Lyapunov, thesis on the stability of motion (1892). In *Landmark Writings in Western Mathematics 1640-1940*, pages 664–676. Elsevier Science, January 2005. `doi:10.1016/B978-044450871-3/50132-7`.

[21] P. Giesl and S. Hafstein. Review on computational methods for Lyapunov functions. *Discrete and Continuous Dynamical Systems - B*, 20(8):2291–2331, October 2015. `doi:10.3934/dcdsb.2015.20.2291`.

[22] W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998. `doi:10.1016/S0005-1098(98)00019-3`.

[23] S. Sasaki and K. Uchida. Presentation slides for ”Nonlinear H-infinity Control System Design via Extended Quadratic Lyapunov Function” (nolcos’98), August 2025. `doi:10.5281/zenodo.16663314`.

[24] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory.* Society for Industrial and Applied Mathematics, 1994. `doi:10.1137/1.9781611970777`.

[25] L. El Ghaoui, F. Delebecque, and R. Nikoukhah. *LMITOOL : A User-Friendly Interface for LMI Optimization , User’s Guide.* 1995.

[26] S. Prajna, A. Papachristodoulou, and P. Parrilo. Introducing SOSTOOLS: A general purpose sum of squares programming solver. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, volume 1, pages 741–746 vol.1, February 2002. `doi:10.1109/CDC.2002.1184594`.

[27] S. Prajna, A. Papachristodoulou, and P. A. Parrilo. SOSTOOLS - Sum of Squares Optimization Toolbox for MATLAB User’s guide, 2002.

[28] Y. Meng, Y. Li, M. Fitzsimmons, and J. Liu. Smooth converse Lyapunov-barrier theorems for asymptotic stability with safety constraints and reach-avoid-stay specifications. *Automatica*, 144:110478, October 2022. doi:10.1016/j.automatica.2022.110478.

[29] C. Dawson, Z. Qin, S. Gao, and C. Fan. Safe Nonlinear Control Using Robust Neural Lyapunov-Barrier Functions. In *5th Annual Conference on Robot Learning*, June 2021. URL: https://openreview.net/forum?id=8K5kisAnb_p.

Data Availability

This study is theoretical and does not involve any empirical datasets. All arguments are based on conceptual synthesis and historical analysis.

Funding

This research received no external funding.

Conflicts of Interest

The author declares no conflicts of interest.

A Note on Methodology and Transparency

This article proposes a synthesis across historically fragmented research traditions in stability theory. Consistent with that integrative approach, the manuscript was refined through iterative dialogue with multiple AI language models (Microsoft Copilot, OpenAI ChatGPT, Anthropic Claude, Google Gemini), which assisted with linguistic clarity and structural coherence.

All conceptual arguments, historical interpretations, and structural decisions originated from the author, grounded in over three decades of research in nonlinear control theory. The AI tools served as reflective interlocutors for refining expression and structure—a role documented here in recognition of evolving practices in scholarly communication. The author takes full responsibility for all content.