

DATA-DRIVEN ANALYSIS AND CONTROL OF DYNAMICAL SYSTEMS VIA RECURRENT LYAPUNOV FUNCTIONS

by
Roy N. Siegelmann

A dissertation submitted to The Johns Hopkins University in conformity
with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
August 2025

© 2025 Roy N. Siegelmann
All rights reserved

Abstract

Lyapunov stability theory has long served as a cornerstone in the study of dynamical systems and controller design, providing rigorous mathematical tools for certifying stability of equilibria. Lyapunov’s Direct Method, despite its ubiquity across engineering disciplines, requires the construction of a Lyapunov function whose sub-level sets are invariant — a condition that is both restrictive and computationally prohibitive in nonlinear settings. While numerous computational approaches, such as PDE solvers, linear and semidefinite programs, and sum-of-squares methods, have been proposed to automate Lyapunov function discovery, their scalability remains limited by exponential complexity. Furthermore, recent efforts to relax conditions on Lyapunov functions, including LaSalle-Krasovskii extensions, higher-order derivative conditions, and discretization methods, continue to preserve many of the original limitations.

This dissertation introduces a fundamental relaxation of invariance through the notion of recurrence. A set is defined to be τ -recurrent if every trajectory that leaves the set returns within time τ . Building on this concept, we develop the theory of Recurrent Lyapunov Functions (RLFs), which replace the invariance requirement with recurrence, thereby decoupling function design from the geometry of level sets. We establish general stability theorems demonstrating that τ -recurrence suffices to guarantee stability, asymptotic stability, exponential stability, and ultimate boundedness, and prove norm-agnostic converse results that highlight recurrence as a fundamental property of stability. To complement the theory, we propose a scalable, GPU-parallelized method for data-driven verification of recurrence conditions, enabling stability certification

without explicit Lyapunov function construction.

Extending these ideas to control, we introduce Recurrent Control Lyapunov Functions (R-CLFs) and develop Non-Parametric Chain Policies (NCP), a correct-by-design, data-driven approach to exponential stabilization. Unlike classical control Lyapunov methods, R-CLFs require only that the candidate function decrease sufficiently often along trajectories. NCP controllers achieve stabilization by chaining together verified control seeds, yielding scalable sample-complexity guarantees and lifelong refinement of the region of attraction. Numerical experiments on nonlinear benchmarks—including unicycle, bicycle, and inverted pendulum models—demonstrate the practicality of the approach.

By reframing stability around recurrence, this dissertation provides both theoretical foundations and computational tools for scalable, data-driven stability analysis and control, advancing the applicability of Lyapunov theory to modern dynamical systems.

Keywords: Dynamical Systems, Lyapunov Theory, Recurrence, Control Theory, Sampling, Non-Parametric

Primary Reader and Thesis Advisor

Dr. Enrique Mallada
Associate Professor
Department of Electrical and Computer Engineering
Johns Hopkins University, Baltimore MD

Secondary Readers

Dr. Donniell Fishkind
Associate Professor and Director of Undergraduate Studies
Department of Applied Mathematics and Statistics
Johns Hopkins University, Baltimore MD

Dr. Louis Whitcomb
Professor
Department of Mechanical Engineering
Johns Hopkins University, Baltimore MD

Dr. Richard Brown
Teaching Professor and Director of Undergraduate Studies
Department of Mathematics
Johns Hopkins University, Baltimore MD

Acknowledgement

These past five years at Johns Hopkins University have been a considerable journey. Now, as the end of this journey approaches, I would like to take this opportunity to thank all those who helped and supported me throughout this process and before without whom I would never have been able to make it this far.

First and foremost, my utmost gratitude to my advisor, Enrique Mallada. Beyond simply introducing me to the problem of the relaxation of Lyapunov functions, his continual guidance and direction have been invaluable throughout this entire process. His recommendations, suggestions, and commentary throughout my doctorate have helped guide the path I have been able to walk. From his insight into difficult problems and their disentangling to his endless patience with me when I struggle, from teaching me how to make good PowerPoints and present them well to forgiving my disorganization, I am endlessly thankful. I simply could not have done it without Enrique.

I am deeply grateful and indebted to two other professors here at Johns Hopkins. Carey Priebe was the reason that I came to Johns Hopkins in the first place, and his endless wisdom and compassion have provided a guiding light for me throughout my time here. My love of mathematics would likely have floundered years ago if not for the boundless joy Carey infuses into the subject. Next, to my friend and mentor Donniell Fishkind. I find it hard to imagine anyone more selfless and willing to drop everything to help at a moment's notice. From 3am Zoom calls for course recommendations

Acknowledgement

to listening to my endless ramblings about one new research direction or another, Donniell provided me with more than I can hope to repay.

My thanks and appreciation go out to two professor emeriti I have had the pleasure of interacting with over the years. First, to Monroe Rabin of the University of Massachusetts, Amherst. Monroe was the professor of my first university-level class and first physics class. He taught me how to approach problems and, perhaps more importantly, that I am capable of solving them through hard work and determination. Second, to Jim Yorke of the University of Maryland. For as long as I can remember, with a twinkle in his eye Jim would teach me to love mathematics and approach it with a sense of wonder and joy. Beyond inspiring me to start on this journey, his guidance has kept me excited throughout all these years. I am honored and privileged to have been able to continue in his footsteps in the study of dynamical systems.

To my ninth-grade English teacher, Daniel Page, I owe my skills in writing as well as my ability to think deeply about problems from many distinct angles, without which I could not conduct research. Beyond this, as my model for a Renaissance man, I owe to him my drive to pursue a wide variety of interests, without which I would not be who I am today. I can still hear his voice in my head when I make silly mistakes and would have it no other way.

To my wonderful friends and colleagues who through friendship and kindness, through aid and guidance, have inspired me and kept me sane throughout these many long years. Though there are too many to name, I appreciate each and every one of them beyond measure – they form the stars by whose constellations in the night sky I

Acknowledgement

navigate.

I would not be able to finish my acknowledgements without acknowledging my grandparents, to whom I owe everything. First, to my grandfather, Joseph Elijah Siegelmann, may his memory be a blessing. He was a father to me, who taught me diligence and determination, who taught me to take responsibility for my actions and obligations, and who is absolutely the best man I have ever known or will ever know. Next, to my grandmother, Hanna Siegelmann. Her boundless love and warmth cannot be understated, giving me the environment and opportunity to grow, and her deep well of knowledge across many disciplines has never failed to provide me with something novel and exciting to learn.

Finally, greatest of them all, to my mother, Hava Tova Siegelmann. Every person I have named to whom I am grateful has helped shape who I am today, but my mother made me who I am today. She has provided unconditional love and support while also correcting my path when I am lost. She is beyond a shadow of a doubt the most brilliant and insightful person I know, guiding my academics and research alike. She has given everything to be my mother, my teacher, my mentor, my closest confidante, my protector, and my support. Every moment of every day she has been there - and continues to be there - to provide what I need, even if I do not yet know what that is myself. Everything I know and everything I have accomplished and will accomplish I owe to her. There are no number of words I can write that can ever express even a sliver of my gratitude.

Order is not imposed from above but discovered from within.

– LORD RABBI JONATHAN SACKS

Table of Contents

Abstract	ii
Acknowledgement	v
Epigraph	viii
List of Tables	xi
List of Figures	xii
Chapter 1 Introduction	1
Chapter 2 Background	8
2.1 Dynamical Systems	8
2.2 Control Theory	21
Chapter 3 iuuRelated Work in Dynamical Systems and Control .	29
3.1 Relaxations of Lyapunov Functions	29
3.2 Explicit Model Predictive Control	33
Chapter 4 Relaxations of Lyapunov Theory via Recurrence . . .	36
4.1 Preliminaries	36
4.2 Recurrence	38
4.3 Recurrent Lyapunov Functions	47
4.4 Asymptotic Stability	50
4.5 Exponential Stability	56
4.6 Norm-Based Converse Theorems	63
4.7 Verification of Exponential RLFs	66
4.7.1 Trajectory-based Verification of a Neighborhood	67

Table of Contents

4.7.2	Trajectory-based Verification of a Region	69
4.8	Numerical Methods	74
4.8.1	Supporting Methods and Algorithms	74
4.8.2	Verification of a Region	78
4.8.3	Exponential Region of Attraction Mining	79
4.8.4	Numerical Validation of Algorithm 3	80
4.8.5	Numerical Validation of Algorithm 4	83
Chapter 5	Non-Parametric Chain Policies for Stabilizing Control	86
5.1	Preliminaries	86
5.1.1	Practical Exponential Stabilizability	87
5.1.2	Recurrent Lyapunov Functions	88
5.2	Recurrent Control Lyapunov Functions	90
5.3	Non-Parametric Chain Policies	96
5.3.1	Convergence Guarantees of NCPs	99
5.3.2	Existence and Sample Complexity of NCPs	103
5.3.3	Incremental Learning of NCPs	106
5.4	Numerical Experiments	110
5.4.1	Unicycle	112
5.4.2	Inverted Pendulum	113
Chapter 6	Conclusion	117
	Bibliographic references	122

List of Tables

Table 4.1	Parameter values and performance comparison between our algorithm and the SOSTOOLS for system (4.59)	83
Table 4.2	Parameter values and performance comparison between our algorithm and the SOSTOOLS for system (4.60)	83

List of Figures

Figure 4.1	Splitting a Ball according to Algorithm 2	75
Figure 4.2	Illustration of Initial Grid Setup.	77
Figure 4.3	Phase Portrait of System (4.59)	82
Figure 4.4	Sizes of blocks resulting from applying Algorithm 3 to system (4.59).	82
Figure 4.5	Phase Portrait of System (4.61), 1 Split.	84
Figure 4.6	Phase Portrait of System (4.61), 6 Splits	85
Figure 4.7	Area of region Verified by Algorithm 4 for system 4.61 and Associated Computational Time Requirements.	85
Figure 5.1	Trajectories of Unicycle NCP. Phase plots of (x, y) for eight evenly distributed points. The black icons depict the initial facing of the unicycle. Plot (a) contains trajectories from NCP trained to minimize V_1 , which results in sharp turns, while (b) is trained to minimize V_2 , which results in softer turns and smoother overall behavior. Plots (c) and (d) show the development of V_1 and V_2 over time respectively. Both converge exponentially to the equilibrium, with at least $\alpha = 0.01$. We have $\tau_{\max} = 5, \varepsilon = 0.01, L = 1$ and $c = \varepsilon(1 + L\tau_{\max}e^{L\tau_{\max}}) \simeq 0.613$, represented by the dotted line.	114
Figure 5.2	Incremental Learning of Unicycle Policy. Extending the state space from the previously learned region in the y -direction. Subfigure (a) contains the phase plot before learning, while subfigure (b) contains the phase plot after. The new region is learned without forgetting, such that parts of the trajectory in the old region use previously designed controls.	115

Figure 5.3 **Additional Data Refinement Facilitates Improved NCP Performance.** Plot (a) contains the balls used to verify the region $(\theta, \dot{\theta}) \in (-\pi, \pi] \times [-5\pi, 5\pi]$ for the inverted pendulum. Plot (b) is a refinement of plot (a), wherein all balls were split once more and re-verified. The minimum verified rate of convergence for trajectories α goes from 0.003 to 0.0145, and the average verified α goes from 1.815 to 3.149. Plot (c) demonstrates the average norm over time of 400 sample trajectories under each schema. We have $\tau_{\max} = 1.5$, $\varepsilon = 0.01$, $L = 1$, and $c = \varepsilon(1 + L\tau_{\max}e^{L\tau_{\max}}) \simeq 0.072$ 116

Chapter 1

Introduction

*Beginnings are the most important of times, because they
carry within them the seeds of all that will follow.*

– LORD RABBI JONATHAN SACKS

Lyapunov stability theory plays a central role in the study of dynamical systems. It provides a rigorous mathematical framework for qualitatively analyzing system solutions and has heavily influenced systems theory and engineering over the past century. A fundamental tool derived from this theory is the so-called Lyapunov direct method, a.k.a. Lyapunov’s second method [74], which states conditions on a function $V(x)$ (non-increasing along trajectories and proper) that can certify stability of an equilibrium point. Since first proposed in 1892, Lyapunov’s direct method has found ubiquitous applications across multiple branches of engineering, including aerospace, electrical, mechanical, and chemical, among others [62, 96, 102, 112].

A critical step in the application of Lyapunov’s direct method is finding the function V that indeed satisfies all the conditions stated by the theory. Unfortunately, while such a function is known to exist via converse theorems, e.g., [79], manually finding a Lyapunov function is often tricky and relies on ingenuity and deep domain knowledge. To circumvent this step, a variety of computational methods have been proposed for finding Lyapunov functions [44], e.g., via the use of partial differential equation (PDE) solvers to solve Zubov’s Equation [50, 121], linear programs (LPs) to find piece-wise linear Lyapunov functions[58], and semidefinite programs (SDPs) to solve

linear matrix inequalities (LMIs) [45] or sum of square (SoSs) problems [95]. However, the computational complexity is known to exponentially increase with not only the dimension of the state space but also the parameterization of the Lyapunov function [3, 44].

This has led to multiple investigations into relaxing the conditions required for V , and in particular, its time derivative \dot{V} . Such relaxations can be broadly divided into three groups. The first group seeks LaSalle-Krasovskii type of conditions by relaxing the negative definiteness of \dot{V} , i.e., only requiring $\dot{V} \leq 0$; see [62, 99] and its generalization [78, 84]. The second group further relaxes the strict negative definite condition Lyapunov method by allowing $\dot{V} > 0$ on some regions of the state space. This is implicitly done by using generalizations of the comparison lemma [46] to impose conditions on higher-order time derivatives of V that still ensure convergence of $V \rightarrow 0$ while allowing $\dot{V} > 0$ for some regions of the state space. The third group uses the so-called discretization method, which considers a fixed parameter $T > 0$ and leverages the net decrement of V across any trajectory $x(t)$, i.e., $V(x(t+T)) - V(x(t))$, to reason about stability [1, 35]. Unfortunately, despite such efforts, the basic principle can still be traced back to the (indirect) construction of a Lyapunov function whose sub-level sets are invariant [2, 59], which still needs to be verified either analytically or via the solution of a convex program, rendering similar verification challenges as before.

The challenge of finding such functions lies in the fact that Lyapunov's direct method implicitly constrains its shape by requiring every sub-level set to be an invariant set. The goal of this dissertation is to relax this condition by replacing the invariance of

sub-level sets with a weaker notion known as recurrence. We say that a set is (τ -)recurrent if every trajectory that starts in the set returns to it (within τ seconds). Such relaxation has been recently shown to provide a versatile mechanism for estimating regions of attractions of stable equilibrium points [103] as well as verifying the safety of a dynamical system [104]. Moreover, from an information theoretical standpoint, (control) recurrence can be achieved at lower data rates than invariance [105] and can often be enforced using a finite number of trajectory data [106].

In Chapter 4, we seek to explore the role of recurrence in certifying different notions of stability of an equilibrium point. The contributions of our work are several:

- *Recurrent Lyapunov Functions:* We introduce the concept of Recurrent Lyapunov Functions (RLFs), which generalize classical Lyapunov functions by replacing the invariance condition on sub-level sets with a more flexible recurrence condition. This relaxation decouples the geometry of the trajectories from the geometry of the level sets.
- *Stability Guarantees:* We establish rigorous stability theorems demonstrating that τ -recurrence of a sequence of compact sub-level sets is sufficient to guarantee stability, asymptotic stability, exponential stability, and ultimate boundedness. These results provide an alternative framework for proving stability without requiring strict invariance conditions.
- *Norm-agnostic Converse Theorems:* We show that, under mild assumptions, any norm is guaranteed to satisfy our RLF conditions, provided one is willing to look for a slightly weaker stability property. This generality highlights the fundamental

role of recurrence in stability analysis, and opens the door for the development of stability verification methods that do not require the computation of a Lyapunov function.

- *Data-Driven Verification:* We develop a computationally efficient, GPU-parallelized method to verify whether a given function, e.g., a norm, satisfies the proposed RLF conditions on a compact domain. Our approach enables scalable data-driven verification of dynamical systems, making stability analysis feasible even when an explicit Lyapunov function is difficult to construct.

In Chapter 5, we extend the new theory to encapsulate stabilization of control systems. *Data-driven control methods* offer a novel paradigm for synthesizing controllers directly from trajectory observations, potentially bypassing the need for accurate system models while reducing computational burden and conservativeness of classical control synthesis [112, 135]. Recent years have witnessed significant progress in data-driven control. These approaches, as well as their level of maturity, depend considerably on the underlying system properties. For *linear systems*, the field has substantially matured: LMI-based formulas [38] and convex programs [19, 36] can transform trajectories into stabilizing feedback controllers with robustness [18, 19, 38, 120], performance [19, 36, 39] and sample complexity guarantees [27, 126].

For *nonlinear systems*, several approaches have been proposed, with methods highly dependent on the implicit assumptions made on the nonlinear system class and the control synthesis methodology. One prolific line of works considers dynamics formed from dictionary-based hypothesis classes—e.g., using polynomials [37, 47, 61], fractions

of polynomials [114], or general nonlinear functions [88]—and formulate semidefinite programs that render policies with a wide variety of guarantees, including contraction-based stability [54, 93, 119] or robustness [61]. Other methods employ general learning techniques to learn models or policies and leverage intrinsic system properties to provide different guarantees, e.g., Koopman operator methods that exploit spectral properties [55], sample complexity analysis for stochastic dynamics [28], and conformal prediction approaches for statistical robustness [53].

Despite the effectiveness of these methods in synthesizing controllers with guarantee, many questions remain unanswered. First, sample complexity guarantees are typically technique-dependent and do not provide clear understanding of how data requirements scale with explicit system properties, such as state dimension or attainable performance levels, or the specific hypothesis class considered. Second, computational complexity of optimization-based methods scales poorly with dictionary size and state dimension. Third, incorporating new data necessitates resolving the underlying optimization problem, often requiring complete recomputation and discarding previous work. As a result, there remains a need for flexible data-driven approaches that can adapt to new information without structural assumptions while providing transparent performance-data trade-offs.

To address these challenges, we introduce the concept of *Nonparametric Chain Policies* (NCPs), a data-driven approach that requires only Lipschitz assumptions on the system dynamics while providing explicit sample complexity guarantees for practical stabilization. NCPs employ a normalized nearest-neighbor rule to assign finite-duration control signals from a stored library of verified trajectories, enabling direct use of data

without parametric modeling or optimization re-solving when new data arrives. Our theoretical guarantees build on the framework of Recurrent Lyapunov Functions [109, 110], which we extend here for the control setting by introducing here the notion of Recurrent Control Lyapunov Functions (RCLFs, Section 5.2).

Contributions. Our approach offers three key advantages over existing methods:

1. **Explicit sample complexity:** NCPs achieve practical exponential stabilization using $\mathcal{O}\left(\left(3/\rho\right)^d \log(R/\delta)\right)$ sample trajectories, with transparent scaling in dimension d , target radius R , precision δ , and a system-dependent parameter ρ .
2. **Incremental learning:** The nonparametric nature of NCPs allows for new verified data to be seamlessly incorporated to expand a certified region, or improve performance, without discarding previous guarantees or re-solving optimization problems.
3. **Performance-complexity trade-offs:** The framework explicitly controls the trade-off between sample requirements and performance through a user-specified parameter ρ that relates best achievable performance and the performance guaranteed by the NCP.

The remainder of this thesis is organized as follows: Chapter 2 covers the background of the mathematical foundations of our primary contributions. It is divided into Section 2.1, which contains an extensive introduction to the field of dynamical systems for the layman, followed by Section 2.2, which gives the same treatment to control theory. We recommend that those familiar with either discipline skip these subsections and

continue to the next section. Topic-specific information more relevant to individual contributions are contained in the introduction sections for the relevant chapters. Chapter 3 covers related work in dynamical systems and control theory in Sections 3.1 and 3.2 respectively. Chapter 4 contains our primary contribution to the realm of dynamical systems, namely the introduction of Recurrent Lyapunov Functions, their necessary and sufficiency for stability analysis, and a fully parallelizable algorithm for stability verification along with experiments demonstrated its utility. Chapter 5 extends our dynamical system results and theory to the realm of control and introducing the paradigm of NCP control. We offer our concluding thoughts in Chapter 6, followed by the bibliography of references cited.

Chapter 2

Background

To be intelligible to others, we must first learn to speak in the language of others. Only then can we begin the conversation.

– LORD RABBI JONATHAN SACKS

This section is dedicated to providing a complete background, not only to the specific topic of Lyapunov functions, but more broadly to the basics of dynamical systems and control theory, as to enable readership from throughout the scientific and engineering communities, including outside of mathematics. We aim to increase intuition and understand the current difficulties which this thesis is targeted at solving.

2.1 Dynamical Systems

Dynamical systems are a key notion in mathematics and a powerful tool for analysis in many engineering and scientific disciplines. We can intuitively consider a dynamical system as "a system which changes deterministically", i.e. one where we know based on the current state what the state will be at every future time. Consider John Doe, a man who commutes daily to work via the same path, and before traffic, so that given a location and time, he knows more or less where he will be at any subsequent time along the path. In this example, his position along his route forms a dynamical system.

Dynamical systems are comprised of two key components, a **state space** and a **mapping**. The state space is - as the name suggest - simply a grouping of all possible

states of the system. A state space can be a one-dimensional system (what we call a scalar system), such as the number of animals of a given sort in a population, the representation of an electric circuit, or the distance of a tightrope walker along the string. A state space can also just as easily be a higher-dimensional system (what we call a vector system), such as the current stock of seventy different items in a grocery store (70 dimensions), the position of ten different objects in three-dimensional space along with their velocities and accelerations (90 dimensions), or the relative angles of fifteen joints in an under-actuated robot arm (15 dimensions). In Johns's case, we consider the state-space to be the collection of all physical positions along his route to work as a two-dimensional system, i.e. his coordinates on the map. The mapping, usually denoted by the Greek letter ϕ , is a function which accepts an amount of time and an initial state, and outputs the final state of the system assuming that we start at the initial state and let the prescribed amount of time pass. If we let t be the time, x be the initial state, and y be the final state, the particular relation using this mapping can be expressed as $y = \phi(t; x)$.

Back to John Doe, where he has empirically derived the mapping through trial and error. An example case of the mapping is as follows: Beginning at a given location, e.g., the corner of Main Street and Park Avenue, he knows that ten seconds later he will be at another particular location, e.g., the corner of Main Street and Oak Drive. As such, this could be written as

$$\phi(10; \text{Main and Park}) = \text{Main and Oak}$$

A key point to note is that in dynamical systems, the mapping must be well-defined at

every point of the state space. This is why we defined the state space in John's case as only the points along his regular route rather than the entirety of the state, country, or world; if there is an unexpected detour due to construction and he is forced to drive along another road, the mapping begins to fall apart. Due to lack of knowledge regarding traffic light timings on the side street and an upheaval in regular traffic patterns, he is no longer able to deterministically predict where he will be at each time in the future along this ride.

The eagle-eyed reader may have already noticed an issue with the determinism of the mapping. While John Doe is driving toward work, ten seconds from Location 1 he will be at Location 2. But when he is driving home, ten seconds from the Location 1 he will be at Location 3. To disentangle this issue we can include the time of day as part of the state space: if it is around 8am he is on the way to work, while if it is around 5pm he is on the way home. We can include a t_0 term in our mapping expression which denotes the initial time, i.e. the time which the state is x , such that $y = \phi(t; x, t_0)$. This lets us write that

$$\phi(10; \text{Location 1}, 8\text{am}) = \text{Location 2}$$

whereas

$$\phi(10; \text{Location 3}, 5\text{pm}) = \text{Location 3}$$

The former system, wherein the initial time does not impact the behavior of the mapping, is called a **time independent system**. The latter system, wherein the initial time alters the behavior of the mapping, is naturally referred to as a **time**

dependent system. By default and in the rest of this thesis, we will assume that systems are time independent. There are some who refer to time independent systems as autonomous systems and time dependent systems as nonautonomous systems. However, we will avoid doing so to avoid ambiguity and abuse of terminology since we will later use the term "autonomous" to refer to another type of dynamical system as is conventional. As a side note, we will use the following common mathematical notation in this manuscript:

- The capital letter S in this section will denote the state space of a dynamical system
- The letter \mathbb{R} will refer to the set of all real numbers, with \mathbb{R}^+ referring to the set of all positive real numbers, and $\mathbb{R}_{\geq 0}$ referring to the set of all non-negative real numbers
- The expression \mathbb{R}^k will refer to the set of k -entry vectors with elements drawn from \mathbb{R}
- $x \in S$, read as " x in S ", denoting the set from which element x is drawn
- $f : S_1 \rightarrow S_2$, read as " f from S_1 to S_2 ", denoting a function f which maps elements of the set S_1 to the set S_2
- \dot{f} , read as " f dot", denoting the derivative of a differentiable function f
- $S_1 \subset S_2$, read as " S_1 subset S_2 ", denoting that the set S_1 is a subset of set S_2 , i.e. that S_1 is a set comprised only of elements which belong to S_2

- $\exists x$, read as “there exists an x ”, denoting the existence of an element satisfying certain properties
- $\forall x$, read as “for all x ”, denoting that all elements in a particular set satisfy certain properties
- $x \mid$, read as “ x , such that”, denoting the restriction of elements to those which satisfy a certain property
- $\|x\|$, read as “norm of x ”, denoting the magnitude of vector x (the absolute value for scalars). If we instead insert a difference expression, e.g. $\|y - x\|$, this represents the distance from y to x .

These will be combined in such ways as $\exists t > 0 \mid \forall x \in S, \phi(t; x) \in S$, to be read as “there exists a t greater than zero such that for all x in set S , the solution of x after time t is contained in set S ”.

Now, let us consider only the forward direction of John Doe’s journey, i.e. that he begins at home, ends his trip at work, and stays there. Thus once he arrives at work, he will not continue moving in any direction. This fact makes his workplace an **equilibrium point** or **fixed point** of the dynamical system (along with other, less-recommended equilibrium points such as falling into the river), i.e. a point of the state space which does not change with time. Mathematically, we can write that a point x^* of the state space is an equilibrium point if for all times t , $\phi(t; x^*) = x^*$. There are many kinds of equilibrium points, each with distinct associated properties, but they all have the property of remaining there forever. In many systems, such as

our example with John Doe, the equilibrium points can form targets, with the 'goal' being to reach a fixed point and stop moving. John needs to get to work and stay there, so his work building is an equilibrium point. Under this framework, and in the field of Dynamical Systems in general, a number of questions about (and corresponding properties of) the system immediately jump out at us:

1. How far away does the system drift from an equilibrium point, assuming it starts sufficiently close?
2. Is the system guaranteed to converge to the equilibrium point, achieving its target?
3. Assuming the system converges to equilibrium, how fast does it do so?

After the formalization dynamical system by the "The Last Universalist" and polymath par excellence Henri Poincare, the first person to consider these questions of equilibria was Aleksandr Lyapunov. His scientific career is quite interesting: Lyapunov was born in 1857 to an astronomer father, and at 19 entered the university as a physics student, following in his father's footsteps. However, he was inexorably drawn to mathematics and changed departments in just over a month (from the Physico-Mathematical Department to the Mathematics Department) to study under the stellar statistician and Father of Russian Mathematics Pafnuty Chebyshev, although he never gave up on his love for physical systems, particularly fluid bodies and hydrostatics. In 1892, Lyapunov introduced formalizations of the questions about equilibrium points, designing algorithms and providing theoretical guarantees for a wide variety of systems under many different conditions. Lyapunov referred to the first property - not going

too far away from an equilibrium point - as **stability**. An equilibrium point is called stable if for any $\varepsilon > 0$ we can define δ such that trajectories which start within δ of the equilibrium point will never grow to a distance greater than ε from the same equilibrium point. Mathematically, for an equilibrium point $x^* \in S$ we can write

$$\forall \varepsilon > 0 \exists \delta > 0 \mid \forall x \in S \left(\|x - x^*\| < \delta \Rightarrow \|\phi(t; x) - x^*\| < \varepsilon, \forall t \geq 0 \right).$$

Lyapunov called the second property - converging to an equilibrium point - **asymptotic stability**. An equilibrium point is said to be asymptotically stable if it is stable and attractive. **Attractivity** is defined to be such that for every sufficiently close starting point, will eventually draw arbitrarily close to the equilibrium point. Mathematically, for an equilibrium point $x^* \in S$ we can write

$$\exists \delta > 0 \mid \forall x \in S \left(\|x - x^*\| < \delta \Rightarrow \forall \varepsilon > 0 \exists T > 0 \forall t > T : \|\phi(t; x) - x^*\| < \varepsilon \right),$$

or equivalently and more succinctly

$$\exists \delta > 0 \mid \forall x \in S \left(\|x - x^*\| < \delta \Rightarrow \lim_{t \rightarrow \infty} \|\phi(t; x) - x^*\| = 0 \right).$$

Finally, Lyapunov named the third property - having a minimum rate at which convergence occurred - **exponential stability**. An equilibrium point is said to be exponentially stable with rate α if every point not only converges to the equilibrium, but converges to the equilibrium point x^* faster than the exponential function $e^{-\alpha t}$.

Mathematically, for an equilibrium point $x^* \in S$ we can write

$$\exists \delta > 0, k \geq 1, \alpha > 0 \mid \forall x \in S \left(\|x - x^*\| < \delta \Rightarrow \|\phi(t; x) - x^*\| \leq ke^{-\alpha t} \|x - x^*\|, \forall t \geq 0 \right).$$

These three conditions are called **local** stability conditions. Similar conditions which are more stringent are called **global** stability conditions, which require the aforementioned conditions to hold for all points, rather than just those which start sufficiently close to the equilibrium point.

As mentioned above, Lyapunov not only studied these properties of equilibria, but gave conditions which can be used to guarantee such properties. Indeed, checking that each and every one of uncountably infinitely many points in a set - or even in a small ball around the equilibrium - satisfies a condition is impossible. Therefore, relying on verifiable alternative conditions which give the desired results is absolutely imperative. To discuss Lyapunov's theory, we must first introduce some definitions. A function $f : D \subset S \rightarrow \mathbb{R}$ is called **positive semidefinite** (respectively **negative semidefinite**) over a set $D \subset S$ if $\forall x \in D, f(x) \geq 0$ (respectively $f(x) \leq 0$). Such a function is called **positive definite** (respectively **negative definite**) over a set $D \subset S$ if $f(0) = 0$ and $\forall x \in D \setminus \{0\}, f(x) > 0$ (respectively $f(x) < 0$). If a function is neither positive nor negative semidefinite, it is referred to as being **indefinite**. As an aside, going forward we will assume without loss of generality that x^* , the equilibrium point of interest, is at the origin, i.e. at zero. While this may seem like an overly strong assumption, no information is actually lost - we can simply reparametrize the system such that the equilibrium point is at the origin, essentially defining the origin as our equilibrium point of interest. As such, we will use the term origin to refer to

our equilibrium point of interest going forward.

Lyapunov had two primary methods for determining the stability of a system. The first method, known as **Lyapunov's Indirect Method**, relies on taking a linear approximation around the origin and analyzing the resulting linearized system for stability. This is quite intuitive - given a linear mapping of the form $\dot{x} = Ax$, if $A < 0$ (analogously, in higher dimensions, the real parts of the eigenvalues of matrix A are all negative) then the trajectory will always converge to origin, and the system will be asymptotically stable. Similarly, if $a > 0$ (analogously, in higher dimensions, the real parts of the eigenvalues of matrix A are all positive) then the trajectory will always diverge from the origin, and the system will be unstable. Therefore the origin original system will be locally asymptotically stable (respectively unstable) if same property is satisfied by the origin of the linear system. While quite easy to perform, this method is limited to asymptotic stability, being unable to make claims about stability and exponential stability. Furthermore, this method is limited to local stability, being unable to make global claims. Finally and perhaps most unfortunate is the fact that in higher-order systems, many systems have eigenvalues on both sides of zero, which results in a failure of the Indirect Method to yield any claims whatsoever. This method is thus incredibly easy to implement, but equally limited in its potential use cases.

The second method, known as **Lyapunov's Direct Method** (commonly referred to as Lyapunov's Theorem) offers significantly more flexibility of use cases and benefits in terms of variety of possible claims. This method is quite straightforward, relying only on designing a **Lyapunov function**, a function $V(x) : D \subset S \rightarrow \mathbb{R}$ satisfying certain properties key to the method. The procedure is as follows: select a continuously

differentiable (continuous with a continuous derivative) function $V(x) : D \subset S \rightarrow \mathbb{R}$ which is positive definite on D . Then, the following are true:

1. If $\dot{V}(x)$ is negative semidefinite on D , then $V(x)$ is a stability Lyapunov function and the equilibrium point at the origin is stable
2. If $\dot{V}(x)$ is negative definite on D , then $V(x)$ is an asymptotic Lyapunov function and the equilibrium point at the origin is asymptotically stable
3. If there exists an $\alpha > 0$ such that $\dot{V}(x) \leq -\alpha V(x)$ for all $x \in D$, then $V(x)$ is an exponential Lyapunov function and the equilibrium point at the origin is exponentially stable with rate α .

In each of these cases, the stability assertions are about local stability. However, they can be made about global stability by taking $D = S$ and requiring the condition of V being **radially unbounded**, i.e. that $V(x) \rightarrow \infty$ as $x \rightarrow \infty$.

Let us look at an example to illustrate the utility of Lyapunov's Theorem. Consider a pendulum with friction, which can be written as the dynamical system

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -a \sin(x_1) - bx_2 \end{bmatrix}$$

for $a, b > 0$, wherein $S = (-\pi, \pi) \times \mathbb{R}$. Consider the candidate Lyapunov function $V(x) = a(1 - \cos(x_1)) + \frac{1}{2}x_2^2$. We know that $V(0) = 0$ since

$$V(0) = a(1 - \cos(0)) + \frac{1}{2} \cdot 0^2 = a(1 - 1) + 0 = a \cdot 0 + 0 = 0 + 0 = 0$$

Additionally, if $x_1 \neq 0$, $\cos(x_1) < 1$ so $a(1 - \cos(x_1)) > 0$. Similarly, if $x_2 \neq 0$, $\frac{1}{2}x_2^2 > 0$. Therefore, $V(x) > 0$ everywhere but the origin, so V is positive definite. We have by differentiation that

$$\dot{V}(x) = a \sin(x_1) \cdot \dot{x}_1 + x_2 \cdot \dot{x}_2 = a \sin(x_1)x_2 + x_2(-a \sin x - bx_2) = -bx_2^2$$

This is negative semidefinite, since the square of a number is always nonnegative, but is not negative definite since $V(x) = 0$ at all points where $x_2 = 0$, not only at the origin. Thus, we can conclude that the origin is stable, but not that it is asymptotically stable.

However, the conditions of Lyapunov's Theorem not being satisfied implies lack of evidence of stability, not evidence proving lack of stability - they are sufficient, but not necessary. To show this, let us select a more appropriate candidate for a Lyapunov function, namely

$$\begin{aligned} V(x) &= a(1 - \cos(x_1)) + \begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} b^2/2 & b/2 \\ b/2 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ &= a(1 - \cos(x_1)) + \frac{1}{2} \left(\frac{b^2}{2} \cdot x_1^2 + bx_1x_2 + x_2^2 \right) \end{aligned}$$

Through some Linear Algebra and the entries of the matrix, we know that the second term of $V(x)$ is positive definite and therefore $V(x)$ is positive definite, as the sum of positive definite functions is positive definite and we have already demonstrated the positive definiteness of the first term of $V(x)$ above. Taking the derivative, and

simplifying we have

$$\dot{V}(x) = -\frac{b}{2}(ax_1 \sin(x_1) + x_2^2)$$

which is negative definite, since $x_1 \sin(x_1) > 0$ for all nonzero $x_1 \in (-\pi, \pi)$ and $x_2 > 0$ for all $x_2 \in \mathbb{R}$. As such, simply by changing the Lyapunov function, we can now assert that the origin is asymptotically stable. This highlights a key problem with Lyapunov's Theorem. Finding an appropriate Lyapunov function can be incredibly difficult, and designing an appropriate one generally requires the system to have a known parametrization with properties that can be utilized in the function design process.

Recall the illustration using John Doe. One might be drawn to think that Euclidean distance to the origin - in this case his work building - forms a good Lyapunov function, as this distance is clearly positive definite along his route. However, we encounter a problem: While he advances toward his job along the selected path, the Euclidean distance may actually grow between him and his workplace (such as driving from his house to get onto the main street) so using this candidate Lyapunov function would fail due to its derivative being positive on the first leg of his drive. Perhaps a better choice would be Manhattan distance, i.e. the number of intersections between two points. This is certainly positive definite and non-increasing, so would seemingly form a good Lyapunov function. However, let us extend the state space to include a second viable route, namely taking the highway to work. In order to drive to the highway, he needs to actually drive away from work, increasing the Manhattan Distance, but will eventually reach work and in fact will reach work faster than taking the normal route. Thus, while in truth the origin of the system is asymptotically stable, we cannot even

prove stability with the candidate Lyapunov functions.

Beyond the difficulty of constructing a Lyapunov function in the general case, Lyapunov Theory has two major drawbacks: lack of robustness and lack of extendability. Changing the underlying dynamics (the mapping) or the state space even by the smallest amount can entirely invalidate a previously acceptable Lyapunov function, sending the designer back to the drawing board. Lyapunov gave a converse theorem, one which guarantees that a stable (respectively asymptotically stable, exponentially stable) equilibrium point must have a Lyapunov function fulfilling the necessary conditions. Despite this and despite many tips and tricks for finding Lyapunov functions based on specific domain knowledge developed over the past thirteen decades by Lyapunov and his disciples, there is no broadly-applicable method for the design of Lyapunov function in the general case. A major field of research for the past century has been in relaxing the constraints of Lyapunov functions, e.g. by allowing $\dot{V}(x)$ to be positive in certain small areas, by requiring conditions on higher-order derivatives, or by discretizing the system, we will discuss in the Related Work section. However, these methods tend to suffer from the same key issue: how do we design a function which satisfies certain conditions in the general case of a dynamical system wherein the dynamics are complex, or perhaps even not fully parametrized?

The difficulties with Lyapunov function design are what led to the first key contribution of the research presented in this manuscript. What if we could eliminate the requirement of designing a Lyapunov function altogether? This novel method, utilizing what we call Recurrent Lyapunov Functions (RLFs) is discussed in-depth in Chapter 4. This method presents an *alternative* to Lyapunov function design which can verify

the stability properties of a system entirely computationally without necessitating extensive trial-and-error, guesswork, or particular domain knowledge and expertise, and can be applied *uniformly* to *any* dynamical system. Without ever resorting to painstakingly studying the parametrization of the dynamical system and designing a Lyapunov function such that its properties along trajectories align with the geometry of the system, we start from the goal and work backwards. By extending Lyapunov's Theorem such that the conclusions apply with significantly more lenient requirements, we prove that we can take effectively any distance metric and use it in place of a Lyapunov function. The process then consists of proving that the metric selected has the desired properties when applied to the dynamical system, which becomes achievable due to the easy-to-verify requirements and their computational efficiency.

A key benefit of our method is its practicality: as no function was designed around a given state space, the method is robust to changes in the underlying system - or lack of parametrization entirely. Furthermore, the method is extendable - any conclusions derived will hold even if the state space is later expanded. This method is also perfectly suited to the computerized modern world, replacing the difficult problem of analysis and design with a straightforward computational algorithm.

2.2 Control Theory

We next turn to discuss what are known as **dynamical systems with control**, or **control systems** for short. Generally, dynamical systems which do not depend on an input are also known as **autonomous systems** as they are fully deterministic insofar as they do not depend on any signal external to the system itself. A control system is

simply a dynamical system with the addition of an input term, commonly denoted by the lowercase letter u . For example, the dynamical system described by $f(x) = x^2$ can be turned into a control system by adding a control, such as $f(x, u) = x^2 + u$. An example can be John Doe's driving of the car by pressing the gas pedal and moving the steering wheel to direct it. The behavior the car exhibits - the trajectories of its motion - are dependent upon the input which it receives, making it a classic example of a control system. As alluded to by the name, control systems generally have an additional component - the target behavior - which can be a single state, a single trajectory, or some more complicated state sequences. In the case of John Doe, the task is to bring him from the starting state to work.

In general, Control Theory can be quite complicated and requires a high level of mathematical rigor, and the goal of developing a specific control can be quite daunting. However, the fundamental notions which govern the nature of control systems are very intuitive, which lead to us performing control problems in our brains without any difficulty. Trajectory planning, i.e. how to move oneself from one place to another, is something which every toddler begins to become proficient in from an early age, as is the task of walking without falling over (something which expert roboticists and control theorists continue to struggle with regarding their robots). Baseball batters perform incredibly precise control problems, moving dozens upon dozens of muscles with precision to time the arrival of their bat to align with the location of the ball to the nearest few millimeters and milliseconds. Everything that relies on balance from riding a bike to walking a tightrope needs a delicate touch of control, lest one become overcome by gravity. Any activity that includes the moving of tools, from digging

a sandcastle to eating with a fork, is governed by similar requirements. The ability which elevated humans over animals - the ability to throw objects with precision - is the prototypical example of object trajectory planning. In short, there can be nothing more intuitive or fundamental to the human experience than control, the careful design of inputs to extract the desired output.

Many of the above examples have illustrated a concern beyond simply moving from one state to another, that of **constraints** and **optimality**. When John Doe drives the car, there are a number of constraints which dictate what comprises a valid control solution. First, there are constraints governed by possibility and the laws of physics - the car cannot begin floating or sink into the ground, nor can it be driven at more than 300mph. Then there are constraints of safety and legality: on each road, the speed of the car should not exceed the posted speed limit, and the car should be entirely stopped at red lights and stop signs. Furthermore, the location of the car should not be superimposed with any other vehicles or any pedestrians, nor should it be on the wrong side of the road, rather it should be on the right side of the road - or the left side, in countries such as England or Japan. In fact, it is imperative that the car always remain on the road and not cross through any building, or even cross sidewalks or lawns though it may be shorter. Constraints govern the **viability** of inputs: any control solution which violates these constraints is nonviable, and cannot be considered whatsoever. This is to be contrasted with the concern of optimality: solutions which are not optimal are still viable but are less desirable.

A solution is said to be optimal if it minimizes some **cost function**. The cost function is chosen to optimize some choice of target metric. If we want to optimize for time,

driving fastest is the best. However, if we want to optimize for distance, we can choose an unpaved dirt road, as the speed does not matter. We could alternatively elect to optimize for some combination of distance and time - weighting time more heavily would prioritize taking a longer path with a higher speed limit such as the highway, whereas weighting distance more heavily would prioritize taking a shorter path with a lower speed limit. However, what if we want to minimize the number of times that he needs to press the brakes? Perhaps we want to minimize gasoline expenditure or maximize tire longevity due to rising inflation? Each consideration or weighted combination of considerations will yield a different cost function, and therefore a different optimal control.

A fundamental question about control theory is how we pick sequences of inputs, also known as **control policies**. Returning to the human examples of control, it is unlikely that someone would decide a priori precisely the sequence of actions to take. For example, as the acrobat begins walking, he may begin leaning too much to the left. Since his current position is too far to the left, and his target position is at the center. Instead of pre-designing his walk step-by-step, he will use this feedback to course correct by leaning to the right, counteracting the current discrepancy in position. This is the idea behind **feedback control** or **closed loop control** (to distinguish from **open loop control**, which does not receive feedback), one of the most common and broadly applicable forms of control. The key defining feature of feedback control is that the state can be observed every so often, and the discrepancy between the current trajectory and the target state (or trajectory) will be used to determine the appropriate input for the next stage.

Each of these two systems of control come with its own distinct benefits and drawbacks, but they are ultimately each subject to difficulties. Under the system of open-loop control, there is no need for continual sensors and updates about the current state, as everything is neatly designed ahead of time. This comes with an issue wherein such a paradigm requires immense domain knowledge such that the designer can know a priori how the system will behave at each point to design a control policy that yields a trajectory converging to the desired goal. This is highly dependent on the initial state, as there is no reason why a sequence of inputs which drives a point on one side of the target to converge would drive a point on another side to converge to the same target. Therefore, each starting point requires an immense amount of design time and computational cost before any control policy is enacted. The issue of a priori computation is somewhat ameliorated in the case of feedback control. Starting with an educated guess for a good initial input, we can continually update this with each state observation based on the discrepancy between the current and desired states. However, while there is no longer a need for a high amount of computational time and cost up-front, this is replaced by continual computation which constantly drains resources and adds time delay with every observation.

As such, in either case there is a significant delay in real-time action, either before the first step or between steps. Equally problematic is the issue of computational cost, which derives from the non-reusability of general control policies. Feedback control makes no attempt at reusing policies across different starting states to begin with - computations are done in real-time based on the current state. As such, each initial point requires the full amount of computational cost, regardless of how many

other iterations of the feedback control process have previously been run from other initializations. Similar issues also hold for open-loop control, where in all but the simplest of systems, each initial state must receive a full design cycle for its control policy. Such problems result from a lack of robustness, wherein any error in the initial state or in the domain-knowledge-based design process will continue percolating and being amplified, as errors cannot be addressed without receiving and adjusting to feedback. As such, for either open-loop or feedback control policies to succeed, they must be designed on an initialization-by-initialization basis.

Modern research has focused on the problem of stabilizing control through a variety of different methods, but generally fail along one of two axes. Methods which provide theoretical guarantees require simultaneous controller design and external verification of correctness, and are highly prone to robustness problems by relying on explicit parametrizations of the underlying system. Data-based methods are more efficient in this sense, but result in lack of theoretical guarantees and significant latency in online deployment.

As such, we propose a method which is highly scalable, has near-instantaneous real-time execution, requires minimal computation, provides firm theoretical guarantees, and is independent of the parametrization of the control system. We will draw inspiration from the following attributes of control policies which humans affect in their day-to-day lives.

1. **Memorizing Landmark Paths:** Humans use discretized landmarks to guide their path. When receiving directions, it is common to hear “turn left at the

second stop sign” or “go straight until the clock tower”. Controls do not need to be designed individually for each point, but can generally reuse controls from a nearest landmark, or ‘seed’ point.

2. **Non-Parametric:** Humans don’t have an explicit functional expression for how to move from one point all the way to an end point. Instead, they know the direction in which they wish to go, and simply exert movement in that direction.
3. **Continually Expanding:** Humans are constantly expanding their horizons - updating their control to encompass more of the world. If someone travels to a location they have never been, they derive new controls based on their new environment to bring them closer to the goal, before re-using previously derived controls.
4. **Adaptive Precision:** Humans change the desired precision of their control decisions (and the rate at which they update them) based on their distance from the target. If someone is in New York City and wishes to get to Washington DC, the control need only be vaguely “drive South”, whereas if they are only a couple hundred meters from their destination, an ever-more precision of required as not to overshoot.

Our designed methodology incorporates these key features of human-inspired control. This algorithmic category, which we dub Non-Parametric Chain Policies (NCP), is discussed in-depth in Chapter 5 and forms one of the key contributions of this research.

NCP control relies on designing individual controls for points either offline or online

and then reusing them for all points which are sufficiently close, backed by strong theoretical guarantees. The controls used at a point in time depend on the region of the state space. Importantly, the space is subdivided into easy-to-find hypercubes which have efficient search algorithms associated with them. Sizes of containing hypercubes expand with distance from the target to reduce the number of required controls, specifically only a number that scales logarithmically with the radius. NCP Control presents significant benefits over current state-of-the-art methods:

- Stabilizing-by-design controls which provide firm theoretical guarantees of exponential stability without the need for separate verification.
- Better scaling, as the required number of control seeds grows only logarithmically with the size of the state space.
- Lifelong learning, allowing new data to expand the stabilizable region without discarding previously verified controllers.
- Improved robustness from lack of parametrization, given that parametric models can be incredibly sensitive to tuning of the parameters.
- Improved speed of execution compared to other non-parametric methods.
- Norm-agnostic design, improving interpretability and simplifying controller synthesis for highly nonlinear systems.

Chapter 3

~~iii~~ Related Work in Dynamical Systems and Control

Before we can change the world, we must understand it.

And to understand it, we must listen before we speak.

– LORD RABBI JONATHAN SACKS

3.1 Relaxations of Lyapunov Functions

The first part of the research presented in this thesis is based on forming a general relaxation of Lyapunov's Direct Method [75]. The intuition behind Lyapunov's Direct Method for asymptotic stability is quite straightforward (analogously for stability, exponential stability): one constructs a Lyapunov function $V(x)$, a continuously differentiable function that is positive definite, with a negative definite derivative. The fact that $V(x)$ is continuous and is positive definite implies that there is a small enough neighborhood of the origin such that for any $\varepsilon > 0$, we have $V_{\leq \varepsilon} \subset B_{\delta}(0)$, with δ decreasing uniformly as $\varepsilon \rightarrow 0$. Since $\dot{V}(x) < 0$, necessarily $V(\phi(t; x)) \rightarrow 0$ and thus $\|\phi(t; x)\| \rightarrow 0$. Both the requirements and the proof of Lyapunov's Direct Method were constructed in a way as to be maximally intuitive along this line of logic.

As with any intuitive method which consists of necessary-but-not-sufficient conditions for an implication, the requirements are easy to understand and connect with the result, but are not necessarily minimal. Due to this fact combined with the utility of Lyapunov theory in a broad realm of disciplines, relaxations of the conditions of Lyapunov's Direct Method has been an active field of research for well over half a

century. Following the logic of the intuition behind the method, much of this research has focused on replacing the negative definiteness of $\dot{V}(x)$ with other conditions which guarantee that $V(\phi(t; x)) \rightarrow 0$, which by the positive definiteness of $V(x)$ will guarantee that $\|\phi(t; x)\| \rightarrow 0$, i.e. that the origin is asymptotically stable. Such relaxations of Lyapunov's conditions can be broadly divided into three categories:

1. Allowing $\dot{V}(x) \geq 0$, but requiring higher-order conditions, i.e. placing some conditions on a function u such that $u(V^{(q)}, V^{(q-1)}, \dots, \ddot{V}, \dot{V}, V)$
2. Requiring $\dot{V}(x) < 0$ on most of the state space, but allowing specific regions where $\dot{V}(x) \geq 0$ subject to certain conditions
3. Providing a discretized Lyapunov condition, i.e. placing inequality constraints on the sequence $\{V(\phi(nT, x))\}_{n \in \mathbb{N}}$ instead of placing general constraints on $\dot{V}(x)$

Placing higher-order conditions on $V(x)$ has long been the most prolific of such approaches and particularly popular early on, with a large number of papers in this direction between 1962-1971. This approach was first considered in [66], where if one requires the negative definiteness of \ddot{V} instead of the negative definiteness of $\dot{V}(x)$, it can be guaranteed that all trajectories either diverge to infinity or converge to the origin. Similar results are derived and then generalized in [130], wherein for any compact, invariant subset of the state space $K \subset S$ wherein $S - K$ remains simply connected, then no solution can have limit points in $S - K$. Under such a framework, the results of [66] become a corollary by taking K to be the origin. This is expanded and simplified by considering the third derivative in addition to the previous ones in [22]. The negative definiteness of the weighted sum of the first three derivatives,

i.e. $g(x) = a_2\ddot{V}(x) + a_1\dot{V}(x) + \dot{V}(x)$ is used in place of the negative definiteness of $\dot{V}(x)$ to guarantee asymptotic stability, a stronger result than is achievable by only considering the first two derivatives.

From specific order of derivatives, research moved on to the general case of higher-order derivatives. In [46], the first m derivatives of $V(x)$ are taken for an arbitrary $m \geq 1$ and a comparison function $u(V^{(q)}, V^{(q-1)}, \dots, \ddot{V}, \dot{V}, V)$ is considered. If $V(x)$ is bounded above and below by class \mathcal{K} functions of the norm, solutions of the comparison function are bounded above, and the m -th derivative of $V(x)$ is bounded above by the comparison function, the origin can be proven to be stable. Around four decades later in [85–87] (all by the same two researchers), negative definiteness of expressions in the form of weighted sum of arbitrary number of derivatives, i.e. $h(x) = V^{(m)}(x) + a_{m-1}V^{(m-1)}(x) + \dots + a_1\dot{V}(x)$ can be used in place of negative definiteness of $\dot{V}(x)$ to guarantee asymptotic stability. There is significant discussion of what the coefficients a_i must be, but eventually over the three papers this is weakened to simply requiring that a_i all be nonnegative. In fact, in [4] it is proven that if this function $h(x)$ is negative definite, then the integral of $h(x)$ forms a traditional asymptotic Lyapunov function.

The second direction is that of limiting the regions where $\dot{V}(x)$ is nonnegative, which forms the bulk of more recent research in the field. The simplest of such cases, wherein $\dot{V}(x)$ is only negative semidefinite instead of negative definite, is considered by LaSalle's Invariance Principle (alternatively known as the Barbashin-Krasovskii Theorem) in [65, 67], and principles of construction of such functions are clearly delineated in [78, 84]. These works prove that if it can be guaranteed that no trajectories accumulate in

the zero-level set of $\dot{V}(x)$ (i.e. in the set $\{x \mid \dot{V}(x) = 0\}$) other than at the origin is used to prove asymptotic stability of the origin. Guarantees that trajectories must not accumulate anywhere beside the origin were provided by [80], wherein auxiliary functions are used to guarantee that the zero-level set must be left at all non-origin points. Results based on extending Matrosov's results by either relaxing the constraints on the auxiliary functions, nesting multiple applications of auxiliary functions, or giving strong guarantees for specific types of systems are provided in [8, 68, 71, 73, 115, 117, 124].

Incisive analysis of the case where the $\dot{V}(x)$ is negative semidefinite is given by [76, 82], who demonstrates that if $\dot{V}(x)$ is negative semidefinite but zero only periodically (i.e. is not zero for any non-empty open intervals of time) that the origin is asymptotically stable. Such methods are most clearly and concisely further refined and generalized to the time-dependent case in [132–134] by requiring that $\dot{V}(t, x) \leq \mu(t)V(t, x)$ for some function $\mu(t)$. Based on the properties of $\mu(t)$, one can ascertain asymptotic or exponential stability of the origin, being either local or global depending on further properties of $\mu(t)$. Similar results, requiring a variety of conditions of the form $\dot{V}(t, x) \leq \mu(t)g(V(t, x))$ to provide Lyapunov stability results (although primarily focusing on input-to-state stability) for a variety of specific categories of systems can be found in [9, 17, 24, 25, 49, 60, 77, 83, 91, 100, 101]. Finally, in [72], a bound on where $\dot{V}(x) \geq 0$ is permitted is given not based on the amount of time that trajectories spend in this condition, but based on the volume of the state space wherein this occurs.

The last of the popular Lyapunov relaxation categories is that of discretization, i.e.

considering discrete-time conditions on the Lyapunov function in continuous time. In practice, this consists of replacing the negative definite condition on $\dot{V}(x)$ (which can be considered a condition of the form $\lim_{t \rightarrow 0} V(\phi(t, x)) - V(x) < 0$) with conditions on $V(\phi(t, x)) - V(x)$ for fixed $t > 0$. This has the added benefit of not requiring differentiability of $V(x)$, which is particularly significant in light of the higher-order derivative conditions which in the limit require smoothness. For the purpose of control, [35] indirectly proves that a system which is asymptotically stable in the discretized sense, i.e. $V(\phi(t, x)) - V(x) < 0$, is also asymptotically stable in the continuous sense. This is expanded to uniform stability in the time-dependent case by [1], extending [89] who previously proved this only for systems wherein $\dot{V}(x)$ was already known to be negative semidefinite.

3.2 Explicit Model Predictive Control

The second part of the research presented in this thesis is based on refining current methods of building control policies throughout the state space with significant practical applications and strong theoretical guarantees. The most natural analog to this work is found in the realm of Model Predictive Control [42, 81, 90, 98]. Model Predictive Control (MPC) is a method of control wherein at a given start time t , a piecewise control policy $u(t)$ is designed with time steps of Δt for a fixed horizon $[t, t + T]$ by estimating the effect of the control on the model. The first part of $u(t)$ is applied to the system, and then after Δt time passes, another observation of the system is taken and a new control is computed for the receding horizon, state, and prediction. For linear MPCs, i.e. regarding models which are truly linear or in

approximating nonlinear models with a fixed-horizon linearization, a number of strong theoretical guarantees for computational time, cost, and ultimate system stability can be provided [20, 40, 64]. Nevertheless, as the system becomes sufficiently complex or precision is of particular importance, nonlinear MPC systems must be utilized [6, 41]. A major drawback of these nonlinear MPCs is that computation can be incredibly expensive, particularly as the system and its interaction with the input control become increasingly complicated.

A popular method which can mitigate some of the drawbacks of the 'online' MPC is that of 'offline' Explicit MPC [5, 12, 13, 21, 52, 92]. Explicit Model Predictive Control (eMPC) is a method which replaces the iterated computation of a control for a specific trajectory during execution time by the off-line computation of piecewise affine (PWA) controls for the entire state space along with their associated regions, which can be chained together during runtime by simply observing the current state and determining the correct region. Such PWA controls have even been incorporated into deep neural network architectures for increased interpretability of reinforcement learning policies [123]. For the case of linear control, each PWA control policy is associated with a convex polytopic region, which can be explicitly parametrized as a subset of the state space [63]. However, the number of controls and regions grows exponentially with attributes of the system including the number of states, which results in a drastic increase in both the memory requirements - of storing the regions and trajectories - and the computational time during the algorithm's execution for determining which polytope contains the current state.

In fact, discovering the correct polytope, known as the point location problem or

set membership problem (or even more broadly the nearest polytope problem), is an active research direction in its own right [11, 15, 118]. For some special cases, such as a linear cost function, the problem can be shown to devolve to a Voronoi diagram which can be solved by a weighted nearest-neighbor algorithm [57]. For the general case, there are a wide variety of algorithms used to find the nearest polytope. There are research angles which try to reduce the problem to that of reachability and use reference points to simplify the process [113, 116]. Others, as in [125], give an explicit lattice representation to these regions, which can reduce the computation time and memory requirements, and can be combined with a binary search tree for a unified algorithm as in [10]. The research in [131] focuses on the trade-off between online and offline computation and integrates components of both to balance the MPC computation requirements with the polytope-finding computation requirements. Understanding the difficulty of finding the polytopes, some turn to approximation methods which exchange computation time and cost for suboptimality, including both a wide variety of different approximate mathematical methods [14, 16, 23, 43, 51, 56] and neural networks [26]. The work in [29] relies on using bounding boxes to find the nearest polytope. A sampling of algorithms including the bounding box methods, binary search tree methods, and seeded reference points can be found in [129].

Chapter 4

Relaxations of Lyapunov Theory via Recurrence

*Greatness is not about never falling, but about
falling and learning how to rise again.*

– LORD RABBI JONATHAN SACKS

This chapter is adapted from the following article:

Siegelmann, R., Shen, Y., Paganini, F. & Mallada, E. *A recurrence-based direct method for stability analysis and gpu-based verification of non-monotonic lyapunov functions in 2023 62nd IEEE Conference on Decision and Control (CDC) (2023)*, 6665–6672. doi:[10.1109/CDC49753.2023.10383373](https://doi.org/10.1109/CDC49753.2023.10383373).

4.1 Preliminaries

We consider a continuous-time dynamical system

$$\dot{x} = f(x), \tag{4.1}$$

where $x \in D \subset \mathbb{R}^n$ is the state, and the map $f : D \rightarrow \mathbb{R}^n$ is a continuous function defined over an *open domain* D . Given an initial state x , we use $\phi(t, x)$ to denote the solution of (4.1). Throughout the paper, we make the following assumption about the vector field and its solutions.

Assumption 4.1. *The vector field $f(x)$ in (4.1) is locally Lipschitz. That is, for any compact set $S \subset D$, there exists a constant $\bar{L}_S \in \mathbb{R}_{\geq 0}$ such that*

$$\|f(y) - f(x)\| \leq \bar{L}_S \|y - x\|, \quad \forall x, y \in S.$$

The local Lipschitz nature of the vector field implies that solutions must exist for some amount of time, which we will denote by the following:

Definition 4.1 (Interval of Existence). *For $x \in D$, the **maximal interval of existence** $I(x) \subset \mathbb{R}$ is the largest open interval around $t = 0$ such that $\phi(t, x)$ exists for all $t \in I(x)$. The trajectory is said to be **forward complete** if $I(x) \supset [0, \infty)$.*

Whenever the initial condition is understood from the context, we will use $x(t) := \phi(t, x)$. Whenever the set S is understood from context, we will use \bar{L} instead of \bar{L}_S . For simplicity, we refer to the dynamical system (4.1) as the dynamical system f .

We next introduce the core building blocks of Lyapunov Stability Theory.

Definition 4.2 (Stability). *An equilibrium x^* is **stable** if for any $\varepsilon > 0$, $\exists \delta > 0$, such that if $\|x - x^*\| \leq \delta$ then $\|\phi(t, x) - x^*\| \leq \varepsilon \forall t \geq 0$.*

Definition 4.3 (Attractivity). *An equilibrium x^* is **attractive on the set S** if every $x \in S$, $\|\phi(t, x) - x^*\| \rightarrow 0$ as $t \rightarrow \infty$.*

Definition 4.4 (Asymptotic Stability). *An equilibrium x^* is **asymptotically stable on the set S** , if it is stable, and attractive on S .*

Definition 4.5 (Exponential Stability). *An equilibrium x^* is **exponentially stable on the set S** if there exists constants $K > 0$, $\lambda > 0$ such that if $x \in S$, then*

$$\|\phi(t, x) - x^*\| \leq K e^{-\lambda t} \|x - x^*\|, \quad \forall t \geq 0. \quad (4.2)$$

It will also be useful throughout our presentation to define sets that are of general use to characterize transient, as well as asymptotic behavior.

We start by characterizing the set of points that can be reached within a finite interval of time.

Definition 4.6 (Reachable Tube). *For the dynamical system f , a time $\tau > 0$, and a set $S \subset D$, we denote the τ -reachable tube from S within τ units of time by*

$$\mathcal{R}^\tau(S) = \bigcup_{x \in S, t \in [0, \tau] \cap I(x)} \{\phi(t, x)\}.$$

Next, we formally define positive invariant sets.

Definition 4.7 (Positively Invariant Sets). *A set $S \subseteq \mathbb{R}^n$ is positively invariant w.r.t.*

(4.1) if and only if:

$$x \in S \implies \phi(t, x) \in S, \quad \forall t \in \mathbb{R}_{\geq 0}. \quad (4.3)$$

Since we only consider here positive invariant sets, as opposed to negative invariant sets, we will often refer to them as plainly invariant sets. As mentioned before, the notion of positive invariance is a fundamental building block of Lyapunov Theory. By trapping trajectories on compact sub-level sets of a function one can guarantee boundedness of trajectories, stability, and asymptotic stability via a gradual reduction of the Lyapunov function value.

4.2 Recurrence

To relax the notion of invariance, one must allow trajectories to temporarily leave a set. However, in order to still be able make statements about asymptotic behavior, our first condition requires trajectories to return infinitely often.

Definition 4.8 (Recurrent Set). *A set $S \subseteq \mathbb{R}^n$ is recurrent w.r.t. (4.1), if for any $x \in S$, and $t \geq 0$,*

$$\exists t' > t, \quad \text{s.t.} \quad \phi(t', x) \in S. \quad (4.4)$$

Since trajectories are allowed to leave S , in our development, it will be useful to keep track of the time intervals where a trajectory $\phi(t, x)$ lies within a given set S for a given initial point $x \in D$.

Definition 4.9 (Containment Times). *Given a set $S \subset D$ and a point $x \in D$, we define the set of containment times, $T_S(x)$, as the set of times t for which the trajectory $\phi(t, x) \in S$, i.e.,*

$$T_S(x) := \{t \in \mathbb{R}_{>0} \mid \phi(t, x) \in S\}.$$

Given constants a, b , we write

$$T_S(x; a, b) := T_S(x) \cap (a, a + b].$$

For convenience we also write $T_S(x; b) := T_S(x; 0, b)$. Finally, when the set S is clear from context, we may omit the subscript entirely.

The notion of recurrent sets introduced here is related to classical Poincare recurrence [97], and in particular, Poincare recurrent sets [7, Def. 2.4.1], which constitutes the union of Poincare recurrent points; a point x is Poincare recurrent if its backward and forward flows, i.e., $\{\phi(-t, x)\}_{t \geq 0}$ and $\{\phi(t, x)\}_{t \geq 0}$, get arbitrarily close to x , *infinitely often*. In fact, one can show that any open subset S of a Poincare Recurrent Set is a Recurrent Set according to Definition 4.8.

Definition 4.8, under appropriate assumptions for S , implies (see [103]) that part of the ω -limit set of f must be contained within the recurrent set S . While this property suggests attractivity to some set that intersects S , for stability analysis we further require control on how far the trajectory may depart from S , therefore we introduce the following stronger notion of recurrence:

Definition 4.10 (τ -Recurrent Set). *A set $S \subseteq D$ is τ -recurrent w.r.t. (4.1), if for any $x \in S$ and $t \geq 0$,*

$$\exists t' > t, \quad \text{with } t' - t \in (0, \tau] \quad \text{and } \phi(t', x) \in S. \quad (4.5)$$

We further say that S is strictly τ -recurrent, if for any $x \in S$, and $t \geq 0$,

$$\exists t' > t, \quad \text{with } t' - t \in (0, \tau] \quad \text{s.t. } \phi(t', x) \in S \setminus \partial S. \quad (4.6)$$

Implicit in the above definition is that trajectories which start in S are *forward complete*, since the flow is defined for arbitrarily large times. Trajectories that start in S will visit it *infinitely often* (again and again), and *forever* (there is always a future time when it is visited again), and moreover there is a bound τ on the lengths of excursions outside S .

While definition 4.10 is sufficient for the development that follows, its conditions for far-off times are hard to verify. The next lemma shows that such verification is simpler when the set S is compact.

Lemma 4.1 (Characterization of τ -Recurrent Compact Sets). *Let $S \subset D$ be a compact set, and consider the system (4.1) under Assumption 4.1, and $\tau > 0$. The following*

conditions are equivalent:

(i) S is τ -recurrent.

(ii) For any $x \in S$, $\exists t \in (0, \tau]$ with $\phi(t, x) \in S$.

(iii) For any $x \in S$ there is a sequence $\{t_n\}_{n \in \mathbb{N}}$ satisfying,

$$\lim_{n \rightarrow \infty} t_n = \infty, \quad \text{with } t_{n+1} - t_n \in (0, \tau], \quad (4.7)$$

and $\phi(t_n, x) \in S \forall n$.

Proof. (i) \implies (ii): Follows from Definition 4.10 and choosing $t = 0$.

(ii) \implies (iii): Given $x \in S$, we build the sequence $\{t_n\}_{n \in \mathbb{N}}$ satisfying (4.7) and $\phi(t_n, x) \in S$ by induction. For the base case, let $t_0 := 0$, $x_0 = \phi(t_0, x) = x$, and define:

$$t_1 = \max\{t \in (0, \tau] \mid \phi(t, x_0) \in S\};$$

the above set of times is non-empty by hypothesis. Its supremum is actually a maximum due to the compactness of S and the continuity of $\phi(t, x)$. By construction, $t_1 \in I(x)$.

The inductive construction proceeds in a similar manner: given $t_1 < t_2 < \dots < t_n$, with

$x_n := \phi(t_n, x) \in S$, define:

$$t_{n+1} = t_n + \max\{t \in (0, \tau] \mid \phi(t, x_n) \in S\}. \quad (4.8)$$

Note that t_{n+1} exists analogously to the above, and satisfies $t_{n+1} - t_n \in (0, \tau]$ as required. Further, $t_{n+1} - t_n \in I(\phi(t_n, x))$, which implies that $t_{n+1} \in I(x)$.

It remains to show that $t_n \rightarrow \infty$, which we argue by contradiction. If, instead, the strictly increasing sequence of times was bounded, we would have $t_n \uparrow t^*$. If t^* was the supremum of the maximal interval $I(x)$, then a standard result in differential equations [30] implies that $\phi(t, x)$ should exit any compact set as $t \uparrow t^*$; this would contradict the fact that $x_n := \phi(t_n, x) \in S \forall n$, and S compact. Therefore $[0, t^*] \subset I(x)$, $x^* = \phi(t^*, x)$ exists, and by continuity of $\phi(\cdot, x)$, it follows that $\lim_{n \rightarrow \infty} x_n = x^* \in S$.

We now choose n large enough, say n^* , such that $t^* < t_{n^*} + \tau$. Since $\phi(t^*, x) \in S$, then t^* is a candidate value for the induction step (4.8) and satisfies $t^* > t_{n^*+1}$; this contradicts the fact t_{n^*+1} is the maximum such value. Thus, $t_n \rightarrow \infty$, as desired. In particular, $I(x) \supset [0, \infty)$.

(iii) \implies (i): Given $x \in S$ and $t \geq 0$, let n^* be the largest n s.t. $t_n \leq t$, then, it follows that $t_{n^*+1} - t \in (0, \tau]$, $\phi(t_{n^*+1}, x) \in S$. Definition 4.10 is satisfied with $t' = t_{n^*+1}$. \square

Note as a consequence of this Lemma, that condition (ii) over the time interval $(0, \tau]$, when valid over a compact set S implies the forward completeness of trajectories initiating in this set.

We now proceed to bound the distance a trajectory can travel away from a τ -recurrent set. To that end, we recall here that the vector field (4.1) is assumed locally Lipschitz (Assumption 4.1). While such property suffices, it will prove convenient to obtain tighter bounds via locally one-sided Lipschitz constants.

Definition 4.11 (One-sided Lipschitz). *For system 4.1 under Assumption 4.1, the one-sided Lipschitz constant over the compact set $S \subset D$ is defined as the smallest $L_S \in \mathbb{R}$ such that*

$$(y - x)^T(f(y) - f(x)) \leq L_S \|y - x\|^2, \quad \forall x, y \in S.$$

Note that $L_S \leq \bar{L}_S$.

We will also introduce a notation for the maximum norm of the vector field on a (compact) set S ,

$$F_S := \max_S \|f(x)\|.$$

When S is a ball of radius ε centered around a fixed point x^* , i.e., $S = B_\varepsilon(x^*)$, we will use $F_\varepsilon := F_{B_\varepsilon(x^*)}$ for simplicity. Note that $F_\varepsilon \rightarrow 0$ when $\varepsilon \rightarrow 0$. Using these definitions, we now derive a bound on how far trajectories can go from a compact set in which they start.

Lemma 4.2 (Containment Lemma). *Consider system (4.1) under Assumption 4.1. Let $S \subset D$ be a compact set such that solutions starting in S are forward complete. Then, $\mathcal{R}^\tau(S)$ has compact closure, and defining $L := L_{\text{cl}\mathcal{R}^\tau(S)} < \infty$, we have*

$$\max_{t \in [0, \tau]} d(\phi(t, x), S) \leq F_S h(\tau; L), \tag{4.9}$$

where

$$h(\tau; L) := \begin{cases} \frac{e^{L\tau}-1}{L}, & L \neq 0, \\ \tau, & L = 0. \end{cases} \quad (4.10)$$

Proof. Compactness of $\text{cl}\mathcal{R}^\tau(S)$ is proved in Proposition 5.1 of [70].

Let $x \in S$ and let $u(t) := \|\phi(t, x) - x\|$, with $t \in [0, \tau]$. Observe that since $x \in S$, $d(\phi(t, x), S) \leq u(t)$. Thus, bounding $u(t)$ will be sufficient. Differentiating $u(t)^2$ w.r.t to time gives

$$\begin{aligned} \frac{d}{dt}u(t)^2 &= \frac{d}{dt}\langle \phi(t, x) - x, \phi(t, x) - x \rangle \\ &= 2\langle \phi(t, x) - x, f(\phi(t, x)) \rangle \\ &= 2\langle \phi(t, x) - x, f(\phi(t, x)) - f(x) \rangle \\ &\quad + 2\langle \phi(t, x) - x, f(x) \rangle \\ &\leq 2Lu(t)^2 + 2F_S u(t), \end{aligned}$$

where the first step uses $\frac{d}{dt}\phi(t, x) = f(\phi(t, x))$, the second adds and subtracts $f(x)$, and the last inequality follows from the definition of L (note $\phi(t, x) \in \mathcal{R}^\tau(S)$), F_S and $u(t)$. It follows then that

$$2u(t)\dot{u}(t) \leq 2Lu(t)^2 + 2F_S u(t).$$

If x is an equilibrium of (4.1), then $u(t) \equiv 0$ and the bound (4.9) is trivial; otherwise, since $u(0) = 0$, we must have an interval $t \in (0, \delta)$ where $u(t) > 0$. Extend this interval

maximally (but no further than τ), defining:

$$t_1 = \max\{\delta \in (0, \tau] \mid u(t) > 0 \text{ in } (0, \delta)\}.$$

We have $t_1 > 0$, and

$$\dot{u}(t) \leq Lu(t) + F_S, \quad \forall t \in (0, t_1);$$

applying Grönwall's inequality (c.f [62], Lemma A.1) yields

$$u(t) \leq \frac{F_S}{L}(e^{Lt} - 1) \quad \text{or} \quad u(t) \leq F_S t,$$

respectively if $L \neq 0$ or $L = 0$, so

$$u(t) \leq F_S h(t; L) \leq F_S h(\tau; L) \quad \text{for } 0 < t < t_1, \quad (4.11)$$

where we used the fact that $h(\cdot, L)$ is increasing and $t_1 \leq \tau$. If $t_1 = \tau$, then $F_S h(\tau; L)$ is a bound on the $u(t)$ across the entire interval $(0, \tau]$. If, instead, $t_1 < \tau$, then necessarily $u(t_1) = 0$ and the system has a periodic orbit of period t_1 ; (4.11) still gives the bound $F_S h(\tau; L)$ over the entire period, and hence for all time. (4.9) follows. \square

The Containment Lemma, which provides containment guarantees for a finite time, can be combined with the recurrence property of Definition 4.10 and Lemma 4.1 to provide trajectory bounds for all positive times.

Corollary 4.1 (Boundedness of Trajectories). *Let S be a compact τ -recurrent set.*

Then it follows that for any $x \in S$,

$$d(\phi(t, x), S) \leq F_S h(\tau; L), \quad \forall t \geq 0,$$

where $L := L_{\text{cl}\mathcal{R}^\tau(S)}$. Moreover, the τ -reachable tube $\mathcal{R}^\tau(S)$ is invariant.

Proof. Applying Lemma 4.1, we have a sequence of times $t_n \rightarrow \infty$ such as $x_n = \phi(t_n, x) \in S$, and $t_{n+1} - t_n \leq \tau$. Also trajectories starting in S are forward complete. Applying Lemma 4.2 starting from $x_n \in S$ we conclude that:

$$d(\phi(t, x), S) \leq F_S h(\tau; L) \text{ for } t \in [t_n, t_{n+1}],$$

from which the bound holds for all time.

For the second claim, observe first that if $x \in S$, then $\phi(t, x) \in \mathcal{R}^\tau(S)$ for all $t \geq 0$. This follows by placing t in an interval $t \in [t_n, t_{n+1})$, and noting that $\phi(t, x) = \phi(t - t_n, x_n)$ with $x_n \in S$ and $t - t_n \leq \tau$.

Now, if $y \in \mathcal{R}^\tau(S)$, then $y = \phi(t', x)$ and therefore $\phi(t, y) = \phi(t + t', x) \in \mathcal{R}^\tau(S)$, as required for invariance. \square

We finalize this section, noting that Corollary 4.1 imbues compact τ -recurrent sets with the same functional property of compact invariant sets, i.e., bounding trajectories. This provides the cornerstone to the development of a recurrence-based stability theory.

4.3 Recurrent Lyapunov Functions

Having established the ability to bound trajectories using τ -recurrent sets, we now introduce the modified conditions on a function $V : D \rightarrow \mathbb{R}_{\geq 0}$, that relax the standard Lyapunov conditions for stability. In contrast to the classical counterpart, we do not require V to be monotonically non-increasing along trajectories. Rather, for any given initial $x \in D$, we allow $\tau(x)$ units of time to elapse before requiring the function to meet any requirements on its value. This leads to the proposed definition of Recurrent Lyapunov Functions.

Definition 4.12 (Recurrent Lyapunov Function (RLF)). *Given an equilibrium point $x^* \in D$ of (4.1), a set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$ and $\tau > 0$. We say that a continuous function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is a τ -**Recurrent Lyapunov Function** over the set S if the following properties hold:*

(i) V is **positive definite** around x^* , that is,

$$V(x) > 0, \forall x \neq x^*, \text{ and } V(x^*) = 0. \quad (4.12)$$

(ii) V is τ -**recurrent** over S , that is,

$$\min_{s \in T_S(x; \tau)} V(\phi(s, x)) \leq V(x), \quad \forall x \in S. \quad (4.13)$$

We make the following remarks about Definition 4.12. First, the minimum in (4.13) is taken over the non-necessarily closed set $T_S(x; \tau) = T_S \cap (0, \tau]$. For the min to be finite

as required, there must exist $t \in (0, \tau]$ with $\phi(t, x) \in S$. Second, we only require V to be continuous; while classical Lyapunov theory can be developed for non-differentiable functions, it usually requires increased complexity in the analysis. Our results can be readily stated for only continuous V . Finally, the τ -recurrent property (4.13) acts as a substitute to the standard differential inequality: $\dot{V} = \nabla V(x)^T f(x) \leq 0$. As we will see next, this condition allows us to substitute the standard invariance property with the more relaxed notion of recurrence.

Lemma 4.3. *Given any $c \geq 0$ and a compact set $S \subseteq D$. If $V : D \rightarrow \mathbb{R}_{\geq 0}$ is continuous and τ -recurrent over S (c.f. (ii) in Definition 4.12), then, the following holds:*

(i) *The set S is τ -recurrent.*

(ii) *The set $V_{\leq c} \cap S = \{x \in S \mid V(x) \leq c\}$ is τ -recurrent.*

Proof. We start by noting that since V is continuous, it has a finite maximum over compact S ; hence, there exists c large enough such that $S \cap V_{\leq c} = S$. As a result, property (i) follows directly from (ii).

To prove (ii) we use the characterization (ii) of Lemma 4.1. By hypothesis, for any $x \in S \cap V_{\leq c}$, one can find $t' \in (0, \tau]$ such that $\phi(t', x) \in S \cap V_{\leq c}$. Since $S \cap V_{\leq c}$ is compact, by Lemma 4.1 it is τ -recurrent. \square

We are now ready to present the main result of this section, which states that the existence of an RLF is sufficient to guarantee the stability of the associated equilibrium point.

Theorem 4.1 (Stability). *Consider system (4.1) under Assumption 4.1, with an equilibrium point $x^* \in D$, and a compact set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$. Then, if $V : D \rightarrow \mathbb{R}_{\geq 0}$ is a τ -RLF over S , the equilibrium x^* is stable.*

Proof. It suffices to show that for any $\varepsilon > 0$, there exists set $I \subset B_\varepsilon(x^*)$, I invariant under (4.1) and with $x^* \in \text{int}(I)$. Without loss of generality, assume $B_\varepsilon(x^*) \subset S$.

Note first that from Lemma 4.3, S is τ -recurrent, so trajectories starting in S are forward complete; furthermore, by Lemma 4.2 there is a finite $L = L_{\text{cl}(\mathcal{R}^\tau(S))}$.

Find $\varepsilon' > 0$ small enough such that

$$\varepsilon' + F_{\varepsilon'} h(\tau; L) \leq \varepsilon. \quad (4.14)$$

Now let $\alpha = \min_{\varepsilon' \leq \|x - x^*\| \leq \varepsilon} V(x)$; by construction, $\alpha > 0$. Select β such that $0 < \beta < \alpha$ and introduce the compact set

$$\Omega_\beta := \{x \in B_{\varepsilon'}(x^*) : V(x) \leq \beta\}. \quad (4.15)$$

Claim 1: Let $I := \mathcal{R}^\tau(\Omega_\beta)$. Then $x^* \in \text{int}(I)$, $I \subset B_\varepsilon(x^*)$.

Given (4.12), $x^* \in \text{int}(\Omega_\beta)$; also $\Omega_\beta \subset I$, so $x^* \in \text{int}(I)$. To establish $I \subset B_\varepsilon(x^*)$, apply the Containment Lemma 4.2 to Ω_β : if $x \in \Omega_\beta$, then for every $t \in (0, \tau]$ we have:

$$d(\phi(x, t), \Omega_\beta) \leq F_{\varepsilon'} h(\tau, L). \quad (4.16)$$

Since $\Omega_\beta \subset B_{\varepsilon'}(x^*)$, the triangle inequality gives the norm bound $\|\phi(x, t) - x^*\| \leq \varepsilon' + F_{\varepsilon'} h(\tau; L) \leq \varepsilon$, as claimed.

Claim 2: Ω_β is τ -recurrent.

Given $x \in \Omega_\beta$, by hypothesis there exists $t' \in (0, \tau]$ such that $x' = \phi(t', x) \in S$ and $V(x') \leq \beta$. Since $x' \in \mathcal{R}^\tau(\Omega_\beta)$ we have $\|x' - x^*\| \leq \varepsilon$ by Claim 1. However, $V(x') < \alpha$ so $\|x' - x^*\|$ cannot be in the interval $[\varepsilon', \varepsilon]$; so $x' \in B_{\varepsilon'}(x^*) \cap V_{\leq \beta} = \Omega_\beta$, and this set is τ -recurrent.

Now apply Corollary 4.1 to conclude that $I = \mathcal{R}^\tau(\Omega_\beta)$ is an *invariant* set. So I satisfies the requirements set up at the beginning of the proof. \square

4.4 Asymptotic Stability

Now that we have proven stability with RLFs, we wish to expand the theory to incorporate asymptotic stability. Similarly to Lyapunov's Direct Method, the extension essentially consists of strengthening the condition from a non-strict inequality to a strict one.

Definition 4.13 (Strict τ -Recurrent Lyapunov Function (S-RLF)). *Given an equilibrium $x^* \in D$ and a set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$. We say that a continuous function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is a **Strict Recurrent Lyapunov Function** over the set S , if the following holds:*

- (i) V is a τ -**RLF**, i.e., Definition 4.12.

(ii) V is **strictly τ -recurrent** over S , that is,

$$\min_{s \in T_{\text{int}(S)}(x; \tau)} V(\phi(s, x)) < V(x), \quad \forall x \in S \setminus \{x^*\}. \quad (4.17)$$

Since a Strict RLF is also an RLF, all the properties from our previous section hold for S-RLFs. In particular, for compact S we have: sets $V_{\leq c} \cap S$, are τ -recurrent, solutions starting in S are forward complete, and Theorem 4.1 on equilibrium stability.

The additional requirement of strict inequality in (4.17), plays the role of the typical Lyapunov condition $\dot{V}(x) < 0$, $\forall x \in S \setminus \{x^*\}$ which implies the strict decrease of $V(\phi(t, x))$ in time. In contrast, our condition amounts to decrements every so often, as the next Lemma shows.

Lemma 4.4. *Let $x^* \in D$ be an equilibrium and $S \subseteq D$ a compact set satisfying $x^* \in \text{int}(S)$. Let $V : D \rightarrow \mathbb{R}_{\geq 0}$ be a Strict τ -RLF over S . Then: given $x \in S \setminus \{x^*\}$, there exists a sequence of times $\{t_n\}_{n \in \mathbb{N}}$, such that*

$$\lim_{n \rightarrow \infty} t_n = \infty \quad t_{n+1} - t_n \in (0, \tau], \quad (4.18)$$

with $\phi(t_n, x) \in S \setminus \{x^*\}$, and for all $n \geq 1$,

$$V(\phi(t_{n+1}, x)) < V(\phi(t_n, x)) < V(x). \quad (4.19)$$

Proof. Consider $x \in S \setminus \{x^*\}$. Note first that $\phi(t, x)$ is defined for all $t \geq 0$ (forward completeness), and $\phi(t, x) \neq x^*$ by uniqueness of solutions, since x^* is an equilibrium.

The sequence $\{t_n\}$ is constructed inductively, similar to Lemma 4.1. For the base case, let $t_0 = 0$ and choose

$$t_1 = \max\{\arg \min_{t \in T_{\text{int}(S)}(x; \tau)} V(\phi(t, x))\}. \quad (4.20)$$

Note that by hypothesis (4.17) there are times $t \in (0, \tau] \cap T_{\text{int}(S)}(x; \tau)$ where $V(\phi(t, x)) < V(x)$, so the above minimum is well defined and smaller than $V(x)$; if there are multiple minimizing times, t_1 is defined as the largest. By construction $V(\phi(t_1, x)) < V(x)$.

In a similar manner: given $t_1 < t_2 < \dots < t_n$, with $x_n := \phi(t_n, x) \in S$, define

$$t_{n+1} - t_n = \max\{\arg \min_{s \in T_{\text{int}(S)}(x_n; \tau)} V(\phi(s, x_n))\}. \quad (4.21)$$

Note that $t_{n+1} - t_n \in (0, \tau]$ as required. Also,

$$x_{n+1} := \phi(t_{n+1}, x) = \phi(t_{n+1} - t_n, x_n) \in S \setminus \{x^*\}, \quad (4.22)$$

and satisfies $V(x_{n+1}) < V(x_n)$ by (4.17), so we verify (4.19).

It remains to show that $t_n \rightarrow \infty$, which we argue again by contradiction. Consider the alternative $t_n \uparrow t^* < \infty$. By continuity of $\phi(\cdot, x)$ and compactness of S , $x_n = \phi(t_n, x) \rightarrow \phi(t^*, x) \in S$; and since $\phi(t, x) \neq x^*$ for all $t \geq 0$, $\phi(t^*, x) \in S \setminus \{x^*\}$. By continuity of V we have:

$$v_n := V(\phi(t_n, x)) \xrightarrow{n \rightarrow \infty} V(\phi(t^*, x)) =: v^*.$$

Since $\{v_n\}$ is strictly decreasing we conclude that $v^* < v_n$ for all $n \in \mathbb{N}$. Now pick n such that $t_n \geq t^* - \tau$. Since $\phi(t^*, x) \in S$, it follows that $s^* := t^* - t_n \in (0, \tau]$ is in the feasible set for the minimization in (4.21), which by definition gives as minimum v_{n+1} , achieved at $t_{n+1} - t_n$. Now, since $v^* = V(\phi(s^*, x_n)) < v_{n+1}$, this means s^* achieves a smaller solution in (4.21) than $t_{n+1} - t_n$, which contradicts with the definition of $t_{n+1} - t_n$. Thus the sequence must be divergent. \square

By utilizing Lemma 4.4 and Theorem 4.1 we are able to demonstrate asymptotic stability.

Theorem 4.2 (Asymptotic Stability). *Consider system (4.1) under Assumption 4.1, with an equilibrium point $x^* \in D$ of (4.1) and a compact set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$. Then, if $V : D \rightarrow \mathbb{R}_{\geq 0}$ is an S-RLF over S , the equilibrium x^* is asymptotically stable on the set S .*

Proof. The stability requirement is already established by Theorem 4.1 and the fact that an S-RLF is also an RLF. Thus, we are only left to show the attractivity of x^* on the set S .

Consider $x \in S \setminus \{x^*\}$. Following Lemma 4.4, construct a sequence of times $\{t_n\}$ and points $x_n = \phi(t_n, x) \in S \setminus \{x^*\}$ satisfying (4.18) and (4.19), and accordingly defined by (4.20), (4.21), and (4.22). Now, let $v_n := V(x_n)$ be the *strictly* decreasing sequence that follows from (4.19), and \bar{v} be its limit, which exists since $v_n > 0$ for all n . Then, $v_n > \bar{v} \geq 0$ for all n . Since $\{x_n\} \subset S$ compact, we may take a convergent subsequence $x_{n_k} \xrightarrow{k \rightarrow \infty} \bar{x} \in S$. By continuity, $\bar{v} = V(\bar{x})$.

Suppose $\bar{v} > 0$, so $\bar{x} \neq x^*$. Then, by the strict τ -recurrence of V (c.f. (4.17)), there exists $\bar{s} \in (0, \tau]$ satisfying $V(\phi(\bar{s}, \bar{x})) < \bar{v}$ and $\phi(\bar{s}, \bar{x}) \in \text{int}(S)$; In fact, we must have $\phi(\bar{s}, \bar{x}) \in \text{int}(S) \setminus \{x^*\}$ by uniqueness and the fact that $\bar{x} \neq x^*$. Note that by continuity

$$V(\phi(\bar{s}, x_{n_k})) \xrightarrow{k \rightarrow \infty} V(\phi(\bar{s}, \bar{x})) < \bar{v}, \quad (4.23)$$

and since $\phi(\bar{s}, \bar{x}) \in \text{int}(S)$, for k large enough we must have $\phi(\bar{s}, \bar{x}) \in \text{int}(S)$.

However, by construction of the sequence according to Lemma 4.4, we have that

$$v_{n_k+1} = \min_{s \in T_{\text{int}(S)}(x_{n_k}; \tau)} V(\phi(s, x_{n_k})) \leq V(\phi(\bar{s}, x_{n_k})). \quad (4.24)$$

Combining (4.23) and (4.24), and taking larger k if needed we have

$$v_{n_k+1} \leq V(\phi(\bar{s}, x_{n_k})) < \bar{v}$$

which leads to a contradiction since by construction of the sequence v_n we must have $v_{n_k+1} > \bar{v}$. Therefore we have shown that $v_n = V(\phi(t_n, x)) \xrightarrow{n \rightarrow \infty} 0$.

An immediate consequence of (4.12) is that $x_n = \phi(t_n, x) \xrightarrow{n \rightarrow \infty} x^*$. Indeed, $x_n \in S \forall n$, and for any $\varepsilon > 0$ s.t. $B_\varepsilon(x^*) \subset S$, the minimum of $V(x)$ in $\text{cl}(S \setminus B_\varepsilon(x^*))$ is positive, so x_n must exit such a set in a finite number of steps, satisfying $\|x_n - x^*\| < \varepsilon$ afterward.

Now, let $L = L_{\mathcal{R}\tau(S)}$, and $r_n := \|x_n - x^*\|$, and choose n large enough s.t. $B_{r_n} \subset S$, for all n onward. Applying the Containment Lemma to B_{r_n} , we have that for any

$t \in (t_n, t_{n+1}]$,

$$\|\phi(t, x_0) - x^*\| \leq r_n + F_{r_n} h(\tau; L).$$

Since the right-hand side goes to zero in n , and $t_n \rightarrow \infty$, we have that $\lim_{t \rightarrow \infty} \phi(t, x_0) = x^*$. □

Theorem 4.2 demonstrates, how by enforcing the required property (strict decrement the values of $V(\phi(t, x))$) on an infinite sequence of times with bounded difference ($t_n, t_{n+1} - t_n \in (0, \tau]$), it is possible to conclude properties of the entire trajectory. This general principle will be further exploited in the next section to prove exponential stability.

Finally, we further point out that while we require S to be compact in Theorem 4.2 the results can extend to a global setting, as follows.

Corollary 4.2 (Global Asymptotic Stability). *Let Assumption 4.1 hold. Consider an equilibrium point $x^* \in D$ of (4.1). Then, if $V : D \rightarrow \mathbb{R}_{\geq 0}$ is an S-RLF over D , and has **compact sub-level sets** $V_{\leq c} \subset D, \forall c \geq 0$, then the equilibrium x^* is globally asymptotically stable.*

Proof. Pick any $x \in D$ and let $c := V(x)$. Since V is S-RLF over D , property (4.17) implies there exists $t' \in (0, \tau]$ satisfying

$$V(\phi(t', x)) < V(x) = c.$$

In fact, the same property holds for all $x \in V_{\leq c} \subset D$. It therefore follows that V is an S-RLF over $S := V_{\leq c}$, and since $V_{\leq c}$ is compact we can apply Theorem 4.2 to claim

stability of $x^* \in V_{\leq c}$ and attractivity of $\phi(t, x)$. Finally, since x was chosen arbitrarily within D , the result follows. \square

4.5 Exponential Stability

In the previous section, we showed that strictly τ -recurrent functions sequentially constrain trajectories, enforcing the desired condition (attractivity) at discrete times. We further showed that this was sufficient to ensure the same condition along the entire trajectory (for all subsequences). We now move towards exponential stability. As before, we seek conditions on the function V that enable us to enforce exponential convergence at discrete, recurrent times.

Typically, exponential stability in classical Lyapunov analysis is verified by ensuring an exponential decrease in the Lyapunov function along trajectories. More precisely, a common integral form condition is given by:

$$V(\phi(t, x)) \leq e^{-\alpha t} V(x), \quad \forall t \geq 0, \quad (4.25)$$

for some positive constant α . Such a condition tightly couples the geometry of V to trajectories, significantly complicating its verification and limiting practical applicability. To alleviate this, we now introduce a relaxed definition—*Exponential Recurrent Lyapunov Functions (E-RLFs)*—which relaxes exponential Lyapunov functions.

Definition 4.14 (Exponential Recurrent Lyapunov Function (E-RLF)). *Given an equilibrium point $x^* \in D$ of (4.1) and a set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$. We say that a continuous function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is an **Exponential Recurrent Lyapunov***

Function over the set S if the following properties hold:

- (i) V is **positive definite and linearly contained** around x^* , that is, there exist constants $a_1, a_2 > 0$ such that

$$a_1 \|x - x^*\| \leq V(x) \leq a_2 \|x - x^*\|, \quad \forall x \in S. \quad (4.26)$$

- (ii) V is **α -exponentially τ -recurrent** over S , that is, there exist a locally bounded function $\tau : S \rightarrow \mathbb{R}_{>0}$ and a positive constant α such that

$$\min_{s \in T_S(x; \tau(x))} e^{\alpha s} V(\phi(s, x)) \leq V(x), \quad \forall x \in S. \quad (4.27)$$

Note that an α -exponentially τ -recurrent function is always strictly τ -recurrent, but not the other way around. We will use (4.27) to control (exponentially decreasing) upper and lower bounds of $V(\phi(t, x))$ as $t \rightarrow \infty$.

Theorem 4.3 (Exponential Stability). *Consider an equilibrium point $x^* \in D$ of (4.1), and a compact set $S \subseteq D$ satisfying $x^* \in \text{int}(S)$. Suppose Assumption 4.1 holds, and let $V : D \rightarrow \mathbb{R}_{\geq 0}$ be an Exponential Recurrent Lyapunov Function over the set S . Then, the equilibrium x^* is exponentially stable with rate α on the set S . In particular, for every $x \in S$ and every $t \geq 0$, it holds that*

$$\|\phi(t, x) - x^*\| \leq C e^{-\alpha t} \|x - x^*\|, \quad (4.28)$$

with $C := \frac{a_2}{a_1} e^{\alpha \tau} (1 + \bar{L}h(\tau; L))$, $\tau = \sup_{x \in S} \tau(x)$, $L := L_{R^\tau(S)}$, and $\bar{L} := \bar{L}_{R^\tau(S)}$.

Proof. Let τ , L , and \bar{L} be as defined in the theorem statement, and pick any $x \in S$. Since E-RLF \Rightarrow S-RLF \Rightarrow RLF, it follows from Lemma 4.4, that $\phi(t, x)$ is bounded and forward complete. Moreover, a similar construction to that of Lemma 4.4 also defines a sequence $\{t_n\}_{n \in \mathbb{N}}$, with $t_0 = 0$, $\lim_{n \rightarrow \infty} t_n = \infty$ and $t_{n+1} - t_n \in (0, \tau]$, $\forall n$, such that

$$e^{\alpha t_{n+1}} V(\phi(t_{n+1}, x)) \leq e^{\alpha t_n} V(\phi(t_n, x)) \leq V(x), \quad n \geq 1. \quad (4.29)$$

Using (4.26) and (4.29) we deduce that, for $n \geq 1$ we

$$\|\phi(t_n, x) - x^*\| \leq \frac{V(\phi(t_n, x))}{a_1} \leq \frac{e^{-\alpha t_n}}{a_1} V(x).$$

Let now $r_n := \frac{e^{-\alpha t_n}}{a_1} V(x)$ and $B_n := B_{r_n}(x^*) \cap S$. It follows from applying Lemma 4.2 on the compact set $B_n \subset S$ that

$$\|\phi(t, x) - x^*\| \leq r_n + F_{r_n} h(\tau; L), \quad \forall t \in (t_n, t_{n+1}].$$

Furthermore, since by Assumption 4.1, f is \bar{L} -Lipschitz on $B_n \subset S$, and $f(x^*) = 0$ we have $F_{r_n} \leq \bar{L} r_n$, leading to

$$\|\phi(t, x) - x^*\| \leq r_n(1 + \bar{L} h(\tau; L)) \leq \frac{e^{-\alpha t_n}}{a_1} (1 + \bar{L} h(\tau; L)) V(x) \quad (4.30)$$

for all $t \in (t_n, t_{n+1}]$. Note, further, that $t \leq t_{n+1} \leq t_n + \tau$, therefore $-t_n \leq \tau - t$ so $e^{-\alpha t_n} \leq e^{\alpha \tau} e^{-\alpha t}$, leading to

$$\|\phi(t, x) - x^*\| \leq e^{\alpha \tau} \frac{e^{-\alpha t}}{a_1} (1 + \bar{L} h(\tau; L)) V(x).$$

Moreover, since the last bound is independent of n , and n was chosen arbitrarily, it must hold for all $t \geq 0$. Finally, applying the upper bound $V(x) \leq a_2 \|x - x^*\|$ we establish (4.28) for any $x \in S$. \square

The above theorem demonstrates the exponential stability of an equilibrium point x^* on a compact set S . Note though, that while the constant C is independent of x , as usually required, it does depend on the compact set S through τ , L and \bar{L} . This makes the extension for global exponential stability slightly more involved. We provide here a restricted extension under additional assumptions.

Corollary 4.3 (Global Exponential Stability). *Consider the system (4.1), and suppose the following hold:*

- (i) *The domain $D \subseteq \mathbb{R}^n$ is forward invariant, and the vector field f is globally Lipschitz on D with Lipschitz constant \bar{L} .*
- (ii) *The function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is an Exponential Recurrent Lyapunov Function over D with compact sublevel sets.*
- (iii) *The recurrence time function $\tau : D \rightarrow \mathbb{R}_{> 0}$ is globally bounded: $\sup_{x \in D} \tau(x) \leq \bar{\tau} < \infty$.*

Then the equilibrium point x^ is globally exponentially stable. In particular, for all $x \in D$ and $t \geq 0$,*

$$\|\phi(t, x) - x^*\| \leq C e^{-\alpha t} \|x - x^*\|, \quad (4.31)$$

where the constant $C := \frac{a_2}{a_1} e^{(\bar{L} + \alpha)\bar{\tau}} > 0$.

Proof. Let $x \in D$ be arbitrary, and define the compact sublevel set $S := V_{\leq V(x)} \subset D$. Since V is an Exponential Recurrent Lyapunov Function over D , it satisfies the ERLF conditions over its sublevel set S as well.

By Theorem 4.3, the exponential stability bound holds on S with:

$$\|\phi(t, x) - x^*\| \leq C_S e^{-\alpha t} \|x - x^*\|,$$

where the constants $C_S := \frac{a_2}{a_1} e^{\alpha \tau_S} (1 + \bar{L}_S h(\tau_S; L_S))$ and $\tau_S := \sup_{x \in S} \tau(x) \leq \sup_{x \in D} \tau(x) \leq \bar{\tau}$.

By assumption, since f is globally Lipschitz on D with constant \bar{L} , both the standard and one-sided Lipschitz constants over $R^{\tau_S}(S) \subseteq D$ satisfy $L_S \leq \bar{L}$ and $\bar{L}_S \leq \bar{L}$.

Moreover, the function $h(\tau, L) := \frac{e^{L\tau} - 1}{L}$ is strictly increasing in both τ and L for $\tau > 0$.

Therefore, we conclude that

$$C_S \leq \frac{a_2}{a_1} e^{\alpha \bar{\tau}} (1 + \bar{L} h(\bar{\tau}; \bar{L})) = C.$$

Since this bound is independent of the particular choice of $x \in D$, we obtain:

$$\|\phi(t, x) - x^*\| \leq C e^{-\alpha t} \|x - x^*\|, \quad \forall x \in D, \forall t \geq 0.$$

□

We conclude this section by relaxing the requirement in Theorem 4.3 that $x^* \in \text{int}(S)$,

which leads to an ultimate boundedness condition that we now formally define.

Definition 4.15 (Exponential Ultimate Boundedness). *The solutions of system (4.1) are said to be exponentially ultimately bounded over a set $S \subset D$ with rate $\alpha > 0$ and bound $\varepsilon' > 0$ if there exists $C > 0$ such that for every $x \in S$ and all $t \geq 0$,*

$$\|\phi(t, x) - x^*\| \leq Ce^{-\alpha t}\|x - x^*\| + \varepsilon'. \quad (4.32)$$

To show this new result we are required to define a slight variation for the exponential τ -recurrent condition in (4.27).

Definition 4.16 (α -Exponential τ -Recurrence over S relative to S'). *Let $S, S' \subseteq D$, with $S \subseteq S'$. We say that a function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is α -exponentially τ -recurrent over S relative to S' if there exists a locally bounded function $\tau : D \rightarrow \mathbb{R}_{> 0}$ and a constant $\alpha > 0$ such that:*

$$\min_{s \in T_{S'}(x; \tau(x))} e^{\alpha s} V(\phi(s, x)) \leq V(x), \quad \forall x \in S.$$

Definition 4.17 (Relative Exponential Recurrent Lyapunov Function). *Let $S \subseteq S' \subseteq D$. We say that a continuous function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is an Exponential Recurrent Lyapunov Function over S relative to S' if:*

(i) *V is linearly contained, i.e., (4.26), for all $x \in S'$,*

(ii) *V is α -exponentially τ -recurrent over S relative to S' .*

Theorem 4.4 (Ultimate Boundedness). *Consider an equilibrium point $x^* \in D$ of (4.1). Suppose Assumption 4.1 holds, and let $V : D \rightarrow \mathbb{R}_{\geq 0}$ be an E-RLF over a*

compact set $S \subseteq D$ relative to $S \cup B_\varepsilon(x^*)$, for some $\varepsilon > 0$ such that

$$x^* \in \text{int}(S \cup B_\varepsilon(x^*)), \quad \text{and} \quad \partial B_\varepsilon(x^*) \subseteq S.$$

Then, the solutions of (4.1) are exponentially ultimately bounded over $S \cup B_\varepsilon(x^*)$, i.e., (4.32), with rate α and bound

$$\varepsilon' := \varepsilon + F_\varepsilon h(\bar{\tau}; L), \tag{4.33}$$

where $\bar{\tau} := \frac{1}{\alpha} \log(C)$, $C := \frac{a_2}{a_1} e^{\alpha\tau} (1 + \bar{L}h(\tau; L))$, $\tau = \sup_{x \in S} \tau(x)$, $L := L_{R^\tau(S)}$, and $\bar{L} := \bar{L}_{R^\tau(S)}$.

Proof. Let $x \in S \cup B_\varepsilon(x^*)$. Assume first $x \in S$. Since V is an E-RLF over S relative to $S \cup B_\varepsilon(x^*)$, it follows from Theorem 4.3 that the trajectory satisfies

$$\|\phi(t, x) - x^*\| \leq C e^{-\alpha t} \|x - x^*\|, \quad \forall t \in [0, t'] \tag{4.34}$$

where t' is the first time instance when $\phi(t', x) \in B_\varepsilon(x^*)$.

Now, given any $x' \in B_\varepsilon(x^*)$ since $\partial B_\varepsilon(x^*) \subset S$, if $\phi(t, x')$ leaves $B_\varepsilon(x^*)$, it must come back to it. In fact, using again (4.34), at most, it will take $\bar{\tau} = \frac{1}{\alpha} \log(C)$ to get back to $B_\varepsilon(x^*)$. It follows then from Lemma 4.1 that $B_\varepsilon(x^*)$ is $\bar{\tau}$ -recurrent. By Lemma 4.2, any excursion from $B_\varepsilon(x^*)$ remains within distance ε' of x^* , where ε' is given in (4.33).

Thus, any trajectory starting from S either decays exponentially while in S , or remains

within the uniform bound ε' once it enters $B_\varepsilon(x^*)$. Hence, for all $t \geq 0$,

$$\|\phi(t, x) - x^*\| \leq Ce^{-\alpha t}\|x - x^*\| + \varepsilon',$$

establishing exponential ultimate boundedness over $S \cup B_\varepsilon(x^*)$. \square

4.6 Norm-Based Converse Theorems

We are now ready to show that the recurrence conditions developed in the previous sections are naturally satisfied by standard norms under classical stability assumptions. In particular, we demonstrate that if a system is asymptotically or exponentially stable, then any norm satisfies the corresponding recurrence inequality over compact subsets of the domain. We begin by showing that any norm satisfies the τ -recurrence condition over compact subsets of the domain of attraction when the system is asymptotically stable.

Theorem 4.5 (Asymptotic Stability Implies Norm is S-RLF). *Given system (4.1). Let $x^* \in D$ be an asymptotically stable equilibrium on a compact set $S \subseteq D$ (Definition 4.4) satisfying $x^* \in \text{int}(S)$. Let $\|\cdot\|$ be any norm on \mathbb{R}^n . Then, the function $V(x) := \|x - x^*\|$ is a Strict Recurrent Lyapunov Function (S-RLF) over S .*

Proof. The proof follows closely the results in Remark 2.4 of [59]. Let $V(x) := \|x - x^*\|$. Since x^* is asymptotically stable on S , standard converse Lyapunov arguments (e.g., Lemma 4.5 in [62]) guarantee the existence of a class \mathcal{KL} function β such that for all $x \in S$:

$$V(\phi(t, x)) \leq \beta(V(x), t), \quad \forall t \geq 0.$$

Now pick an arbitrary constant $\mu \in (0, 1)$ and $x \in S$. Since $x^* \in \text{int}(S)$ and is asymptotically stable on S , there exist $\delta > 0$ satisfying $B_\delta(x^*) \subseteq S$, and a finite time $t_0(x) > 0$, such that

$$\phi(t, x) \in B_\delta(x^*), \quad \forall t \geq t_0(x).$$

Now, for each $x \in S \setminus \{x^*\}$, define explicitly:

$$\tau(x) := \min\{t \geq t_0(x) : \beta(V(x), t) \leq \mu V(x)\}.$$

Since $t_0(x) > 0$, the function $\beta(s, t)$ is continuous, strictly decreasing in t , and satisfies $\lim_{t \rightarrow \infty} \beta(s, t) = 0$, it follows that this minimum exists, is unique, and strictly positive. Furthermore, continuity of β combined with compactness of S ensures $\tau(x)$ is locally bounded on S .

Clearly, by construction, we have for all $x \in S \setminus \{x^*\}$:

$$V(\phi(\tau(x), x)) \leq \beta(V(x), \tau(x)) \leq \mu V(x) < V(x),$$

thus satisfying the strictly decreasing recurrence condition (4.17). Therefore, the function $V(x) = \|x - x^*\|$ is strictly τ -recurrent over the compact set S . Finally, since by definition a norm is always proper, it follows from Definition 4.13 that $V(x) = \|x - x^*\|$ is an S-RLF. \square

Having established that asymptotic stability implies strict τ -recurrence of standard norms, we now turn our attention to exponential stability. As mention, exponential

stability will lead to norms satisfying an E-RLF condition, providing one gives some slack on the exponential rate that can be verified.

Theorem 4.6 (Exponential Stability Implies Norm is E-RLF). *Consider system (4.1), and let $x^* \in D$ be a λ -exponentially stable equilibrium on a compact set S , i.e., Definition 4.5, satisfying $x^* \in \text{int}(S)$. Then for any $0 < \alpha < \lambda$, the function $V(x) := \|x - x^*\|$ is an Exponential Recurrent Lyapunov Function (E-RLF) on S for any l.b. function $\tau(x)$ satisfying*

$$\tau(x) \geq \tau := \frac{1}{\lambda - \alpha} \ln \left(K \frac{a_2}{a_1} \right), \quad \forall x \in S, \quad (4.35)$$

where K, λ are given in (4.2) and a_1, a_2 are positive constants satisfying: $B_{a_1}(x^*) \subseteq S \subseteq B_{a_2}(x^*)$.

Proof. Let $V(x) := \|x - x^*\|$, where $\|\cdot\|$ is the norm of Definition 4.5. By hypothesis, the equilibrium x^* is λ -exponentially stable over the compact set S , implying that for all $x \in S$ and all $t \geq 0$, $\|\phi(t, x) - x^*\| \leq K e^{-\lambda t} \|x - x^*\|$.

Therefore it follows that for all $t \geq \tau$

$$\|\phi(t, x) - x^*\| \leq K e^{-\lambda t} \|x - x^*\| = K e^{-\lambda \frac{1}{\lambda - \alpha} \ln \left(K \frac{a_2}{a_1} \right)} \|x - x^*\| \quad (4.36)$$

$$\leq K e^{-\lambda \frac{1}{\lambda} \ln \left(K \frac{a_2}{a_1} \right)} \|x - x^*\| \leq \frac{a_1}{a_2} \|x - x^*\| \leq a_1 \quad (4.37)$$

implying that $\phi(\tau, x) \in B_{a_1}(x^*) \subset S$.

Analogously, for all $x \in S$ and τ as in (4.35) we have:

$$e^{\alpha\tau} \|\phi(\tau, x) - x^*\| \leq K e^{(\alpha-\lambda)\tau} \|x - x^*\| \leq \frac{a_1}{a_2} \|x - x^*\| \leq \|x - x^*\| \quad (4.38)$$

It therefore follows that for any locally bounded function $\tau(x)$ satisfying (4.35) we must have

$$\min_{t \in T_S(x; \tau(x))} e^{\alpha t} \|\phi(t, x) - x^*\| \leq e^{\alpha\tau} \|\phi(\tau, x) - x^*\| \leq \|x - x^*\|,$$

which means that $V(x) = \|x - x^*\|$ is an E-RLF over S , as desired. \square

4.7 Verification of Exponential RLFs

So far, we have defined E-RLFs and provided theoretical guarantees of exponential stability, under the assumption that the exponential τ -recurrence condition (4.27) is satisfied. Furthermore, we established that standard norms themselves are E-RLFs. We now leverage these insights to provide a practical mechanism to verify condition (4.27) explicitly using trajectory data. Our development in this section is closely related to the work on topological entropy of dynamical systems as well as its extensions to control.

In particular, in order to certify a specific behavior (α -exponential τ -recurrence) over a set S , we will first focus on using a trajectory $\phi(t, x)$ of fixed duration $[0, \tau]$ to certify such behavior over a neighborhood $B_\varepsilon(x)$ (Section 4.7.1). We will then extend our method to verify said behavior on a set S . For reasons that will become clear later,

our focus will be on sets S satisfying the E-RLF conditions for exponential ultimate boundedness of trajectories around an equilibrium x^* .

4.7.1 Trajectory-based Verification of a Neighborhood

We start by deriving conditions to verify condition (4.27) around a neighborhood of a trajectory.

Theorem 4.7 (Trajectory-based Verification of E-RLF). *Consider the system (4.1), an equilibrium point $x^* \in D$, a compact set $S' \subseteq D$, and constants $\varepsilon > 0$, $\alpha > 0$, and $\tau > 0$. Assume that $L = L_{\mathcal{R}^\tau(S')}$ and that for some $x \in S'$, with $B_\varepsilon(x) \subseteq S'$, there exists $t \in (0, \tau]$ satisfying simultaneously:*

$$e^{\alpha t} (\|\phi(t, x) - x^*\| + \varepsilon e^{Lt}) \leq \|x - x^*\| - \varepsilon, \quad (4.39a)$$

$$\text{sd}(\phi(t, x), S') + \varepsilon e^{Lt} \leq 0, \quad (4.39b)$$

where $\text{sd}(\cdot, S')$ is the signed distance w.r.t S' based on the same norm $\|\cdot\|$. Then, the function $V(x) = \|x - x^*\|$ is an E-RLF over $B_\varepsilon(x)$ relative to S' .

Proof. Let $x \in S'$ be such that $B_\varepsilon(x) \subseteq S'$, and define $V(\cdot) := \|\cdot - x^*\|$. Let $\tau > 0$ and $L := L_{\mathcal{R}^\tau(S')}$, as in the statement.

Since system (4.1) is L -one-sided-Lipschitz on $\mathcal{R}^\tau(S')$, for all $y \in B_\varepsilon(x)$ and all $t \in [0, \tau]$, we have:

$$\|\phi(t, x) - \phi(t, y)\| \leq e^{Lt} \|x - y\| \leq \varepsilon e^{Lt} \quad (4.40)$$

$$\implies \|\phi(t, y) - x^*\| \leq \|\phi(t, x) - x^*\| + \varepsilon e^{Lt}. \quad (4.41)$$

Let $t^* \in (0, \tau]$ be a time that satisfies the verification condition (4.39a)–(4.39b) at point x . In particular,

$$e^{\alpha t^*} \left(\|\phi(t^*, x) - x^*\| + \varepsilon e^{L t^*} \right) \leq \|x - x^*\| - \varepsilon. \quad (4.42)$$

Now, fix any $y \in B_\varepsilon(x)$. Then:

$$\begin{aligned} e^{\alpha t^*} \|\phi(t^*, y) - x^*\| &\leq e^{\alpha t^*} \left(\|\phi(t^*, x) - x^*\| + \varepsilon e^{L t^*} \right) && \text{(by (4.41))} \\ &\leq \|x - x^*\| - \varepsilon \leq \|y - x^*\|, && \text{(by (4.42))} \end{aligned}$$

where the last inequality holds since $y \in B_\varepsilon(x)$. This shows that $\|\cdot - x^*\|$ would satisfy the exponential decrease condition (4.27) at every point $y \in B_\varepsilon(x)$ provided that $\phi(t^*, y) \in S'$.

To complete the argument, assume that $\phi(t^*, y) \notin S'$ for some $y \in B_\varepsilon(x)$. By the assumptions of the theorem, $\text{sd}(\phi(t^*, x), S') \leq 0$ which implies $\phi(t^*, x) \in S'$. Now let $y_p \in \partial S'$ be such that $y_p = \lambda \phi(t^*, x) + (1 - \lambda) \phi(t^*, y)$ for $\lambda \in [0, 1]$. It follows then by Grönwall's inequality (c.f [62], Lemma A.1), that

$$\varepsilon e^{L t} \geq \|\phi(t^*, y) - \phi(t^*, x)\| = \|\phi(t^*, y) - y_p\| + \|\phi(t^*, x) - y_p\| \quad (4.43)$$

$$\geq \inf_{z \in \partial S'} \|\phi(t^*, y) - z\| + \inf_{z \in \partial S'} \|\phi(t^*, x) - z\| \quad (4.44)$$

$$= \text{sd}(\phi(t^*, y), S') - \text{sd}(\phi(t^*, x), S') \quad (4.45)$$

Therefore we conclude that

$$\text{sd}(\phi(t^*, y), S') \leq \text{sd}(\phi(t^*, x), S') + \varepsilon e^{Lt^*} \leq 0, \quad (4.46)$$

contradicting the assumption that $\phi(t^*, y) \notin S'$, since by compactness of S' implies $\text{sd}(\phi(t^*, y), S') > 0$.

Thus, $\phi(t^*, y) \in S'$ for all $y \in B_\varepsilon(x)$, and therefore $V(y) := \|y - x^*\|$ is an Exponential Recurrent Lyapunov Function (E-RLF) over $B_\varepsilon(x)$ relative to S' . \square

Theorem 4.7 provides a mechanism to verify that every point $y \in B_\varepsilon(x)$ will satisfy condition (4.27). We aim to leverage this conditions to verify the Relative E-RLF property of a norm $\|\cdot\|$ over S relative to S' , $S \subseteq S'$. It is therefore natural to consider an ε -cover¹ of such set. However, this approach poses two main challenges. First, condition (4.39a) cannot be satisfied for any point within $B_\varepsilon(x^*)$, which requires ε to be small as one gets closer to x^* and is aligned with the well known fact that exponential stability cannot be verified using a finite number of trajectories [34, Example 2.1]. Second, any ε -cover of S will necessarily require $O((R/\varepsilon)^d)$ number of trajectories, making its verification a computationally difficult problem.

4.7.2 Trajectory-based Verification of a Region

In this section we overcome the above mentioned problems by focusing on the verification of sets of the form $S = \text{cl}(B_R(x^*) \setminus B_\varepsilon(x^*))$ and $S' = B_R(x^*)$, which satisfy the conditions of Theorem 4.4 for exponential ultimate boundedness. This allows

¹An ε -cover of a set S is a collection of balls $B_\varepsilon(x_i)$ s.t. $S \subseteq \cup_i B_\varepsilon(x_i)$

us to limit the radius ε of the balls needed to cover S . Moreover, we will allow the radius of such cover to gradually increase as the center of the balls move away from x^* . The result is a strategy that allows to verify the region $S = \text{cl}(B_R(x^*) \setminus B_\varepsilon(x^*))$ using significantly less number of trajectories.

Theorem 4.8 (Sample Complexity for Verifying E-RLF). *Consider system (4.1), an equilibrium point $x^* \in D \subseteq \mathbb{R}^d$, and constants $R > 0$, $\varepsilon \in (0, R)$, and $\tau > 0$. Let $L := L_{\mathcal{R}^\tau(B_R(x^*))}$ denote the one-sided Lipschitz constant of (4.1) over the τ -reachable set from $B_R(x^*)$. Suppose that x^* is λ -exponentially stable over $B_R(x^*)$ with constant $K \geq 1$, choose any rate $\alpha \in (0, \lambda)$, and define $\beta := \frac{\lambda - \alpha}{\lambda + L}$, $K_\beta := K^{\frac{1-\beta}{\beta}}$, and assume $\beta \in (0, 1)$.*

Then, the function $V(x) := \|x - x^\|$ can be verified to be an Exponential Recurrent Lyapunov Function (E-RLF) over the set $S = B_R(x^*) \setminus B_\varepsilon(x^*)$, relative to $S' = B_R(x^*)$, using at most*

$$N(\varepsilon) = d \ln \left(\frac{R}{\varepsilon} \right) \frac{e}{\beta} \left(2(2 + K_\beta) \right)^d = O \left(\log \left(\frac{R}{\varepsilon} \right) \right). \quad (4.47)$$

trajectory evaluations of duration

$$\tau = \frac{1}{\lambda - \alpha} \ln \left(K(2 + K_\beta) e^{\beta/d} \right).$$

Proof. We construct a layered covering of the region $B_R(x^*) \setminus B_\varepsilon(x^*)$ using concentric annuli A_i centered at x^* , i.e., $A_i := \{x \in B_R(x^*) : R_i \geq \|x - x^*\| \geq R_{i+1}\}$, with $i \in \{0, \dots, N-1\}$, with $R_0 = R$, $R_i = \rho R_{i-1} = \rho^i R$, and $\rho, \mu \in (0, 1)$. We aim to cover each layer with verification balls of radius also given by $\varepsilon_i := \mu R_i$. The key idea

is to progressively chose different radii ε_i so as to limit the total number of balls, while still guaranteeing that $V(x) := \|x - x^*\|$ satisfies the exponential recurrence condition.

We aim to find values of ρ and μ such that for any given $i \in \{0, \dots, n\}$, the function V is an E-RLF over $B_{\varepsilon_i}(x)$, for all $x \in A_i$, relative to $B_R(x^*)$. By Theorem 4.7, it is sufficient to show that:

$$\min_{t \in T_{B_R(x^*)}(x; \tau)} e^{\alpha t} (\|\phi(t, x) - x^*\| + \varepsilon_i e^{Lt}) \leq \|x - x^*\| - \varepsilon_i. \quad (4.48)$$

Since x^* is λ -exponentially stable over $B_R(x^*)$, the left hand side of (4.48) can be upper-bounded for all $x \in A_i$ by:

$$Ke^{(\alpha-\lambda)\tau} \|x - x^*\| + \mu R_i e^{(\alpha+L)\tau} \leq R_i (Ke^{(\alpha-\lambda)\tau} + \mu e^{(\alpha+L)\tau}).$$

Now, using the fact that $\|x - x^*\| \geq R_{i+1} = \rho R_i$ and $\varepsilon_i = \mu R_i$, it follows that inequality (4.48) is satisfied whenever:

$$f(\tau, \mu) := Ke^{(\alpha-\lambda)\tau} + \mu e^{(\alpha+L)\tau} \leq \rho - \mu. \quad (4.49)$$

Choosing τ large enough, i.e.,

$$\tau(\mu) = \frac{1}{\lambda - \alpha} \ln \left(\frac{K}{\mu^\beta} \right) > \frac{1}{\lambda - \alpha} \ln(K), \quad (4.50)$$

ensures the existence of some μ small enough s.t. (4.49) hold. In particular, since

$$f(\tau(\mu), \mu) = \mu^\beta + \mu \left(\frac{K}{\mu^\beta} \right)^{\frac{1-\beta}{\beta}} = (1 + K_\beta) \mu^\beta$$

and $\beta \in (0, 1)$, it is sufficient

$$\mu(\rho) := \frac{\rho^{1/\beta}}{\sqrt[\beta]{2 + K_\beta}} \quad (4.51)$$

for (4.49) to hold, thus ensuring that (4.48) is satisfied for all i .

Next, we compute the number of annuli required to cover the region $B_R(x^*) \setminus B_\varepsilon(x^*)$.

Since the radii follow a geometric progression $R_i = \rho^i R$, we stop when $R_n \leq \varepsilon$, which implies

$$\rho^n R \leq \varepsilon \quad \Rightarrow \quad n = \ln \left(\frac{R}{\varepsilon} \right) \frac{1}{\ln(\rho^{-1})}. \quad (4.52)$$

To estimate the number of balls of radius $\varepsilon_i = \mu R_i$ required to cover the annular region A_i , we leverage properties of covering and packing numbers of a set K , i.e., $\mathcal{N}(K, \varepsilon)$ and $\mathcal{P}(K, \varepsilon)$, respectively. In particular, for a given set, it is easy to show [122]:

$$\mathcal{N}(K, 2\varepsilon) \leq \mathcal{P}(K, \varepsilon) \leq \mathcal{N}(K, \varepsilon) \leq \mathcal{P}(K, \frac{\varepsilon}{2}) \quad (4.53)$$

as well as

$$\mathcal{P}(K, \varepsilon) \leq \frac{\text{vol}(K)}{\text{vol}(B_\varepsilon)} \leq \mathcal{N}(K, \varepsilon) \quad (4.54)$$

It follows then that

$$\mathcal{N}(A_i, \varepsilon_i) \leq \mathcal{N}(B_{R_i}, \varepsilon_i) - \mathcal{P}(B_{R_{i+1}}, \varepsilon_i) \quad (4.55)$$

$$\leq \mathcal{P}(B_{R_i, \frac{\varepsilon_i}{2}}) - \mathcal{N}(B_{R_{i+1}}, 2\varepsilon_i) \leq \frac{\text{vol}(B_{R_i})}{\text{vol}(B_{\frac{\varepsilon_i}{2}})} - \frac{\text{vol}(B_{R_{i+1}})}{\text{vol}(B_{2\varepsilon_i})} \quad (4.56)$$

$$= \frac{4^d \text{vol}(B_{R_i}) - \rho^d \text{vol}(B_{R_i})}{2^d \text{vol}(B_{\varepsilon_i})} \leq \frac{2^d}{\mu(\rho)^d}, \quad (4.57)$$

where in the last step we kept track of constants that depend on d . Summing over all annuli, the total number of trajectories required is

$$n \cdot \mathcal{N}(A_i, \varepsilon_i) \leq \frac{2^d}{\mu(\rho)^d} \frac{\ln\left(\frac{R}{\varepsilon}\right)}{\ln(\rho^{-1})} = \frac{\ln\left(\frac{R}{\varepsilon}\right) 2^d (2 + K_\beta)^d}{\rho^{\frac{d}{\beta}} \ln(\rho^{-1})}. \quad (4.58)$$

where the first step follows from (4.52) and (4.57), and the last from (4.51). Optimizing for the maximum $\rho \in (0, 1)$ in $g(\rho) = \rho^{\frac{d}{\beta}} \ln(\rho^{-1})$ leads to $\rho^* = e^{-\frac{\beta}{d}}$ and $g(\rho^*) = \frac{\beta}{d} \frac{1}{e}$. Substituting ρ^* in (4.58) and $\mu(\rho^*) = \left(\frac{e^{-\frac{1}{d}}}{C_\beta + 1}\right)^{\frac{1}{\beta}}$ in (4.50) finishes the proof. \square

Remark 4.1 (Performance vs Complexity Trade-off). *Theorem 4.8 highlight the intrinsic trade-off between the performance gap $\lambda - \alpha$ and the sample complexity of the verification process. Notably, when $\lambda - \alpha \rightarrow 0^+$, $\beta \rightarrow 0^+$, and $K_\beta \rightarrow \infty$. However, a constant gap enables us to get sample complexity that is exponentially better than the naive $O\left(\left(\frac{R}{\varepsilon}\right)^d\right)$.*

Remark 4.2 (Highly Performing Systems). *When $\lambda \rightarrow \infty$, $\beta, K_\beta \rightarrow 1$, thus leading to*

$$N(\varepsilon) = O\left(d \ln\left(\frac{R}{\varepsilon}\right) 6^d\right).$$

In this setting, since the trajectories converge very, very fast, one can use arbitrarily small $\tau \rightarrow 0^+$.

4.8 Numerical Methods

In this section, we build on Theorem 4.4 to develop parallelizable algorithms for certifying exponential ultimate boundedness around an equilibrium point. Rather than verifying exponential recurrence over a dense set, we use the neighborhood-based condition from Theorem 4.7 to efficiently certify local balls around individual sample points.

We present two complementary algorithms. The first (Section 4.8.2) takes a target region $S = B_R(x^*) \setminus B_\varepsilon(x^*)$ and computes the largest value of α for which $V(x) = \|x - x^*\|$ is an ERLF over S relative to $S' = B_R(x^*)$. The second (Section 4.8.3) takes a fixed decay rate α and incrementally constructs a certified region S where $V(x) := \|x - x^*\|$ is an ERLF relative to $S \cup B_\varepsilon(x^*)$. We validate both algorithms in Sections 4.8.4 and 4.8.5, respectively. In all of our algorithms below, L , τ , are precomputed ahead of time and globally shared. We further assume here that $x^* = 0$, and $V(x) = \|x\| := \max_{i \in [n]} |x_i|$, i.e., the ℓ_∞ -norm.

4.8.1 Supporting Methods and Algorithms

In this section we introduce a set of supporting routines that are used in our methods.

4.8.1.1 Verifying a Ball at a Fixed Point

Given a grid point $x \in \mathbb{R}^d$ and a radius $r > 0$, we first build a routine that determines the largest rate α for which the ball $B_r(x)$ satisfies the exponential recurrence condition (4.39). We formulate this as a simple one-dimensional optimization over α , based on Theorem 4.7. Note that checking (4.39b) is not necessary when S'

is a sub-level set of $V(x)$. Also, failure to certify is implied by $\alpha_{\max}(x, r; S') < 0$, and in particular, when (4.39b) fails $\alpha_{\max}(x, r; S') = -\infty$. For simplicity, we use $\alpha_{\max}(x, r) := \alpha_{\max}(x, r; B_R(x^*))$.

Algorithm 1: $\alpha_{\max}(x, r; S')$ — Maximum α certifying exponential recurrence for $B_r(x)$

Input: Center $x \in \mathbb{R}^d$, radius $r > 0$

Output: Maximum α such that $B_r(x)$ satisfies (4.39)

Find $\alpha^* = \max\{\alpha \mid \exists t \in (0, \tau] \text{ s.t. (4.39a)} \wedge \text{(4.39b)}\}$;

return α^* ;

4.8.1.2 Splitting Failed Points

In cases where Algorithm 1 fails to certify a ball $B_r(x)$ (i.e., $\alpha_{\max}(x, r) < 0$), the failure may result from evaluating the recurrence condition over a region that is too large to satisfy the assumptions of Theorem 4.7. To improve local resolution, we subdivide $B_r(x)$ into 3^d smaller balls of radius $r/3$, each centered on a uniformly spaced grid. This allows us to recursively refine problematic regions while maintaining full parallelizability. The subdivision procedure is described in Algorithm 2, and an illustration is provided in Figure 4.1 for the 2D case.




Figure 4.1: Splitting a Ball according to Algorithm 2

4.8.1.3 Initial Grid Setup

Given a center x^* , inner radius ε , and outer radius R , our goal is to verify exponential recurrence over the region $S := B_R(x^*) \setminus B_\varepsilon(x^*)$. To do this efficiently, we construct

Algorithm 2: Split(x, r) — Splitting a ball $B_r(x)$ into 3^d sub-balls

Input: Center $x = (x_i) \in \mathbb{R}^d$, radius $r > 0$ **Output:** List of sub-balls of radius $r/3$ Let $S \leftarrow \{-\frac{2}{3}r, 0, \frac{2}{3}r\}$; $Y \leftarrow \{(y_i) \in \mathbb{R}^d \mid y_i \in x_i + S, i \in [d]\}$;**splits** $\leftarrow \{(y, r/3) \mid y \in Y\}$;**return splits**;

a layered grid of candidate points. Specifically, we divide the region into m layers, where the radius of the center point in the ℓ -th layer is given by $r_\ell := 3^{\ell-1}\varepsilon$, for $\ell \in \{1, \dots, m\}$. Each layer contributes $3^d - 1$ points (excluding the origin), so the total number of initial grid points is $\mathcal{O}(3^d m)$.

To ensure that the full annular region S is covered, we choose the number of layers m such that:

$$R \leq \varepsilon + \sum_{\ell=1}^m 2r_\ell = \varepsilon \left(1 + 2 \sum_{\ell=1}^m 3^{\ell-1} \right) = 3^m \varepsilon.$$

This construction is exponentially more efficient than a uniform ε -grid, which would require $\mathcal{O}((R/\varepsilon)^d)$ points. A two-layer example is shown in Figure 4.2.

4.8.1.4 Estimation of L and τ

To apply Theorem 4.7, we require a one-sided Lipschitz constant L valid over a set S' that contains all trajectories starting from the region $S = B_R(x^*) \setminus B_\varepsilon(x^*)$. To this end, we seek a conservative outer set of the form $S' := B_{R'}(x^*)$ for some $R' > R$. We begin by constructing a uniform grid $G \subset \partial B_R(x^*)$ with spacing at most $\ell > 0$. From each point $x \in G$, we simulate the trajectory $\phi(t, x)$ over the interval $t \in [0, \tau]$, and verify that it reenters $B_R(x^*)$ within time τ . If any trajectory fails to return, τ must

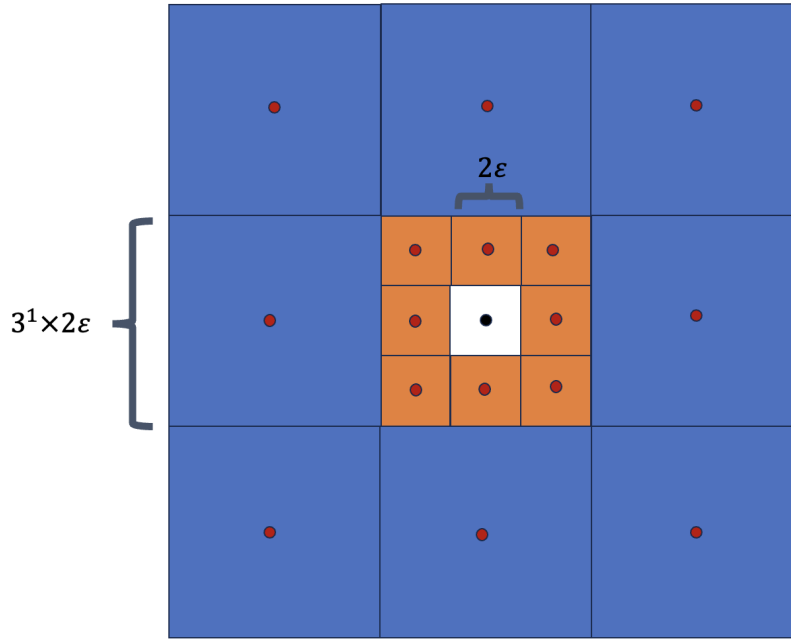


Figure 4.2: Illustration of Initial Grid Setup. We use $R = 3^m \varepsilon$, with $m = 2$ layers. The red dots are the grid points, while the central black dot is x^*

be increased.

Next, we compute the maximum excursion of all trajectories:

$R_{\max} := \max_{x \in G, t \in [0, \tau]} \|\phi(t, x) - x^*\|$, and set $R' := R_{\max} + \delta$, where $\delta > 0$ is a small buffer. Finally, we estimate the one-sided Lipschitz constant $L := \max_{z \in B_{R'}(x^*)} L_z$ using a fine grid. To ensure robustness with respect to perturbations around grid points, we verify the inequality:

$$\max_{t \in (0, \tau]} \max_{x \in G} (R' - \|\phi(t, x)\|) e^{-tL} \geq \ell.$$

This condition guarantees that all trajectories starting within an ℓ -neighborhood of the grid remain safely inside $B_{R'}(x^*)$ for the entire interval $[0, \tau]$. If this fails, we halve

the grid resolution ℓ and repeat the process until the condition is met.

4.8.2 Verification of a Region

We now integrate the routines developed in Section 4.8.1 into a complete algorithm for certifying exponential recurrence over the annular region $S := B_R(x^*) \setminus B_\varepsilon(x^*)$. We begin by computing a valid recurrence time τ and a one-sided Lipschitz constant L over a reachable set $S' \supseteq S$, as described in Section 4.8.1.4 (*Estimation of L and τ*). Using these constants, we construct an initial layered grid over S as explained in Section 4.8.1.3 (*Initial Grid Setup*). Each grid point x is associated with a ball $B_r(x)$, where the radius r scales with the distance to the equilibrium.

To evaluate recurrence at each point, we apply the routine $\alpha_{\max}(x, r)$ from Algorithm 1, which returns the largest value of α for which the ball satisfies condition (4.39). For each point x , we compute both:

- a **lower bound** $\underline{\alpha}(x) := \alpha_{\max}(x, r)$, which certifies condition (4.39) over the full ball,
- an **upper bound** $\bar{\alpha}(x) := \alpha_{\max}(x, 0)$, which certifies condition (4.39) only at the center.

If the worst-case relative $(\bar{\alpha}(x) - \underline{\alpha}(x))/\underline{\alpha}(x)$ the lowest $\underline{\alpha}(x)$ across the grid exceeds a specified threshold θ , we iteratively refine the k lowest-scoring points using the SPLIT routine (Algorithm 2). This process continues for at most m refinement steps.

The full verification procedure is summarized in Algorithm 3. If successful, it returns

Algorithm 3: Find- $\alpha_{\min}(R, \varepsilon, \theta)$ — Find best rate α for region $B_R(0) \setminus B_\varepsilon(0)$ via parallel ball certification

Input: Outer radius $R > 0$, inner radius $\varepsilon \in (0, R)$, threshold $\theta \in (0, 1)$, max number of refinements m

Output: Lower bound on the certified value of α

Construct initial grid $G \leftarrow \{(x_i, r_i)\}$ covering $B_R(x^*) \setminus B_\varepsilon(x^*)$; set counter $\leftarrow 0$;

while counter $\leq m - 1$ **do**

 For all $(x_i, r_i) \in G$, compute $\underline{\alpha}_i \leftarrow \alpha_{\max}(x_i, r_i)$ and $\bar{\alpha}_i \leftarrow \alpha_{\max}(x_i, 0)$;
 // upper and lower bounds

 Let $i^* \leftarrow \arg \min_{i \in [G]} \underline{\alpha}_i$; // lowest estimate of α in region

if $(\bar{\alpha}_{i^*} - \underline{\alpha}_{i^*}) / \underline{\alpha}_{i^*} > \theta$ **then**

 Replace the k lowest- $\underline{\alpha}_i$ balls (x_i, r_i) in G with $\text{SPLIT}((x_i, r_i))$;

 counter \leftarrow counter + 1;

end

else

break;

end

end

return $\underline{\alpha}_{i^*}$

a uniform lower bound α such that $V(x) = \|x\|_\infty$ is an ERLF over S , as guaranteed by Theorem 4.4.

4.8.3 Exponential Region of Attraction Mining

We now consider the complementary task to Section 4.8.2: given a fixed decay rate $\alpha > 0$, identify a maximal subset of the state space over which $V(x) = \|x\|_\infty$ can be certified as an Exponential Recurrent Lyapunov Function (ERLF) with rate at least α .

To that end, we define an initial covering grid $G_0 := \{(x_i, r_i)\}$ over the annular region $B_R(x^*) \setminus B_\varepsilon(x^*)$, where each pair (x_i, r_i) represents a candidate ball $B_{r_i}(x_i)$.

The union of all such balls forms a discrete approximation to the candidate region of attraction. Using this grid, we iteratively apply the ball certification routine $\alpha_{\max}(x_i, r_i)$ (Algorithm 1). Balls that pass the condition $\alpha_{\max}(x_i, r_i) \geq \alpha$ are added to the certified set. Those that fail are recursively subdivided using the SPLIT routine, up to a maximum number of refinements. This growth procedure is encapsulated in Algorithm 4. To improve efficiency, Algorithm 4 supports a Boolean flag `Trim`. When set to `False`, each ball is verified using only local information. Once the region has been expanded, the same algorithm is rerun with `Trim` set to `True`, ensuring that all recurrence trajectories remain entirely within the certified region (as required by Theorem 4.4).

The result is a data-driven, self-consistent inner approximation of the α -region of attraction.

4.8.4 Numerical Validation of Algorithm 3

We continue by providing a preliminary validation of the proposed Algorithm 3. To investigate the efficiency of our proposed method, we consider the following systems:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + B_1 \begin{bmatrix} x_1^2 \\ x_1 x_2 \\ x_2^2 \end{bmatrix}; \quad (4.59)$$

Algorithm 4: Find- α -RoA($R, \varepsilon, \alpha, G_0, \text{Trim}$) — Find certified region for fixed exponential rate α

Input: Target $\alpha > 0$, range parameters $R > 0, \varepsilon \in (0, R)$, initial grid $G_0 = \{(x_i, r_i)\}$, boolean Trim , max splits m

Output: Subset of G_0 certifying $\|x\|_\infty$ as an ERLF with rate at least α

Set $G \leftarrow G_0, \text{Positives} \leftarrow \emptyset, \text{counter} \leftarrow 0$;

while $\text{counter} \leq m - 1$ **and** $G \neq \emptyset$ **do**

Define $\text{Region} \leftarrow \bigcup_{(x,r) \in \text{Positives} \cup G} B_r(x)$; // Used if trimming is enabled

foreach $(x_i, r_i) \in G$ **do**

$\alpha_i \leftarrow \text{Trim } \alpha_{\max}(x_i, r_i; \text{Region}) : \alpha_{\max}(x_i, r_i)$; // Evaluate recurrence

if $\alpha_i \geq \alpha$ **then**

$\text{Positives} \leftarrow \text{Positives} \cup \{(x_i, r_i)\}$; // Add to certified set

end

end

$G \leftarrow \text{SPLIT}(G \setminus \text{Positives})$; // Refine remaining uncertified balls

$\text{counter} \leftarrow \text{counter} + 1$;

end

return Positives ; // Region certified for recurrence at rate α

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} -1 & 0 & 0 \\ 0.5 & -1 & 0 \\ 0.5 & 0.5 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + B_2 \begin{bmatrix} x_1^2 \\ \vdots \\ x_3^2 \end{bmatrix}, \quad (4.60)$$

where $B_1 \in \mathbb{R}^{2 \times 3}$ and $B_2 \in \mathbb{R}^{3 \times 9}$ are drawn independently from a Gaussian distribution, i.e., $[B_1]_{ij}, [B_2]_{ij} \sim \mathcal{N}(0, \sigma)$. We will increase the standard deviation σ as a means to increase the complexity of the dynamics. In our experiments, we choose the ℓ_∞ norm as our choice of $V(x)$ and as the norm used to measure distances between trajectories. Thus, $\alpha_1 = \alpha_2 = 1$. Sample trajectories for the system (4.59) with $\sigma = 0.3$ are shown in Figure 4.3, where we also illustrate the ball of radius R (blue) selected, the

computed ball of radius R' (red), and the small region around the origin (x^*) not certified (black). We also show in Figure 4.4 the verified region and a coloring scheme illustrating the different ball sizes used at different parts of $B_R(x^*) \setminus B_\varepsilon(x^*)$.




Figure 4.3: Phase Portrait of System (4.59). The black box surrounds the region which we do not verify ($B_\varepsilon(x^*)$), the blue box represents surrounds the region which we verify in Algorithm 3 ($B_R(x^*)$), and the red box surrounds the region which trajectories that begin in the blue box do not leave ($B_{R'}(x^*)$).




Figure 4.4: Sizes of blocks resulting from applying Algorithm 3 to system (4.59).

In these experiments we use $R = 0.7$ and $\varepsilon = 0.01$. Table 4.1 and Table 4.2 summarize the results obtained by running Algorithm 3 together with a comparison with SOS-TOOLS. When running our algorithm, we use the Torchode toolbox [69] to compute system trajectories in parallel. We also show the solving times of our algorithm and SOSTOOLS [94]. It can be seen that as σ grows - i.e., the system becomes more

nonlinear - and the dimension grows, our algorithm begins to significantly outperform SOSTOOLS.

2D sys. (4.59)	0.3,	0.6	1
Alg 3 α :	0.470	0.414	0.349
SoS α :	0.360	0.247	0.223
Alg 3 Time:	10.95s	15.01s	15.30s
SoS Time:	0.97s	1.06s	0.94s

Table 4.1: Parameter values and performance comparison between our algorithm and the SOSTOOLS for system (4.59)

3D sys. (4.60)	$\sigma = 0.1,$	0.3	0.4	0.5
Alg 3 α :	0.635	0.449	0.370	0.359
SoS α :	0.309	0.341	0.256	0.213
Alg 3 Time:	140.85s	144.33s	150.07s	154.92s
SoS Time:	54.89s	276.10s	299.77s	632.55s

Table 4.2: Parameter values and performance comparison between our algorithm and the SOSTOOLS for system (4.60)

4.8.5 Numerical Validation of Algorithm 4

We end this section by providing preliminary experiments for Algorithm 4. To do so, we consider the Kuramoto Oscillator with uniform coupling constants, i.e., for an n -dimensional system and some constant k , oscillator θ_i is defined by:

$$\dot{\theta}_i = \frac{k}{n} \sum_{j=1}^n \sin(\theta_j - \theta_i) \quad (4.61)$$

To eliminate the rotational symmetry of this system, we consider the change of variables $\varphi_i = \theta_i - \theta_n$, which effectively reduces the dimensions by one.

We next consider the use of Algorithm 4 to find the α -RoA. First we consider the two-oscillator system with $\alpha = 1$ and $k = 10$. Figures 4.5 and 4.6 are the phase

portraits along with the 1-RoA formed of the verified balls, when usnig maximum split counts of 1 and 6, respectively.

Finally, we investigate how the size of the certified region and the computational cost of verification scale with the ambient dimension. Specifically, we fix the system dynamics and all algorithmic parameters, and vary only the state dimension. Figure 4.7 reports the volume of the region certified by Algorithm 4 and the corresponding computation time, for dimensions ranging from 2 to 6 and maximum split counts from 0 to 6 (represented by the dots).




Figure 4.5: Phase Portrait of System (4.61), 1 Split. The the blue region composed of distinct cells is the verified region resulting from Algorithm 4 with a maximum split count of 1.




Figure 4.6: Phase Portrait of System (4.61), 6 Splits. The blue region composed of distinct cells is the verified region resulting from Algorithm 4 with a maximum split count of 6.




Figure 4.7: Area of region Verified by Algorithm 4 for system 4.61 and Associated Computational Time Requirements. Split counts range from 0 to 6, with each split count corresponding to a dot and dimensions 2 through 6. Time is provided in log-scale.

Chapter 5

Non-Parametric Chain Policies for Stabilizing Control

There is no single heroic narrative. There are a thousand small acts of courage that hold the world together.

– LORD RABBI JONATHAN SACKS

This chapter is adapted from the following article:

5.1 Preliminaries

We consider a nonlinear control system:

$$\dot{x}(t) = f(x(t), u(t)), \tag{5.1}$$

with state $x(t) \in \mathbb{R}^n$ and input $u(t) \in U \subseteq \mathbb{R}^m$. We define

$$\mathcal{U}^{(a,b]} := \{u : (a, b] \rightarrow U \mid u \text{ measurable}\},$$

as the set of admissible control signals on interval $(a, b]$, and set $\mathcal{U} := \mathcal{U}^{(0,\infty)}$. Given $u_0 \in \mathcal{U}^{(0,a]}$ and $u_1 \in \mathcal{U}^{(0,b]}$, their concatenation $u_0u_1 \in \mathcal{U}^{(0,a+b]}$ is defined by

$$(u_0u_1)(t) = \begin{cases} u_0(t), & t \in (0, a], \\ u_1(t), & t \in (a, a + b]. \end{cases}$$

More generally, for a sequence of control signals $u_n \in \mathcal{U}^{(0, \tau_n]}$, with $\tau_n > 0, \forall n \in \mathbb{N}$, we further use $u_{[n]} := u_0 u_1 \dots u_n$, and $u_{[\infty]} = \lim_{n \rightarrow \infty} u_{[n]}$. In some occasions we slightly abuse notation by using u interchangeably to represent instantaneous inputs in U and signals in $\mathcal{U}^{(a, b]}$; the intended meaning will always be clear from context.

For an initial state $x \in \mathbb{R}^n$ and control signal $u \in \mathcal{U}^{(0, a]}$, we denote by $\phi(t, x, u)$ the solution of (5.1) for $t \in (0, a]$. We further assume the following regularity conditions for (5.1).

Assumption 5.1 (Forward Completeness). *The solutions of the control system (5.1) are **forward complete**. Specifically, for each initial condition $x \in \mathbb{R}^n$ and every control signal $u \in \mathcal{U}$, the trajectory $\phi(t, x, u)$ exists and remains bounded for all $t \geq 0$.*

Assumption 5.2 (Uniform Lipschitz Continuity). *The vector field $f(x, u)$ of system (5.1) is locally Lipschitz continuous in x , uniformly with respect to u . More precisely, for every compact set $S \subseteq \mathbb{R}^n$, there exists a constant $L_S \geq 0$ such that*

$$\|f(y, u) - f(x, u)\| \leq L_S \|y - x\|, \quad \forall x, y \in S, \forall u \in U.$$

5.1.1 Practical Exponential Stabilizability

In this work, we aspire to render an equilibrium point $x^* \in \mathbb{R}^n$ practical exponentially stable.

Definition 5.1 (Equilibrium Point). *A point $x^* \in \mathbb{R}^n$ is an **equilibrium point** of system (5.1) if there exists a control input $u^* \in U$ such that $f(x^*, u^*) = 0$.*

Definition 5.2 ((Practical) Exponential Stabilizability). *Let $S \subseteq \mathbb{R}^n$. The equilibrium x^* of system (5.1) is said to be:*

- (i) **Exponentially Stabilizable** on S if for every $x \in S$, there exists a control signal $u \in \mathcal{U}$ satisfying

$$\|\phi(t, x, u) - x^*\| \leq Ke^{-\lambda t}\|x - x^*\|, \quad \forall t \geq 0; \quad (5.2)$$

- (ii) **Practically Exponentially Stabilizable** on S if for every $x \in S$, there exists a control signal $u \in \mathcal{U}$ satisfying

$$\|\phi(t, x, u) - x^*\| \leq Ke^{-\lambda t}\|x - x^*\| + c, \quad \forall t \geq 0, \quad (5.3)$$

for constants $K \geq 1$, $\lambda > 0$, and $c \geq 0$.

It is well known from the topological entropy literature that it is impossible to exponentially stabilize a system, i.e., achieve (5.2), using a finite number of control signals [32]. We will therefore aim to enforce the weaker notion of practical exponential stability, i.e., (5.3), which follows the terminology of [33, 48].

5.1.2 Recurrent Lyapunov Functions

To provide guarantees for our data-driven stabilization framework, we build on the theory of *Recurrent Lyapunov Functions (RLFs)* [110]. Unlike classical Lyapunov functions, which require strict decrease along trajectories, RLFs only require a decrease at a sequence of *recurrent times*. This relaxation broadens the class of certificates available, while still ensuring exponential stability.

We begin with the notion of containment times, which we define for general trajectories of the controlled system.

Definition 5.3 (Containment Times). *Given a set $S \subset \mathbb{R}^n$, an initial state $x \in \mathbb{R}^n$, and an input $u \in \mathcal{U}$, the set of containment times is*

$$T_S(x, u) := \{t \in \mathbb{R}_{>0} \mid \phi(t, x, u) \in S\}.$$

For constants $a, b > 0$ we define

$$T_S(x, u; a, b) := T_S(x, u) \cap (a, a + b],$$

and for convenience $T_S(x, u; b) := T_S(x, u; 0, b)$.

We now recall the definition of an RLF in the *autonomous* case, where the trajectory is uniquely determined by the initial condition. In this case we write $\phi(t, x)$ for the flow of the system.

Definition 5.4 (Recurrent Lyapunov Function). *Let $S \subset \mathbb{R}^n$ be a compact set with $x^* \in \text{int}(S)$. A continuous function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is called a Recurrent Lyapunov Function (RLF) over S with rate $\alpha > 0$ and horizon $\tau > 0$ if*

$$\min_{t \in T_S(x; \tau)} e^{\alpha t} V(\phi(t, x)) \leq V(x), \quad \forall x \in S, \quad (5.4)$$

where $T_S(x; \tau) := \{t \in (0, \tau] \mid \phi(t, x) \in S\}$.

It will also be useful to characterize the set of points that can be reached within a finite interval of time.

Definition 5.5 (Reachable Tube). *For the control system (5.1), a constant $\tau > 0$, and a set $S \subset \mathbb{R}^n$, we denote the τ -reachable tube from S within τ units of time by*

$$\mathcal{R}^\tau(S) = \bigcup_{x \in S, u \in \mathcal{U}, t \in [0, \tau]} \{\phi(t, x, u)\}.$$

5.2 Recurrent Control Lyapunov Functions

As mentioned before, our guarantees rely on the theory of Recurrent Lyapunov Functions (RLFs) from [110]. In this section, we extend this notion to the control setting, introducing Recurrent Control Lyapunov Functions (R-CLFs), and illustrate how they can be used to certify practical stabilizability. Though RLFs and R-CLFs have been shown to certify stability, asymptotic stability and exponential stability [110], our focus here is on practical exponential stability and thus we will use the following definition.

Definition 5.6 (Recurrent Control Lyapunov Function (R-CLF)). *Consider the control system (5.1) with equilibrium $x^* \in \mathbb{R}^n$. Let $S \subseteq \mathbb{R}^n$ be a set satisfying $x^* \in \text{int}(S)$. A continuous function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is a **Recurrent Control Lyapunov Function (R-CLF)** over S if the following conditions hold:*

- (i) **Positive Definiteness and Linear Bounds:** *There exist constants $a_1, a_2 > 0$ such that*

$$a_1 \|x - x^*\| \leq V(x) \leq a_2 \|x - x^*\|, \quad \forall x \in S. \quad (5.5)$$

- (ii) **Control α -Exponential (τ, δ) -Recurrence:** *There exist constants $\tau, \alpha > 0$*

and $\delta \geq 0$ such that for every $x \in S$, there exists $u \in \mathcal{U}^{[0,\tau]}$ satisfying

$$\min_{t \in T_S(x,u;\tau)} e^{\alpha t} (V(\phi(t, x, u)) - \delta) \leq [V(x) - \delta]_+. \quad (5.6)$$

The following lemma characterizes the long term behavior of the control system (5.1) under the controls $u \in \mathcal{U}$ that are build upon concatenation of controls satisfying property (ii) of Definition 5.6.

Lemma 5.1 (Characterization of R-CLF). *Let assumptions 5.1 and 5.2 hold. Consider an equilibrium x^* of (5.1) and a compact set S satisfying $x^* \in \text{int}(S)$. A function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ satisfying (5.5) is a Recurrent Control Lyapunov Function (R-CLF) over S if and only if there exists parameters $\alpha, \tau > 0$ and $\delta \geq 0$ such that for any $x \in S$ there is a sequence $\{t_n\}_{n \in \mathbb{N}}$ and $u \in \mathcal{U}$ satisfying the following conditions:*

$$\lim_{n \rightarrow +\infty} t_n = +\infty, \text{ with } t_{n+1} - t_n \in (0, \tau], \quad (5.7a)$$

$$\phi(t_n, x, u) \in S, \quad \text{and} \quad (5.7b)$$

$$V(\phi(t_n, x, u)) - \delta \leq \begin{cases} e^{-\alpha t_n} (V(x) - \delta), & n \leq \bar{n}, \\ 0, & \text{o.w.}, \end{cases} \quad (5.7c)$$

for a non necessarily finite $\bar{n} \in \mathbb{N} \cup \{\infty\}$.

Proof of Lemma 5.1. We prove each direction separately.

Necessity (\Rightarrow): Suppose that V is a Recurrent Control Lyapunov Function (R-CLF) over the compact set S . By Definition 5.6, there exist constants $\alpha, \tau > 0$ and $\delta \geq 0$

such that, for any $x \in S$, there exists $\bar{u} \in \mathcal{U}^{[0,\tau]}$ satisfying (5.6).

We will build $u \in \mathcal{U}$ and the sequence $\{t_n\}_{n \in \mathbb{N}}$ inductively. Let $t_0 = 0$, $x_0 := x \in S$, and define for $n \geq 0$,

$$\tau_n := \max \left\{ \arg \min_{t \in T_S(x_n, \bar{u}_n; \tau)} e^{\alpha t} (V(\phi(t, x_n, \bar{u}_n)) - \delta) \right\}, \quad (5.8)$$

$$t_{n+1} := t_n + \tau_n, \quad \text{and} \quad x_{n+1} := \phi(\tau_n, x_n, u_n), \quad (5.9)$$

where $\bar{u}_n \in \mathcal{U}^{[0,\tau]}$ is a control satisfying (5.6), u_n is its restriction to the interval $(0, \tau_n]$, and $x_n \in S \forall n \in \mathbb{N}$, by definition. Next, let $u := \lim_{t \rightarrow \infty} u_{[n]}$. Note that $\phi(t_0, x, u) = x_0 \in S$, and whenever for some n , $\phi(t_n, x, u) = x_n \in S$ that by the group property of the flow

$$\phi(t_{n+1}, x, u) = \phi(t_{n+1} - t_n, \phi(t_n, x, u), u_n) \quad (5.10)$$

$$= \phi(\tau_n, x_n, u_n) = x_{n+1} \in S, \quad (5.11)$$

which by induction ensures that (5.7b) holds.

Next, from the recurrence condition (5.6), it follows that as long as $V(\phi(t_n, x, u)) \geq \delta$, then $e^{\alpha \tau_n} (V(\phi(\tau_n, x_n, \bar{u}_n)) - \delta) \leq V(x_n) - \delta$, which implies

$$e^{\alpha t_{n+1}} (V(\phi(t_{n+1}, x, u)) - \delta) \leq e^{\alpha t_n} (V(\phi(t_n, x, u)) - \delta), \quad (5.12)$$

and, in particular,

$$e^{\alpha t_n} (V(\phi(t_n, x, u)) - \delta) \leq V(x) - \delta, \quad \forall n \leq \bar{n}, \quad (5.13)$$

where \bar{n} is the last instance with $V(x_{\bar{n}}) \geq \delta$. It also follows from (5.6) that, when $\bar{n} < \infty$, for all $n \geq \bar{n} + 1$, $V(\phi(t_n, x, u)) \leq \delta$, which completes (5.7c).

To show (5.7a), we first note that by definition, $t_{n+1} - t_n = \tau_n \in (0, \tau]$. Next, we will show that $t^* = \infty$. Suppose not, i.e., $t^* < \infty$. By continuity of $\phi(t, x, u)$ and compactness of S , $\phi(t_n, x, u) \rightarrow \phi(t^*, x, u) \in S$.

Now, let $v_n := V(x_n) - \delta$ and $v^* = V(\phi(t^*, x, u)) - \delta$. If $v^* > 0$, it follows from (5.12) and the continuity of V that $v_n \downarrow v^* := V(\phi(t^*, x, u)) - \delta$, and $e^{\alpha t^*} v^* \leq e^{\alpha t_{n+1}} v_{n+1} \leq e^{\alpha t_n} v_n$. Thus, for large enough n , $t^* \in (t_n, t_n + \tau]$ and $t^* > t_{n+1}$, which implies

$$e^{\alpha t^* - t_n} v^* \leq e^{\alpha \tau_n} v_{n+1} \leq v_n,$$

which contradicts τ_n being the max in (5.8). As similar argument holds when $v^* < 0$, and thus $t^* = \infty$.

Sufficiency (\Leftarrow): To show that V is a R-CLF, it is sufficient for any $x \in S$ to restrict the corresponding $u \in \mathcal{U}$ that satisfies (5.7) to the interval $[0, \tau)$ and choosing t_1 from the sequence (5.7) to show (5.6).

□

We will leverage Lemma 5.1 to prove (practical) exponential convergence of trajectories. To that end, we need to bound how much a trajectory can travel in between instance of exponential convergence (5.7c). The following lemma provides a mechanism to obtain such bounds. The proof is based on Grönwall’s Lemma [62, Lemma A.1] and can be found in [108].

Lemma 5.2 (Containment Lemma). *Let Assumption 5.2 hold. Consider a compact set $S \subset \mathbb{R}^n$ and a constant $\tau > 0$. Then, for any $x \in S, u \in \mathcal{U}$ the following holds:*

$$\max_{t \in [0, \tau]} d(\phi(t, x, u), S) \leq F_S \tau e^{L\tau} \quad (5.14)$$

where $L := L_{\mathcal{R}^\tau(S)}$.

We are now ready to show that R-CLFs as defined in Definition 5.6 guarantee exponential stabilizability.

Theorem 5.1 (R-CLF Implies (Practical) Exponential Stabilizability). *Consider the control system (5.1) with equilibrium $x^* \in \mathbb{R}^n$, and let S be a set satisfying $S \subseteq \mathbb{R}^n$ and $x^* \in \text{int}(S)$. Let Assumption 5.1 and Assumption 5.2 hold, and $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ be a Recurrent Control Lyapunov Function over S , with constants $\alpha, \tau > 0, \delta \geq 0$, and linear bound constants $a_1, a_2 > 0$ from (5.5).*

Then, the equilibrium x^ is (practically) exponentially stabilizable on S (when $\delta > 0$). In particular, for every initial condition $x \in S$, a control signal $u \in \mathcal{U}$ satisfying (5.7) for some sequence of times $\{t_n\}_{n \in \mathbb{N}}$ ensures:*

$$\|\phi(t, x, u) - x^*\| \leq K e^{-\lambda t} \|x - x^*\| + c, \quad \forall t \geq 0, \quad (5.15)$$

where

$$\lambda := \alpha, \quad K := \frac{a_2}{a_1} e^{\alpha\tau} (1 + L\tau e^{L\tau}), \quad \text{and} \quad c := \frac{\delta}{a_1} (1 + L\tau e^{L\tau}), \quad (5.16)$$

with $L := L_{R^\tau(S)}$.

Proof. We will use the control $u \in \mathcal{U}$ from Lemma 5.1 to prove this theorem. Given $x \in S$, by Lemma 5.1, there exists $u \in \mathcal{U}$ and a sequence $t_n, x_n := \phi(t_n, x) \in S$ satisfying (5.7). There are two cases. First, assume that $n \leq \bar{n} \in \mathbb{N} \cup \{\infty\}$. Since $V(x_n) \leq e^{-\alpha t_n} (V(x) - \delta) + \delta$ and $V(x_n) \geq a_1 \|x_n - x^*\|$, it follows that

$$r_n := \|x_n - x^*\| \leq \frac{1}{a_1} V(x_n) \leq \frac{a_2}{a_1} e^{-\alpha t_n} \|x - x^*\| + \frac{\delta}{a_1}.$$

Consider any time $t \in (t_n, t_{n+1}]$ and $B_n := B_n(x^*) \cap S$. By applying the containment lemma on B_n and using the fact that $L \geq L_{\mathcal{R}^\tau(B_n)}$, we get, using the triangle inequality:

$$\begin{aligned} \|\phi(t, x, u) - x^*\| &\leq \|x_n - x^*\| + \|\phi(t, x, u) - x_n\| \\ &\leq r_n + F_{r_n} \tau e^{L\tau} \leq (1 + L\tau e^{L\tau}) r_n \end{aligned}$$

where $r_n := \|x_n - x^*\|$. Then for any $n \leq \bar{n}$ and any $t \in (t_n, t_{n+1}]$ we have,

$$\|\phi(t, x, u) - x^*\| \leq (1 + L\tau e^{L\tau}) r_n \leq (1 + L\tau e^{L\tau}) \left(\frac{a_2}{a_1} e^{-\alpha t_n} \|x - x^*\| + \frac{\delta}{a_1} \right) \quad (5.17)$$

$$\leq K e^{-\alpha\tau} e^{-\alpha t_n} \|x - x^*\| + \frac{\delta}{a_1} (1 + L\tau e^{L\tau}) \quad (5.18)$$

$$\leq K e^{-\alpha t} \|x - x^*\| + c \quad (5.19)$$

where the last step follows, since $t \leq t_{n+1} \leq t_n + \tau$, which implies $-t_n - \tau \leq -t$, so

that $e^{-\alpha\tau}e^{-\alpha t_n} \leq e^{-\alpha t}$.

If $\bar{n} = \infty$ we are done. Otherwise, consider $n > \bar{n}$. We have $r_n = \|x_n - x^*\| \leq \frac{V(x_n)}{a_1} \leq \frac{\delta}{a_1} \forall n > \bar{n}$, and thus by Lemma 5.2 again, for all $t > t_{\bar{n}}$

$$\|\phi(t, x, u) - x^*\| \leq (1 + L\tau e^{L\tau})r_n \leq (1 + L\tau e^{L\tau})\frac{\delta}{a_1} = c.$$

Thus, for all $t \geq 0$,

$$\|\phi(t, x, u) - x^*\| \leq Ke^{-\lambda t}\|x - x^*\| + c,$$

as desired. □

Theorem 5.1 states that the existence of an R-CLF implies that x^* can be made practically exponentially stable. At the core of its proof is the fact that one can find a function V that satisfies the recurrent condition (5.6). A key observation of [110], is that condition (5.6) can be met by a norm, provided τ and α are properly chosen (c.f. [110, Theorem 6]). The the caveat is, however, that in order to make R-CLFs practically useful, one would need to store, for each $x \in S$, a suitable $u : [0, \tau) \rightarrow U$ that ensures (5.6). In the next section we surprisingly show that when $\delta > 0$, only a finite number of such signals are needed.

5.3 Non-Parametric Chain Policies

In the previous section, we introduced Recurrent Control Lyapunov Functions (R-CLFs) to characterize exponential stabilizability via carefully selected control signals. In this

section, we propose *nonparametric chain policies*, a systematic, data-driven approach for generating these stabilizing signals. The proposed method aligns closely with recent developments in topological entropy—a notion quantifying the minimal complexity required to accomplish various control tasks (see, e.g., [31–33]). A distinctive feature of our method is that we do not assume the control signals can be generated online; instead, we explicitly store them in a finite set, called a *control alphabet* [106].

Definition 5.7 (Control Alphabet). *A **control alphabet** is a finite collection of control signals*

$$\mathcal{A} := \{v_i : (0, \tau_i] \rightarrow U\}_{i=0}^M, \quad (5.20)$$

where each v_i is piecewise continuous and defined over a duration $\tau_i > 0$.

The control alphabet provides a library of candidate signals. To deploy them, we assign specific controls to regions of influence within the state space. To aid this task we define an assignment set.

Definition 5.8 (Assignment Set). *An **assignment set** is a finite collection of verification triples*

$$\mathcal{K} := \{(x_i, r_i, v_i)\}_{i=1}^N \subseteq \mathbb{R}^n \times \mathbb{R}_{>0} \times \mathcal{A},$$

where $x_i \in \mathbb{R}^n$ is a center point, $r_i > 0$ is its radius, and $v_i \in \mathcal{A}$ is the control signal assigned to that region. The support of \mathcal{K} is

$$\text{Supp}(\mathcal{K}) := \bigcup_{i=1}^N B_{r_i}(x_i).$$

We denote $N := |\mathcal{K}|$ as the size of the assignment set.

While an assignment set specifies regions of influence, it does not by itself resolve which control to apply when balls overlap, nor what to do when a state lies outside $\text{Supp}(\mathcal{K})$. To address this, we introduce a normalized nearest-neighbor rule with a fall-back option:

$$\iota_{\mathcal{K}}(x) := \begin{cases} \arg \min_{i:(x_i, r_i, v_i) \in \mathcal{K}} \frac{\|x - x_i\|}{r_i}, & r_{\mathcal{K}}(x) \leq 1, \\ 0, & \text{otherwise,} \end{cases} \quad (5.21)$$

where

$$r_{\mathcal{K}}(x) := \min_{(x_i, r_i, v_i) \in \mathcal{K}} \frac{\|x - x_i\|}{r_i},$$

and $\iota_{\mathcal{K}}(x) = 0$ corresponds to selecting the default control v_0 .

Remark 5.1. We designate $v_0 \in \mathcal{A}$ as the default control. Unless otherwise stated, we take $v_0(t) = u^* \in U$, $\forall t \in [0, \tau_0)$, where u^* is the equilibrium control of Definition 5.1.

The index map $\iota_{\mathcal{K}}$ specifies, for any state x , which control from the assignment set (or the default control) should be applied. Building on this rule, we can now formalize the induced feedback policy.

Definition 5.9 (Nonparametric Chain Policy). *Given an assignment set \mathcal{K} and default control v_0 , the **nonparametric chain policy** is given by the map $\pi_{\mathcal{K}} : \mathbb{R}^n \rightarrow \mathcal{A}$:*

$$\pi_{\mathcal{K}}(x) := v_{\iota_{\mathcal{K}}(x)}. \quad (5.22)$$

Remark 5.2. The policy $\pi_{\mathcal{K}}$ induces an infinite-horizon control signal $u_{\mathcal{K},x} \in \mathcal{U}$ through concatenation. Starting with $x_0 = x$ and the empty signal $u_{[0]} = \emptyset$, for each

$n \geq 0$ define

$$u_{[n+1]} = u_{[n]} v_{\iota_{\mathcal{K}}(x_n)}, \quad x_{n+1} = \phi\left(\tau_{\iota_{\mathcal{K}}(x_n)}, x_n, v_{\iota_{\mathcal{K}}(x_n)}\right). \quad (5.23)$$

The resulting control is then

$$u_{\mathcal{K},x} := \lim_{n \rightarrow \infty} u_{[n]}.$$

5.3.1 Convergence Guarantees of NCPs

With the nonparametric chain policy in place, we now turn to its stability properties. The following theorem establishes conditions under which such a policy renders the equilibrium x^* practically exponentially stable on a prescribed region.

Theorem 5.2 (Practical Exponential Stabilization via Chain Policies). *Consider an equilibrium point $x^* \in \mathbb{R}^n$ of (5.1), and let $S \subseteq \mathbb{R}^n$ be a set with $x^* \in \text{int}(S)$. Let $\pi_{\mathcal{K}}$ denote a nonparametric chain policy associated with the assignment set $\mathcal{K} = \{(x_i, r_i, v_i)\}_{i=1}^N$ and a default control $v_0 \in \mathcal{A}$, and define $\tau := \max\{\tau_0, \tau_1, \dots, \tau_N\}$, and let $L := L_{\mathcal{R}^\tau(S)}$. Suppose the following hold:*

(i) **Covering.** *There exists $\varepsilon > 0$ such that*

$$B_\varepsilon(x^*) \subset B_{\varepsilon(1+L\tau e^{L\tau})}(x^*) \subset \text{int}(S), \quad (5.24a)$$

$$\text{cl}(S \setminus B_\varepsilon(x^*)) \subseteq \text{Supp}(\mathcal{K}). \quad (5.24b)$$

(ii) **Verification.** *For each $(x_i, r_i, v_i) \in \mathcal{K}$ with $\tau_i > 0$,*

$$e^{\alpha\tau_i} \left(\|\phi(\tau_i, x_i, v_i) - x^*\| + r_i e^{L\tau_i} \right) \leq \|x_i - x^*\| - r_i, \quad (5.25a)$$

$$\text{sd}(\phi(\tau_i, x_i, v_i), S) + r_i e^{L\tau_i} \leq 0, \quad (5.25b)$$

(iii) **Equilibrium.** For all $t \in [0, \tau_0)$, $\phi(t, x^*, v_0) = x^*$.

Then the equilibrium x^* is practically exponentially stable on S under the policy $\pi_{\mathcal{K}}$, with constants

$$\lambda = \alpha, \quad K = e^{\alpha\tau}(1 + L\tau e^{L\tau}), \quad c = \varepsilon(1 + L\tau e^{L\tau}).$$

Proof. We will show that the control $u_{\mathcal{K},x}$ induced by the nonparametric chain policy $\pi_{\mathcal{K}}$ admits a sequence of times $\{t_n\}_{n \in \mathbb{N}}$ that satisfies the conditions of Lemma 5.1 for the function $V(x) = \|x - x^*\|$ over the set S . This establishes two points: (1) $V = \|\cdot - x^*\|$ is an R-CLF with rate α , and (2) the control $u_{\mathcal{K},x}$ practically stabilizes x^* with exponential rate α over S .

Let $x_0 = x \in S$, $t_0 = 0$, and $u_{[0]} = \emptyset$. Define the sequences $\{x_n\}$, $\{t_n\}$, and $u_{[n]}$ according to (5.23), i.e.,

$$x_{n+1} = \phi(\tau_{\iota_{\mathcal{K}}(x_n)}, x_n, v_{\iota_{\mathcal{K}}(x_n)}), \quad u_{[n+1]} = u_{[n]} v_{\iota_{\mathcal{K}}(x_n)}, \quad (5.26)$$

$$u_{\mathcal{K},x} = \lim_{n \rightarrow \infty} u_{[n]}, \quad t_{n+1} := t_n + \tau_{\iota_{\mathcal{K}}(x_n)}, \quad \forall n \geq 0. \quad (5.27)$$

By construction, for all $n \geq 0$,

$$0 < \min_{i \in \{0, \dots, N\}} \tau_i \leq t_{n+1} - t_n \leq \max_{i \in \{0, \dots, N\}} \tau_i =: \tau,$$

so condition (5.7a) holds. Moreover, by induction one shows that for all $n \geq 1$,

$$\phi(t_n, x, u_{\mathcal{K},x}) = x_n = \phi(\tau_{\nu_{\mathcal{K}}(x_{n-1})}, x_{n-1}, v_{\nu_{\mathcal{K}}(x_{n-1})}).$$

We claim that if $x_n \in S$ then $x_{n+1} \in S$. Suppose first that $x_n \in S \setminus B_\varepsilon(x^*)$. By the covering condition (5.24b), there exists $(x_i, r_i, v_i) \in \mathcal{K}$ such that $x_n \in B_{r_i}(x_i)$ and the verification condition (5.25) holds. In particular, by (5.25b),

$$\text{sd}(x_{n+1}, S) = \text{sd}(\phi(\tau_i, x_n, v_i), S) \tag{5.28}$$

$$\leq \text{sd}(\phi(\tau_i, x_i, v_i), S) + r_i e^{L\tau_i} \tag{5.29}$$

$$\leq 0, \tag{5.30}$$

which implies $x_{n+1} \in S$.

If instead $x_n \in B_\varepsilon(x^*)$, then either $\nu_{\mathcal{K}}(x_n) \neq 0$ and the above argument applies, or $\nu_{\mathcal{K}}(x_n) = 0$, in which case we apply v_0 for time τ_0 . By the containment lemma applied to the ball $B_\varepsilon(x^*)$,

$$\|x_{n+1} - x^*\| = \|\phi(\tau_0, x_n, v_0) - x^*\| \tag{5.31}$$

$$\leq \varepsilon + d(\phi(\tau_0, x_n, v_0), B_\varepsilon(x^*)) \tag{5.32}$$

$$\leq \varepsilon + F_{B_\varepsilon(x^*)} \tau e^{L\tau_0} \tag{5.33}$$

$$\leq \varepsilon(1 + L\tau_0 e^{L\tau_0}), \tag{5.34}$$

so by (5.24a) and $\tau_0 \leq \tau$ we conclude $x_{n+1} \in S$. Thus, $x_n \in S$ implies $x_{n+1} \in S$, i.e.,

condition (5.7b) holds.

Verification of (5.7c). Let $\delta := \varepsilon(1 + L\tau e^{L\tau})$ and $\bar{n} := \inf\{n : \|x_n - x^*\| \leq \delta\}$. If $\iota_{\mathcal{K}}(x_n) = 0$ and $\|x_n - x^*\| \leq \varepsilon \leq \delta$, then by (5.34) we have $\|x_{n+1} - x^*\| \leq \delta$.

If $\iota_{\mathcal{K}}(x_n) = i \neq 0$, then from (5.25a) and $x_n \in B_{r_i}(x_i)$,

$$e^{\alpha(t_{n+1}-t_n)}\|x_{n+1} - x^*\| \leq e^{\alpha\tau_i}(\|\phi(\tau_i, x_i, v_i) - x^*\| + r_i e^{L\tau_i}) \quad (5.35)$$

$$\leq \|x_i - x^*\| - r_i \leq \|x_n - x^*\|, \quad (5.36)$$

which implies

$$e^{\alpha(t_{n+1}-t_n)}(\|x_{n+1} - x^*\| - \delta) \leq \|x_n - x^*\| - \delta.$$

If $\|x_n - x^*\| \leq \delta$ ($n \geq \bar{n}$), then this inequality ensures $\|x_{n+1} - x^*\| \leq \delta$. If $\|x_n - x^*\| > \delta$ ($n < \bar{n}$), iterating yields

$$\|x_n - x^*\| - \delta \leq e^{-\alpha t_n}(\|x - x^*\| - \delta).$$

Hence for $n \geq \bar{n}$, $\|x_n - x^*\| \leq \delta$, while for $n < \bar{n}$ the excess above δ decays exponentially.

This verifies condition (5.7c).

By Lemma 5.1, $V(x) = \|x - x^*\|$ is an R-CLF with rate α over S and parameter $\delta = \varepsilon(1 + L\tau e^{L\tau})$. Therefore, Theorem 5.1 implies that x^* is practically exponentially stable on S under $\pi_{\mathcal{K}}$, with constants

$$\lambda = \alpha, \quad K = e^{\alpha\tau}(1 + L\tau e^{L\tau}), \quad c = \varepsilon(1 + L\tau e^{L\tau}).$$

□

5.3.2 Existence and Sample Complexity of NCPs

Theorem 5.2 establishes that nonparametric chain policies can guarantee practical exponential stability of a region around an equilibrium point that is appropriately covered by data points from \mathcal{K} . However, it is a priori not clear how many data points are needed to construct such policy, or even whether such a policy exists. The next result provides conditions for existence of Chain Policies as well as a bound on the sample complexity of such policies, i.e., the sizes of the assignment set \mathcal{K} and alphabet \mathcal{A} required to construct such policy.

Theorem 5.3 (Existence and Sample Complexity of Chain Policies). *Consider the control system (5.1) with equilibrium $x^* \in \mathbb{R}^n$, and assume x^* is λ -exponentially stabilizable on \mathbb{R}^n with gain $K > 0$. Let $S = B_R(x^*)$ with $R > 0$, and choose ε s.t. $R > \varepsilon > 0$. Fix any $\alpha \in (0, \lambda)$ and choose*

$$\tau > \frac{\ln K}{\lambda - \alpha}, \quad L := L_{\mathcal{R}^\tau(S)}, \quad \rho := \frac{1 - Ke^{-(\lambda-\alpha)\tau}}{1 + e^{(L+\alpha)\tau}}. \quad (5.37)$$

Then there exists a nonparametric chain policy $\pi_{\mathcal{K}}$ built from a finite assignment set of verification points $\{(x_i, r_i)\}_{i=1}^N \subset S$ and associated controls $\{v_i\}_{i=1}^N$ such that:

- (i) **Practical exponential stability.** *For every $x \in S$, the induced closed loop satisfies*

$$\|\phi(t, x, u_{\mathcal{K},x}) - x^*\| \leq C e^{-\alpha t} \|x - x^*\| + c, \quad \forall t \geq 0,$$

with $C = e^{\alpha\tau}(1 + L\tau e^{L\tau})$ and $c = \varepsilon(1 + L\tau e^{L\tau})$.

(ii) **Sample complexity.** The number N of covering centers and controls satisfies

$$N = O\left(\left(\frac{3}{\rho}\right)^d \log \frac{R}{c}\right).$$

Remark 5.3 (Performance–Complexity Trade-off). The definition of ρ in (5.37) reveals two contrasting regimes. When $\lambda - \alpha$ is close to the lower bound $\ln K/\tau$, the numerator $1 - Ke^{-(\lambda-\alpha)\tau}$ approaches zero, so $\rho \approx 0$. In this regime the guaranteed rate α is nearly as fast as the best attainable λ , but the sample complexity bound $O((3/\rho)^d)$ becomes extremely large. At the other extreme, when $\alpha \ll \lambda$, the term $Ke^{-(\lambda-\alpha)\tau}$ vanishes, and ρ approaches $1/(1 + e^{L\tau})$. In this regime, far fewer samples are needed, but the realized performance α is much slower than the system’s intrinsic rate λ .

Thus, ρ quantifies the fundamental trade-off: choosing α close to λ yields strong performance at the cost of high sample complexity, while smaller α reduces sample requirements but sacrifices convergence speed.

Proof. Since x^* is λ -exponentially stabilizable on S with constant $K > 0$, for each grid center $x_i \in S$ we can select a constant control v_i such that

$$\|\phi(t, x_i, v_i) - x^*\| \leq Ke^{-\lambda t} \|x_i - x^*\|, \quad \forall t \geq 0. \quad (5.38)$$

We construct the nonparametric chain policy $\pi_{\mathcal{D}}$ by covering S with finitely many balls $B_{r_i}(x_i)$ and assigning to each x_i the control v_i above. In our construction, we will assume $\|\cdot\|$ is the infinity norm, i.e., $\|x\| = \max_{j=1,\dots,n} |x_j|$.

Step 1: Choice of radii. Fix $\alpha \in (0, \lambda)$ and $\tau > \frac{\ln K}{\lambda - \alpha}$. Let $L := L_{\mathcal{R}\tau(S)}$ and set

$$\rho := \frac{1 - Ke^{-(\lambda - \alpha)\tau}}{1 + e^{(L + \alpha)\tau}} > 0, \quad r_i := \rho \|x_i - x^*\|.$$

By construction $r_i > 0$. From (5.38), for $t = \tau$ we have

$$Ke^{-(\lambda - \alpha)\tau} \|x_i - x^*\| + r_i e^{(L + \alpha)\tau} = \|x_i - x^*\| - r_i,$$

which implies the verification condition

$$\min_{t \in (0, \tau]} e^{\alpha t} \left(\|\phi(t, x_i, v_i) - x^*\| + r_i e^{Lt} \right) \leq e^{\alpha \tau} \left(\|\phi(\tau, x_i, v_i) - x^*\| + r_i e^{L\tau} \right) \leq \|x_i - x^*\| - r_i.$$

Hence each (x_i, r_i) satisfies (5.25a).

Step 2: Covering number bound. We cover the annular region $B_R(x^*) \setminus B_\varepsilon(x^*)$ by n concentric annuli A_1, \dots, A_n of thickness $R_i = 3^{i-1}\varepsilon$. Since the total radial width is $R - \varepsilon$,

$$R - \varepsilon \leq \sum_{i=1}^n 3^{i-1}\varepsilon = \varepsilon \frac{3^n - 1}{2} \quad \Rightarrow \quad n = \left\lceil \log_3 \left(\frac{2R}{\varepsilon} - 1 \right) \right\rceil = O\left(\log \frac{R}{\varepsilon} \right).$$

Each A_i is initially partitioned into $3^d - 1$ hypercubes of side R_i . For any $x \in A_i$, $\|x - x^*\| \geq 3^{i-1}\varepsilon$ and $r_i \leq \rho \|x - x^*\|$. We refine each hypercube by successive splits into 3^d subcubes until the side length is at most $\rho \|x - x^*\|$. This requires

$$3^m \geq \frac{1}{\rho} \quad \Rightarrow \quad m = \lceil \log_3(1/\rho) \rceil$$

splits, producing at most

$$N_{\text{annulus}} = (3^d - 1) 3^{dm} = O\left(\left(\frac{3}{\rho}\right)^d\right)$$

points per annulus.

Step 3: Total number of points. Multiplying by the number of annuli,

$$N = n \cdot N_{\text{annulus}} = O\left(\left(\frac{3}{\rho}\right)^d \log \frac{R}{\varepsilon}\right).$$

Finally, since $c = \varepsilon(1 + L\tau e^{L\tau})$, this yields the claimed sample complexity bound. \square

5.3.3 Incremental Learning of NCPs

The performance–complexity trade-off discussed in Remark 5.3, together with the existence and sample complexity result of Theorem 5.3, suggests a practical methodology for progressively improving performance. By actively sampling trajectories more finely and refining the covering set \mathcal{K} , one can construct NCPs that certify larger rates α (by reducing the effective radius r), thereby reducing the gap between realized performance and the best attainable rate λ . In other words, performance can be systematically enhanced by enriching the assignment set with additional verified points, which readily enables incremental learning.

Beyond improving rates, another key feature of NCPs is their ability to incrementally *expand the certified region*. The next result formalizes this incremental learning property: previously verified assignments can be reused together with new ones to

grow the domain over which stability is guaranteed.

Theorem 5.4 (Incremental Learning of \mathcal{K}). *Consider an equilibrium point $x^* \in \mathbb{R}^n$ of (5.1) and a set $S \subseteq \mathbb{R}^n$ satisfying $x^* \in \text{int}(S)$. Let $\pi_{\mathcal{K}}$ be a nonparametric policy with assignment set $\mathcal{K} = \{(x_i, r_i, v_i)\}_{i=1}^N$ and default control $v_0 \in \mathcal{A}$ satisfying properties (i)–(iii) of Theorem 5.2 with parameters $\alpha, \delta, \tau, L, \varepsilon$. Take $x_j \in \mathbb{R}^n \setminus S$, $r_j > 0$, and $v_j \in \mathcal{U}^{[0, \tau_j]}$ s.t. $B_{r_j}(x_j) \cup S = \emptyset$. Define the enlarged set $S' := S \cup B_{r_j}(x_j)$, and let $L_j = L_{\mathcal{R}^{\tau_j}(B_{r_j}(x_j))}$, and $L' = \max\{L_j, L\}$.*

Whenever the following conditions are satisfied:

(1) **Feasibility of (x_j, r_j, v_j)** : The 3-tuple (x_j, r_j, v_j) with $v_j : [0, \tau_j] \rightarrow U$ satisfy

$$\text{sd}(\phi(\tau_j, x_j, v_j), S) + r_j e^{L_j \tau_j} \leq 0. \quad (5.39)$$

(2) **Either of the following holds**:

(a) **Direct verification at (x_j, r_j, v_j)** : The tuple satisfies decrease condition:

$$e^{\alpha \tau_j} \left(\|\phi(\tau_j, x_j, v_j) - x^*\| + r_j e^{L \tau_j} \right) \leq \|x_j - x^*\| - r_j. \quad (5.40)$$

Set $\alpha' = \alpha$, $\tau' = \max\{\tau, \tau_j\}$, $\delta' = \delta$.

(b) **Bootstrapping**: There is $\hat{\mathcal{K}} \subseteq \mathcal{K}$, such that

(i) $B_{r_j e^{L_j \tau_j}}(\phi(\tau_j, x_j, v_j)) \subseteq \text{Supp}(\hat{\mathcal{K}})$

(ii) There is $\alpha' < \alpha$ such that

$$\max_{(x_i, r_i, v_i) \in \hat{\mathcal{K}}} \frac{e^{-(\alpha-\alpha')\tau_i} \|x_i - x^*\| + r_i}{e^{-\alpha\tau_j} \|x_j - x^*\| - r_j} \leq 1.$$

Set $\tau' = \tau + \tau_j$, $\delta' = \delta$.

Then the augmented assignment set $\mathcal{K}' := \mathcal{K} \cup \{(x_j, r_j, v_j)\}$ and the default control v_0 induce a policy $\pi_{\mathcal{K}'}$ that practically exponentially stabilizes x^* over $S' = S \cup B_{r_j}(x_j)$ with

$$\lambda' = \alpha', \quad K' = e^{\alpha'\tau'}(1 + L'\tau e^{L'\tau'}), \quad c' = \varepsilon(1 + L\tau e^{L\tau}).$$

Proof. We first note that since \mathcal{K} and v_0 satisfy Theorem 5.1 (i)–(iii) on S , and $S \cap B_{r_j}(x_j) = \emptyset$, any initial state $x \in S$ under $\pi_{\mathcal{K}'}$ will only trigger assignments from \mathcal{K} or the default control v_0 . By Theorem 5.2, this ensures that for every $x \in S$ there exists a sequence of times $\{t_n\}$ satisfying the conditions of Lemma 5.1, along which (5.7) is satisfied for $V(x) = \|x - x^*\|$. Hence, whenever $x \in S \subset S'$, the trajectory $\phi(t, x, u_{\mathcal{K}',x})$ is practically exponentially stable with parameters $\lambda = \alpha$, $K = e^{\alpha\tau}(1 + L\tau e^{L\tau})$, and $c = \varepsilon(1 + L\tau e^{L\tau})$. Moreover, since $\alpha \leq \alpha'$, $\tau' \geq \tau$, and $L' \geq L$, the same trajectory also satisfies the practical exponential stability bound with the updated constants λ' , K' , and c' as stated in the theorem.

We next consider the case $x \in S' \setminus S$. Since $S \cap B_{r_j}(x_j) = \emptyset$, it follows that $\iota_{\mathcal{K}'}(x) = j$, so the first control applied by $\pi_{\mathcal{K}'}$ is v_j . We will use the sequence (5.23) induced by $\pi_{\mathcal{K}'}$ to build a sequence of times $\{t_n\}_{n \in \mathbb{N}}$ satisfying the properties in (5.7) of Lemma 5.1 for the control signal $u_{\mathcal{K}',x}$. The result then follows from Theorem 5.1. We will choose δ'

and α' in (5.7) later on. Recall that $\pi_{\mathcal{K}}$ satisfies (5.7) for some α and $\delta = \varepsilon(1 + L\tau e^{L\tau})$.

To simplify notation, we use u to refer to $u_{\mathcal{K}',x}$.

We choose $x_0 = x$ and $t_0 = 0$. We will select t_1 differently, depending on which clause of condition (2) in the theorem's hypothesis hold.

Case (2a). We choose $t_1 = \tau_j$, and accordingly $x_1 = \phi(\tau_j, x, v_j) = \phi(t_1, x, u)$. By condition (1) of the theorem, we have

$$\text{sd}(\phi(\tau_j, x, v_j), S) \leq \text{sd}(\phi(\tau_j, x_j, v_j), S) + r_j e^{L_j \tau_j} \leq 0,$$

which implies $\phi(\tau_j, x, v_j) \in S$. This ensures $\phi(t_1, x, u) = \phi(\tau_j, x, v_j) \in S \subset S'$, i.e., (5.7b) for the set S' . It further follows from (5.40) by a similar argument to Theorem 5.2, that

$$\|\phi(t_1, x, v_j) - x^*\| \leq e^{-\alpha t_1} \|x - x^*\|,$$

so (5.7c) holds for α and any value of δ .

Case (2b). As in case (2a) after choosing v_j , $\phi(\tau_j, x, v_j) \in S$. In fact, by (2b.i), $\phi(\tau_j, x, v_j) \in \text{Supp}(\hat{\mathcal{K}}) \subseteq S$. Let $(x_i, r_i, v_i) \in \hat{\mathcal{K}}$ s.t. $i = \iota_{\mathcal{K}}(\phi(\tau_j, x, v_j))$. We choose

$$t_1 := \tau_j + \tau_i, \tag{5.41}$$

$$x_1 := \phi(\tau_j + \tau_i, x, v_j v_i) = \phi(\tau_i, \phi(\tau_j, x, v_j), v_i). \tag{5.42}$$

Since (5.25) holds for $(x_i, r_i, v_i) \in \mathcal{K}$, $x_1 \in S$, and therefore we have $\phi(t_1, x, u) = x_1 \in S \subset S'$; hence (5.7b) holds for t_1 and S' .

It remains to ensure the decrease in (5.7c) at t_1 with some $\alpha' < \alpha$ and $\delta' = \delta$. We thus consider

$$e^{\alpha' t_1} \|\phi(t_1, x, u) - x^*\| = \quad (5.43)$$

$$= e^{\alpha'(\tau_j + \tau_i)} \|\phi(\tau_j + \tau_i, x, v_j v_i) - x^*\| \quad (5.44)$$

$$= e^{\alpha'(\tau_j + \tau_i)} \|\phi(\tau_i, \phi(\tau_j, x, v_j), v_i) - x^*\| \quad (5.45)$$

$$\leq e^{\alpha'(\tau_j + \tau_i)} e^{-\alpha \tau_i} \|\phi(\tau_j, x, v_j) - x^*\| \quad (5.46)$$

$$= \frac{e^{-(\alpha - \alpha')\tau_i}}{e^{-\alpha' \tau_j}} \left(\frac{\|\phi(\tau_j, x, v_j) - x^*\|}{\|x - x^*\|} \right) \|x - x^*\| \quad (5.47)$$

$$\leq \frac{e^{-(\alpha - \alpha')\tau_i}}{e^{-\alpha' \tau_j}} \left(\frac{\|x_i - x^*\| + r_i}{\|x_j - x^*\| - r_j} \right) \|x - x^*\| \quad (5.48)$$

$$\leq \|x - x^*\|, \quad (5.49)$$

where step one follows from definition of t_1 and u , two from the group property of ϕ , three from (5.25a) on (x_i, r_i, v_i) , and the final step from the bootstrapping condition (2b.ii).

In both cases, after t_1 the subsequent points are generated by \mathcal{K} , so the sequence $\{t_n\}$ continues to satisfy (5.7b)–(5.7c) by Theorem 5.2. Thus $V(x) = \|x - x^*\|$ is an R-CLF on S' , and by Theorem 5.1, $\pi_{\mathcal{K}'}$ renders x^* practically exponentially stable on S' with the claimed constants. \square

5.4 Numerical Experiments

Using the sufficiency of NCPs derived in Theorem 5.2 along with the grid construction of Theorem 5.3, we next introduce an algorithmic methodology to design NCPs to

stabilize a given region. The algorithmic flow is as follows:

- i. Given a region $S \subset \mathbb{R}^n$, select a desired convergence rate $\alpha > 0$, select a τ_{\max} to upper-bound τ_i , and determine the one-sided Lipschitz constant L for the underlying dynamics across $\mathcal{R}^{\tau_{\max}}(S)$ ¹.
- ii. Create a grid of points and radii $G = \{(x_i, r_i)\}$ covering the region S with increasingly large radii per annulus according to Theorem 5.3.
- iii. For each x_i , derive controls v_i for τ_{\max} time (for example, using sampling methods akin to Model-Predictive Path Integrals (MPPI) [127, 128]).
- iv. For each (x_i, r_i, v_i) , if condition (5.25a) is satisfied for some $\tau_i \leq \tau_{\max}$, store the largest τ_i that achieves equality.
- v. Otherwise, split the ball $B_{r_i}(g_i)$ into 3^d smaller balls.
- vi. Repeat until a control is found for α -exponential decrease with slack of at least each ball's radius, or until a pre-defined maximum number of splits is achieved.
- vii. For balls which still fail the conditions, apply condition 2b of Theorem 5.4 to leverage previously derived controls.
- viii. Trim down the verified region only to those trajectories which satisfy (5.25b), save verified $\alpha_i \geq \alpha$ for each cell.

¹Achieved through estimating the reachable tube by simulating samples along the boundary and adding precision-correction terms, and then sampling points in that region for OSL while adding precision-correction terms again.

With this algorithmic method, we present a number of case studies in different classic control stabilization problems, which will demonstrate useful features of NCPs.

5.4.1 Unicycle

Consider the unicycle model moving in the plane, parametrized by its x position, y position, and angle of vehicle θ with respect to the x -axis, with two control inputs - velocity v and angular velocity ω . The dynamics are given by

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} v \cos(\theta) \\ v \sin(\theta) \\ \omega \end{bmatrix}.$$

We bound $u \in [0, 1]$ and $\omega \in [-1, 1]$, and run the method to derive NCPs for two different norms, being $V_1 = \max\{|x|, |y|, |\theta|\}$ and $V_2 = \sqrt{x^2 + y^2 + 0.01\theta^2}$. For either choice of norm (simpler or tied to the classic reward function), the method quickly stabilized the entire region of $(x, y) \in [-20, 20]^2, \theta \in (-\pi, \pi]$, see Figure 5.1.

To demonstrate the incremental growth capabilities of NCP, we do two stages of learning. After learning a control for the previous region, i.e., $(x, y, \theta) \in [-20, 20]^2 \times (-\pi, \pi]$, we expand the state space to include values in the region $(x, y, \theta) \in [-20, 20] \times [20, 25] \times (-\pi, \pi]$. Trajectories fragments starting in the formerly verified region retain the same behavior, while the new behavior (for initial values in the new region) are depicted in Figure 5.2.

5.4.2 Inverted Pendulum

We next analyze the utility of the NCPs on the inverted pendulum. The system consists of a mass m attached at the end of a rigid pendulum of length l , pivoting freely about a fixed point. The dynamics are governed by the torque around the pivot due to both gravity and an external control input. Denoting by θ the angle of the pendulum measured from the vertical (with $\theta = 0$ corresponding to the inverted equilibrium), the equation of motion is given by

$$ml^2\ddot{\theta}(t) = mgl \sin(\theta(t)) + u(t), \quad (5.50)$$

where g is the gravitational acceleration, and $u(t)$ is the external control torque applied at the pivot.

Figure 5.3 demonstrates the refinement capabilities of NCPs, such that by adding data (simulated by splitting all balls once), the rate of convergence achieved is significantly increased.

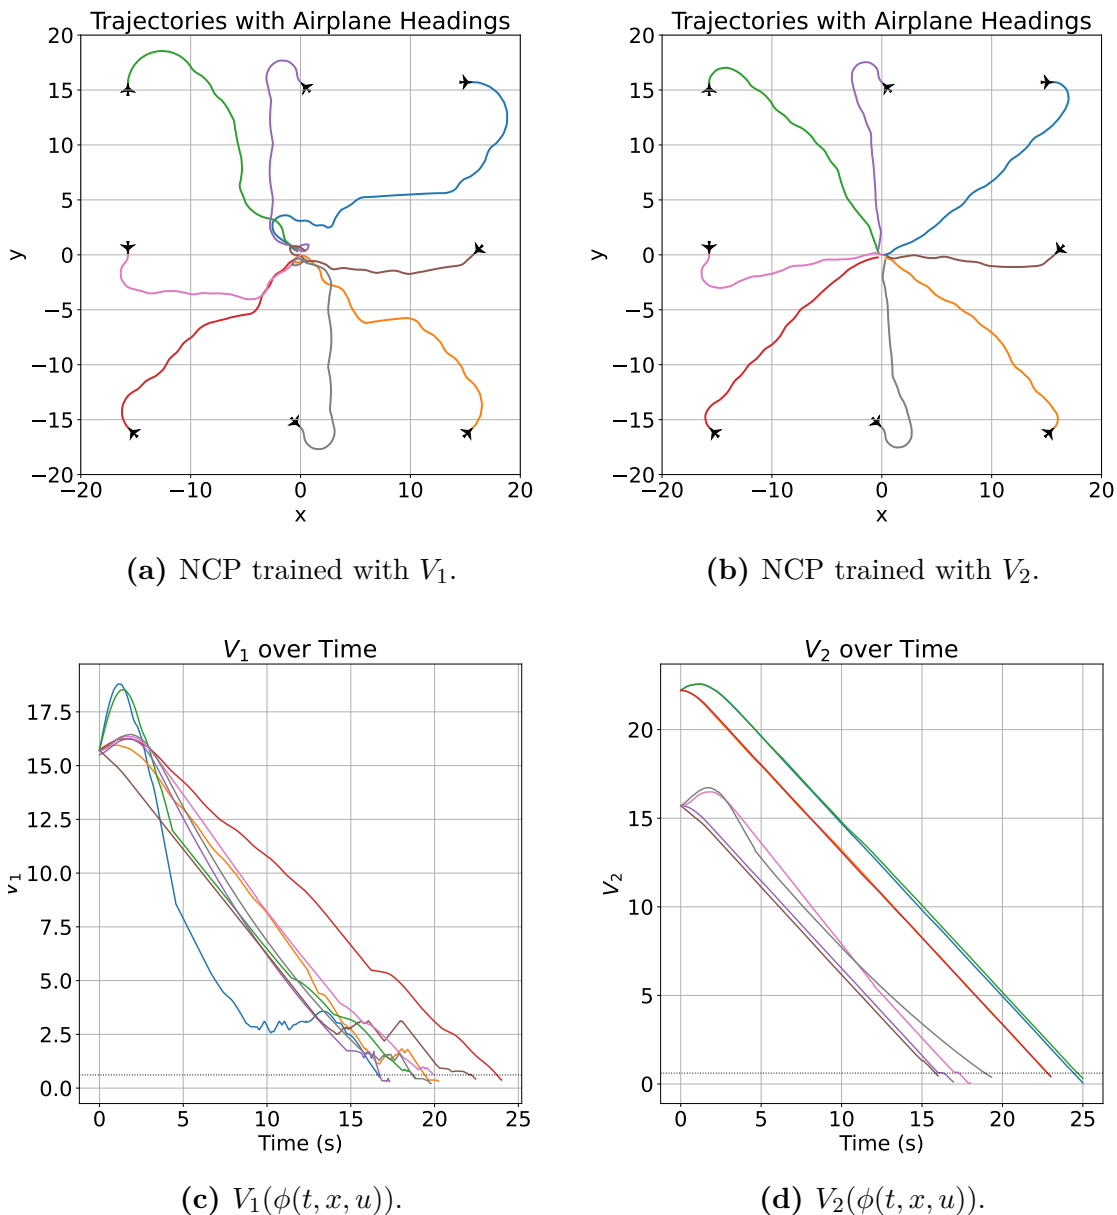


Figure 5.1: Trajectories of Unicycle NCP. Phase plots of (x, y) for eight evenly distributed points. The black icons depict the initial facing of the unicycle. Plot (a) contains trajectories from NCP trained to minimize V_1 , which results in sharp turns, while (b) is trained to minimize V_2 , which results in softer turns and smoother overall behavior. Plots (c) and (d) show the development of V_1 and V_2 over time respectively. Both converge exponentially to the equilibrium, with at least $\alpha = 0.01$. We have $\tau_{\max} = 5, \varepsilon = 0.01, L = 1$ and $c = \varepsilon(1 + L\tau_{\max}e^{L\tau_{\max}}) \simeq 0.613$, represented by the dotted line.




Figure 5.2: Incremental Learning of Unicycle Policy. Extending the state space from the previously learned region in the y -direction. Subfigure (a) contains the phase plot before learning, while subfigure (b) contains the phase plot after. The new region is learned without forgetting, such that parts of the trajectory in the old region use previously designed controls.

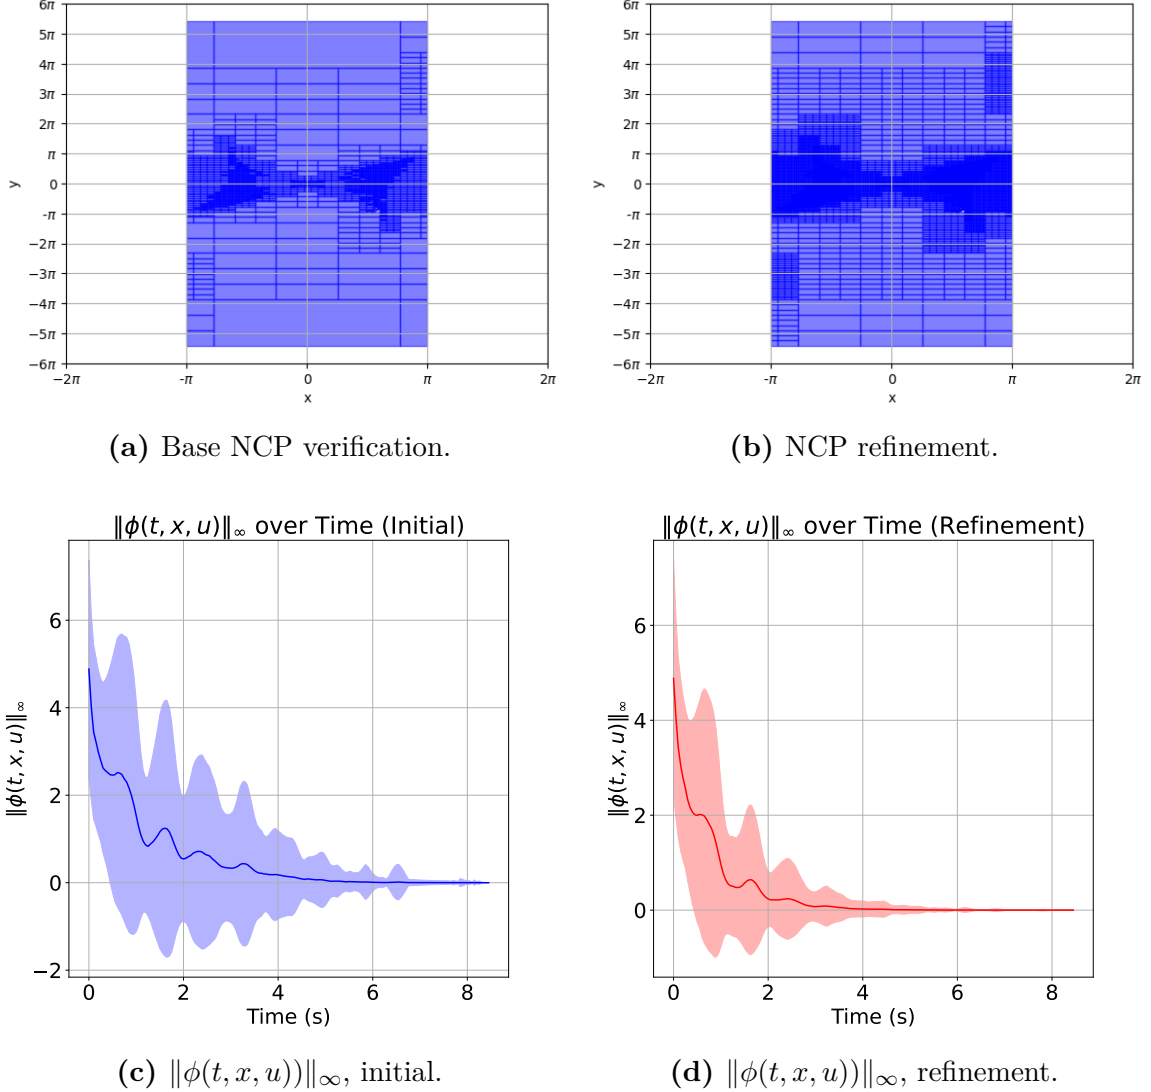


Figure 5.3: Additional Data Refinement Facilitates Improved NCP Performance. Plot (a) contains the balls used to verify the region $(\theta, \dot{\theta}) \in (-\pi, \pi) \times [-5\pi, 5\pi]$ for the inverted pendulum. Plot (b) is a refinement of plot (a), wherein all balls were split once more and re-verified. The minimum verified rate of convergence for trajectories α goes from 0.003 to 0.0145, and the average verified α goes from 1.815 to 3.149. Plot (c) demonstrates the average norm over time of 400 sample trajectories under each schema. We have $\tau_{\max} = 1.5$, $\varepsilon = 0.01$, $L = 1$, and $c = \varepsilon(1 + L\tau_{\max}e^{L\tau_{\max}}) \simeq 0.072$.

Chapter 6

Conclusion

What we begin in life may not be completed in our lifetime. That does not make it meaningless. It makes it honorable.

– LORD RABBI JONATHAN SACKS

Stabilization analysis has long relied on the construction of Lyapunov functions. Unfortunately, they are difficult to design, with re-usability between systems being highly limited. The end result is that there are many systems which are stable—including some which we empirically or theoretically know must be stable—cannot be verified as having stability properties. Furthermore, Lyapunov functions lack robustness; the expansion of the state space may require the complete re-design of these functions. This feels intuitively wrong: what we learn about part of a system should not simply be erased when beginning to look at the system in its entirety.

This problem is exacerbated in the design of controls. Having firm theoretical guarantees not only relies on the construction of Lyapunov functions, but also simultaneous construction of controls which render the selected function valid. This also requires knowledge of the parametrization of the system, which is a significant ask from an information standpoint, and results in further loss of robustness. This problem in particular inspired the creation of modern AI and RL methods for control, which while are non-parametric, come with their own drawbacks. First, there are no guarantees of performance and stabilization, as such schemes are entirely empirical. Furthermore, the real-time execution of such systems presents a latency bottleneck, as controls must

be recomputed at each state.

While not part of this dissertation, we recently investigated similar issues beyond dynamical systems and control in the realms of LLMs and Computational Biology. These other studies facilitated a deeper and more abstract understanding of our contributions in this thesis. In the up-and-coming field of Large Language Models (LLMs), the problem of latency and computational cost is a significant issue in current training and evaluation metrics, which almost universally require generations on a test set to be created and then evaluated. Beyond this, existing metrics tend to be binary, lacking robustness to linguistic structure and testing methodology (such as true or false versus multiple choice, free response, numeric extraction, etc.). Similar to Lyapunov functions, such testing methodologies also require explicit design, and are ill-suited to analyzing performance in open-ended subjects, such as language.

In Computational Biology—as in LLMs—state-of-the-art analytical methods generally rely on binarization, here via thresholds to determine what constitutes ‘biological regions of significance’. Beyond the lack of formalism (and thus apparent arbitrariness) in selection of such thresholds, they inherently introduce the potential of entirely missing biological phenomena which fall partially or mostly beneath their detection thresholds.

In all aforementioned disciplines, there is a fundamental disconnect between the researcher—who had particular targets in mind—and the systems themselves, with their dispassionate data. Our direction of study entails maximal reduction in the conditions set upon the target of analysis, and letting the data speak for itself, aiding

us in disentangling the underlying system and its properties.

In each case, our solution has been slightly different. In dynamical systems, rather than painstakingly design a Lyapunov function to fit the system into the neat box of invariance, solutions are allowed to leave sublevel sets and recur. In a stable system, it is necessary and sufficient for any norm to have this property of τ -recurrence, and thus the solutions themselves can be used to provide stability guarantees without referring to the functional representation. This freedom from parametrization facilitates a significant step in control theory, wherein theoretical guarantees are provided by construction in a framework that offers incredibly fast real-time execution and continual expansion capabilities.

Our solution for LLMs [107], entailed the introduction of the metric of Complementarity, which enables a paradigm of robust, cross-system comparisons that are broadly applicable and free of binarization. By measuring fluency, Complementarity facilitating an examination of true understanding of particular domains rather than the particular style, format, or biased testing seen in many state-of-the-art AI metrics. This freedom from naive binarization forms the backbone of MOTiF - our solution in computational biology [111] - wherein a broad spectrum of different thresholds is utilized to achieve a panoramic view of the genomic structural and functional landscape, and thus a deeper explanation of the underlying biological effects.

Science, at its best, is a discipline of truth. Yet in practice, even the most rigorous inquiry is shaped—subtly and profoundly—by human choice. The selection of tools, the framing of problems, the conditions we impose on proofs—these are not neutral

acts. They reflect assumptions, traditions, and implicit values.

Lyapunov theory has stood for more than a century as a standard by which stability is judged. But like all standards, it is not without its limits. The insistence on invariance has long defined the boundaries of what counts as stability, often at the cost of flexibility or accessibility. This thesis has sought, in some small measure, to widen that horizon: to show that recurrence may serve where invariance has been thought essential, and that stability can be understood not only through what never leaves, but also through what leaves and returns.

The contrast between Lyapunov theory and our own is that of constraint: both methods offer necessary and sufficient conditions. However, whereas Lyapunov's necessity is only in the existence of a **single function** satisfying the strict conditions, ours is a broader statement, that of **any norm** satisfying the simpler, less-restrictive conditions. We therefore believe that RLFs present a more holistic picture of the nature of system stability.

The responsibility of science is not to shape systems in our image, nor to bind them within conditions that suit our convenience. The task is more difficult, and more sacred: to let the system speak, and to listen—to discover forms of order that are already present, even if they do not match the patterns we expected.

If recurrence offers another path, more adaptable to modern systems and more accessible in practice, then this work has served its purpose: to add a small stone to the long road by which science moves closer to understanding. Its purpose is not an end, but a beginning—an invitation to think differently, to question inherited

boundaries, and to seek stability where once it seemed out of reach.

Bibliographic references

1. Aeyels, D. & Peuteman, J. A new asymptotic stability criterion for nonlinear time-variant differential equations. *IEEE Transactions on automatic control* **43**, 968–971 (1998).
2. Ahmadi, A. A. *Non-monotonic Lyapunov functions for stability of nonlinear and switched systems: theory and computation* PhD thesis (Massachusetts Institute of Technology, 2008).
3. Ahmadi, A. A. & Majumdar, A. Dsos and sdsos optimization: more tractable alternatives to sum of squares and semidefinite optimization. *SIAM Journal on Applied Algebra and Geometry* **3**, 193–230 (2019).
4. Ahmadi, A. A. & Parrilo, P. A. *On higher order derivatives of Lyapunov functions* in *Proceedings of the 2011 American Control Conference* (2011), 1313–1314. doi:[10.1109/ACC.2011.5991573](https://doi.org/10.1109/ACC.2011.5991573).
5. Alessio, A. & Bemporad, A. in *Nonlinear Model Predictive Control: Towards New Challenging Applications* 345–369 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009). doi:[10.1007/978-3-642-01094-1_29](https://doi.org/10.1007/978-3-642-01094-1_29).
6. Allgöwer, F. & Zheng, A. *Nonlinear Model Predictive Control* (Birkhäuser Basel, 2000).
7. Alongi, J. M. & Nelson, G. S. *Recurrence and topology* (American Mathematical Soc., 2007).
8. Astolfi, A. & Praly, L. *A lasalle version of matrosov theorem* in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)* (2011), 320–324. doi:[10.1109/CDC.2011.6161174](https://doi.org/10.1109/CDC.2011.6161174).
9. B. Ghanmi, N. H. T. & Hammami, M. Growth conditions for exponential stability of time-varying perturbed systems. *International Journal of Control* **86**, 1086–1097. doi:[10.1080/00207179.2013.774464](https://doi.org/10.1080/00207179.2013.774464). eprint: <https://doi.org/10.1080/00207179.2013.774464> (2013).

Bibliographic references

10. Bayat, F., Johansen, T. A. & Jalali, A. A. *Flexible piecewise function evaluation methods with application to explicit model predictive control* in *2011 IEEE International Conference on Mechatronics* (2011), 218–223. doi:[10.1109/ICMECH.2011.5971284](https://doi.org/10.1109/ICMECH.2011.5971284).
11. Bayat, F., Johansen, T. A. & Jalali, A. A. Using hash tables to manage the time-storage complexity in a point location problem: application to explicit model predictive control. *Automatica* **47**, 571–577. doi:<https://doi.org/10.1016/j.automatica.2011.01.009> (2011).
12. Bemporad, A., Morari, M., Dua, V. & Pistikopoulos, E. *The explicit solution of model predictive control via multiparametric quadratic programming* in *Proceedings of the 2000 American Control Conference. ACC (IEEE Cat. No.00CH36334)* **2** (2000), 872–876 vol.2. doi:[10.1109/ACC.2000.876624](https://doi.org/10.1109/ACC.2000.876624).
13. Bemporad, A. Explicit model predictive control. *Encyclopedia of Systems and Control* (2015).
14. Bemporad, A. & Filippi, C. Suboptimal explicit receding horizon control via approximate multiparametric quadratic programming. *Journal of Optimization Theory and Applications* **117**, 9–38. doi:[10.1023/A:1023696221899](https://doi.org/10.1023/A:1023696221899) (2003).
15. Bemporad, A., Morari, M., Dua, V. & Pistikopoulos, E. N. The explicit linear quadratic regulator for constrained systems. *Automatica* **38**, 3–20. doi:[https://doi.org/10.1016/S0005-1098\(01\)00174-1](https://doi.org/10.1016/S0005-1098(01)00174-1) (2002).
16. Bemporad, A., Oliveri, A., Poggi, T. & Storace, M. Ultra-fast stabilizing model predictive control via canonical piecewise affine approximations. *IEEE Transactions on Automatic Control* **56**, 2883–2897. doi:[10.1109/TAC.2011.2141410](https://doi.org/10.1109/TAC.2011.2141410) (2011).
17. Benjemaa, M., Gouadri, W. & Hammami, M. A. *New results on the uniform exponential stability of non-autonomous perturbed dynamical systems* 2020. arXiv: [1905.10570](https://arxiv.org/abs/1905.10570) [[math.DS](https://arxiv.org/abs/1905.10570)].
18. Berberich, J., Koch, A., Scherer, C. W. & Allgöwer, F. *Robust data-driven state-feedback design* in *2020 American Control Conference (ACC)* (2020), 1532–1538.

Bibliographic references

19. Berberich, J., Köhler, J., Müller, M. A. & Allgöwer, F. Data-driven model predictive control with stability and robustness guarantees. *IEEE transactions on automatic control* **66**, 1702–1717 (2020).
20. Berberich, J., Köhler, J., Müller, M. A. & Allgöwer, F. Linear tracking mpc for nonlinear systems—part i: the model-based case. *IEEE Transactions on Automatic Control* **67**, 4390–4405. doi:[10.1109/TAC.2022.3166872](https://doi.org/10.1109/TAC.2022.3166872) (2022).
21. Borrelli, F., Bemporad, A. & Morari, M. *Predictive Control for Linear and Hybrid Systems* (Cambridge University Press, 2017).
22. Butz, A. Higher order derivatives of Liapunov functions. *IEEE Transactions on automatic control* **14**, 111–112 (1969).
23. Canale, M., Fagiano, L. & Milanese, M. Set membership approximation theory for fast implementation of model predictive control laws. *Automatica* **45**, 45–54. doi:<https://doi.org/10.1016/j.automat.2008.06.015> (2009).
24. Chen, G. & Yang, Y. New stability conditions for a class of linear time-varying systems. *Automatica* **71**, 342–347. doi:<https://doi.org/10.1016/j.automat.2016.05.005> (2016).
25. Chen, G. & Yang, Y. Relaxed conditions for the input-to-state stability of switched nonlinear time-varying systems. *IEEE Transactions on Automatic Control* **62**, 4706–4712. doi:[10.1109/TAC.2016.2625979](https://doi.org/10.1109/TAC.2016.2625979) (2017).
26. Chen, S., Saulnier, K., Atanasov, N., Lee, D. D., Kumar, V., Pappas, G. J. & Morari, M. *Approximating explicit model predictive control using constrained neural networks* in *2018 Annual American Control Conference (ACC)* (2018), 1520–1527. doi:[10.23919/ACC.2018.8431275](https://doi.org/10.23919/ACC.2018.8431275).
27. Chen, X. & Hazan, E. *Black-box control for linear dynamical systems* in *Conference on Learning Theory* (2021), 1114–1143.
28. Chen, Y. & Vaidya, U. *Sample complexity for nonlinear stochastic dynamics* in *2019 American Control Conference (ACC)* (2019), 3526–3531.
29. Christophersen, F. J., Kvasnica, M., Jones, C. N. & Morari, M. *Efficient evaluation of piecewise control laws defined over a large number of polyhedra* in *2007*

Bibliographic references

- European Control Conference (ECC)* (2007), 2360–2367. doi:[10.23919/ECC.2007.7068280](https://doi.org/10.23919/ECC.2007.7068280).
30. Coddington, E. A. *An Introduction to Ordinary Differential Equations* (Dover Publications, New York, 1989).
 31. Colonius, F. & Kawan, C. Invariance entropy for control systems. *SIAM J. Control Optim.* **48**, 1701–1721 (2009).
 32. Colonius, F. Minimal bit rates and entropy for exponential stabilization. *SIAM Journal on Control and Optimization* **50**, 2988–3010 (2012).
 33. Colonius, F. & Hamzi, B. Entropy for practical stabilization. *SIAM Journal on Control and Optimization* **59**, 2195–2222 (2021).
 34. Colonius, F. & Kawan, C. Invariance entropy for control systems. *SIAM Journal on Control and Optimization* **48**, 1701–1721. doi:[10.1137/080713902](https://doi.org/10.1137/080713902). eprint: <https://doi.org/10.1137/080713902> (2009).
 35. Coron, J.-M. & Rosier, L. A relation between continuous time-varying and discontinuous feedback stabilization. *J. Math. Syst., Estimation, Control* **4**, 67–84 (1994).
 36. Coulson, J., Lygeros, J. & Dörfler, F. *Data-enabled predictive control: in the shallows of the deep* in *2019 18th European control conference (ECC)* (2019), 307–312.
 37. Dai, T. & Sznaier, M. A semi-algebraic optimization approach to data-driven control of continuous-time nonlinear systems. *IEEE Control Systems Letters* **5**, 487–492 (2020).
 38. De Persis, C. & Tesi, P. Formulas for data-driven control: stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control* **65**, 909–924 (2019).
 39. De Persis, C. & Tesi, P. Low-complexity learning of linear quadratic regulators from noisy data. *Automatica* **128**, 109548 (2021).

Bibliographic references

40. Degner, M., Soloperto, R., Zeilinger, M. N., Lygeros, J. & Köhler, J. *Adaptive Economic Model Predictive Control for linear systems with performance guarantees* 2024. arXiv: [2403.18398](https://arxiv.org/abs/2403.18398) [[eess.SY](#)].
41. Findeisen, R., Allgöwer, F. & Biegler, L. *Assessment and Future Directions of Nonlinear Model Predictive Control* (2006).
42. García, C. E., Prett, D. M. & Morari, M. Model predictive control: theory and practice, a survey. *Automatica* **25**, 335–348. doi:[https://doi.org/10.1016/0005-1098\(89\)90002-2](https://doi.org/10.1016/0005-1098(89)90002-2) (1989).
43. Gersnoviez, A., Brox, M. & Baturone, I. High-speed and low-cost implementation of explicit model predictive controllers. *IEEE Transactions on Control Systems Technology* **27**, 647–662. doi:[10.1109/TCST.2017.2775187](https://doi.org/10.1109/TCST.2017.2775187) (2019).
44. Giesl, P. & Hafstein, S. Review on computational methods for Lyapunov functions. *Discrete and Continuous Dynamical Systems - B* **20**, 2291–2331 (2015).
45. Goebel, R., Teel, A., Hu, T. & Lin, Z. Conjugate convex Lyapunov functions for dual linear differential inclusions. *IEEE Transactions on Automatic Control* **51**, 661–666. doi:[10.1109/TAC.2006.872764](https://doi.org/10.1109/TAC.2006.872764) (2006).
46. Gunderson, R. A comparison lemma for higher order trajectory derivatives. *Proceedings of the American Mathematical Society* **27**, 543–548 (1971).
47. Guo, M., De Persis, C. & Tesi, P. Data-driven stabilization of nonlinear polynomial systems with noisy data. *IEEE Transactions on Automatic Control* **67**, 4210–4217 (2021).
48. Hamzi, B. & Krener, A. J. in *New Trends in Nonlinear Dynamics and Control and Their Applications* 37–48 (Springer, 2004).
49. Hara, T., Yoneyama, T. & Okazaki, Y. On the boundedness of solutions of perturbed linear systems. *Journal of Mathematical Analysis and Applications* **83**, 188–208. doi:[https://doi.org/10.1016/0022-247X\(81\)90256-0](https://doi.org/10.1016/0022-247X(81)90256-0) (1981).
50. Hassan, M. & C.Storey. Numerical determination of domains of attraction for electrical power systems using the method of zubov. *International Journal of Control* **34**, 371–381. doi:[10.1080/00207178108922536](https://doi.org/10.1080/00207178108922536) (1981).

Bibliographic references

51. Holaza, J., Takács, B., Kvasnica, M. & Di Cairano, S. Nearly optimal simple explicit mpc controllers with stability and feasibility guarantees. *Optimal Control Applications and Methods* **36**, 667–684. doi:<https://doi.org/10.1002/oca.2131>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/oca.2131> (2015).
52. Hovland, S., Gravdahl, J. & Willcox, K. Explicit model predictive control for large-scale systems via model reduction. *Journal of Guidance Control and Dynamics - J GUID CONTROL DYNAM* **31**. doi:[10.2514/1.33079](https://doi.org/10.2514/1.33079) (2008).
53. Hsu, T.-W. & Tsukamoto, H. Statistical guarantees in data-driven nonlinear control: conformal robustness for stability and safety. *IEEE Control Systems Letters* (2025).
54. Hu, Z., De Persis, C. & Tesi, P. Enforcing contraction via data. *IEEE Transactions on Automatic Control* (2025).
55. Huang, B., Ma, X. & Vaidya, U. in *The Koopman Operator in Systems and Control: Concepts, Methodologies, and Applications* 313–334 (Springer, 2020).
56. Johansen, T. & Grancharova, A. Approximate explicit constrained linear model predictive control via orthogonal search tree. *IEEE Transactions on Automatic Control* **48**, 810–815. doi:[10.1109/TAC.2003.811259](https://doi.org/10.1109/TAC.2003.811259) (2003).
57. Jones, C., Grieder, P. & Raković, S. A logarithmic-time solution to the point location problem for parametric linear programming. *Automatica* **42**, 2215–2218. doi:<https://doi.org/10.1016/j.automatica.2006.07.010> (2006).
58. Julian, P., Guivant, J. & Desages, A. A parametrization of piecewise linear Lyapunov functions via linear programming. *International Journal of Control* **72**, 702–715. doi:[10.1080/002071799220876](https://doi.org/10.1080/002071799220876) (1999).
59. Karafyllis, I. Can we prove stability by using a positive definite function with non sign-definite derivative? *IMA Journal of Mathematical Control and Information* **29**, 147–170. doi:[10.1093/imamci/dnr035](https://doi.org/10.1093/imamci/dnr035) (2011).
60. Karafyllis, I. & Tsinias, J. A converse Lyapunov theorem for nonuniform in time global asymptotic stability and its application to feedback stabilization. *SIAM Journal on Control and Optimization* **42**, 936–965. doi:[10.1137/S0363012901392967](https://doi.org/10.1137/S0363012901392967). eprint: <https://doi.org/10.1137/S0363012901392967> (2003).

Bibliographic references

61. El-Kebir, H. & Ornik, M. Sum-of-squares data-driven robustly stabilizing and contracting controller synthesis for polynomial nonlinear systems. *arXiv preprint arXiv:2503.07438* (2025).
62. Khalil, H. K. *Nonlinear systems; 3rd ed.* 2002.
63. Knyazev, A., Zhu, P. & Di Cairano, S. *Explicit model predictive control accuracy analysis* in *2015 54th IEEE Conference on Decision and Control (CDC)* (IEEE, 2015), 2389–2394. doi:[10.1109/cdc.2015.7402565](https://doi.org/10.1109/cdc.2015.7402565).
64. Köhler, J., Andina, E., Soloperto, R., Müller, M. A. & Allgöwer, F. *Linear robust adaptive model predictive control: Computational complexity and conservatism – extended version* 2020. arXiv: [1909.01813](https://arxiv.org/abs/1909.01813) [[eess.SY](https://arxiv.org/abs/1909.01813)].
65. Krasovskii, N. N. Stability of motion. *Stanford University Press, California* (1963).
66. Kudaev, M. The use of ljapunov functions for investigating the behavior of trajectories of systems of differential equations. *Soviet Math* **3**, 1802–4 (1962).
67. LaSalle, J. Some extensions of Liapunov’s second method. *IRE Transactions on Circuit Theory* **7**, 520–527. doi:[10.1109/TCT.1960.1086720](https://doi.org/10.1109/TCT.1960.1086720) (1960).
68. Lee, T.-C., Tan, Y. & Mareels, I. *A generalized matrosov theorem for signal sets on time scales* in *2021 European Control Conference (ECC)* (2021), 1887–1894. doi:[10.23919/ECC54610.2021.9655211](https://doi.org/10.23919/ECC54610.2021.9655211).
69. Lienen, M. & Günemann, S. *Torchode: a parallel ODE solver for pytorch* in *The Symbiosis of Deep Learning and Differential Equations II, NeurIPS* (2022).
70. Lin, Y., Sontag, E. D. & Wang, Y. A smooth converse Lyapunov theorem for robust stability. *SIAM Journal on Control and Optimization* **34**, 124–160. doi:[10.1137/S0363012993259981](https://doi.org/10.1137/S0363012993259981). eprint: <https://doi.org/10.1137/S0363012993259981> (1996).
71. Liu, K.-Z., Wang, X.-F., Teel, A. R., Sun, X.-M. & Liu, J. A matrosov theorem for hybrid systems with memory. *IEEE Transactions on Automatic Control* **66**, 4918–4925. doi:[10.1109/TAC.2020.3042482](https://doi.org/10.1109/TAC.2020.3042482) (2021).

Bibliographic references

72. Liu, S., Liberzon, D. & Zharnitsky, V. Almost Lyapunov functions for nonlinear systems. *Automatica* **113**, 108758 (2020).
73. Loría, A., Panteley, E., Popović, D. & Teel, A. R. A nested matrosov theorem and persistency of excitation for uniform convergence in stable nonautonomous systems. *IEEE Transactions on Automatic Control* **50**, 183–198. doi:[10.1109/TAC.2004.841939](https://doi.org/10.1109/TAC.2004.841939) (2005).
74. Lyapunov, A. M. The general problem of the stability of motion. *International journal of control* **55**, 531–534 (1992).
75. Lyapunov, A. M. The general problem of the stability of motion. *International journal of control* **55**, 531–534 (1992).
76. Malisoff, M. & Mazenc, F. Further remarks on strict input-to-state stable Lyapunov functions for time-varying systems. *Automatica* **41**, 1973–1978. doi:<https://doi.org/10.1016/j.automatica.2005.05.015> (2005).
77. Malisoff, M. & Mazenc, F. Constructions of strict Lyapunov functions for discrete time and hybrid time-varying systems. *Nonlinear Analysis: Hybrid Systems* **2**, 394–407. doi:[10.1016/j.nahs.2006.12.005](https://doi.org/10.1016/j.nahs.2006.12.005) (2008).
78. Malisoff, M. & Mazenc, F. *Constructions of strict Lyapunov functions* (Springer Science & Business Media, 2009).
79. Massera, J. L. On Liapounoff's conditions of stability. *Annals of Mathematics*, 705–721 (1949).
80. Matrosov, V. M. On the stability of motion. *Sibirskij Matematicheskij Zhurnal (Siberian Mathematical Journal)* **4**. English translation available in the same journal, Volume 4, pages 769–790, 1108–1121 (1963).
81. Mayne, D. Q. Model predictive control: recent developments and future promise. *Automatica* **50**, 2967–2986. doi:<https://doi.org/10.1016/j.automatica.2014.10.128> (2014).
82. Mazenc, F. Strict Lyapunov functions for time-varying systems. *Automatica* **39**, 349–353. doi:[https://doi.org/10.1016/S0005-1098\(02\)00233-9](https://doi.org/10.1016/S0005-1098(02)00233-9) (2003).

Bibliographic references

83. Mazenc, F. & Malisoff, M. Extensions of Razumikhin's theorem and Lyapunov, Krasovskii functional constructions for time-varying systems with delay. *Automatica* **78**, 1–13. doi:<https://doi.org/10.1016/j.automat.2016.12.005> (2017).
84. Mazenc, F. & Nesic, D. Strong Lyapunov functions for systems satisfying the conditions of la salle. *IEEE Transactions on Automatic Control* **49**, 1026–1030 (2004).
85. Meigoli, V. & Nikraves, S. K. A new theorem on higher order derivatives of Lyapunov functions. *ISA Transactions* **48**, 173–179. doi:[10.1016/j.isatra.2009.01.001](https://doi.org/10.1016/j.isatra.2009.01.001) (2009).
86. Meigoli, V. & Nikraves, S. K. Y. *Application of higher order derivatives of Lyapunov functions in stability analysis of nonlinear homogeneous systems in* (2009).
87. Meigoli, V. & Nikraves, S. K. Y. Stability analysis of nonlinear systems using higher order derivatives of Lyapunov function candidates. *Systems & Control Letters* **61**, 973–979. doi:<https://doi.org/10.1016/j.sysconle.2012.07.005> (2012).
88. Monshizadeh, N., De Persis, C. & Tesi, P. A versatile framework for data-driven control of nonlinear systems. *IEEE Transactions on Automatic Control* (2025).
89. Narendra, K. S. & Annaswamy, A. M. Persistent excitation in adaptive systems. *International Journal of Control* **45**, 127–160. doi:[10.1080/00207178708933715](https://doi.org/10.1080/00207178708933715). eprint: <https://doi.org/10.1080/00207178708933715> (1987).
90. Nikolaou, M. in, 131–204 (Academic Press, 2001). doi:[https://doi.org/10.1016/S0065-2377\(01\)26003-7](https://doi.org/10.1016/S0065-2377(01)26003-7).
91. Ning, C., He, Y., Wu, M., Liu, Q. & She, J. Input-to-state stability of nonlinear systems based on an indefinite Lyapunov function. *Systems & Control Letters* **61**, 1254–1259. doi:<https://doi.org/10.1016/j.sysconle.2012.08.009> (2012).
92. Oberdieck, R. & Pistikopoulos, E. N. Explicit hybrid model-predictive control: the exact solution. *Automatica* **58**, 152–159. doi:<https://doi.org/10.1016/j.automat.2015.05.021> (2015).

Bibliographic references

93. Oliveira, A., Zheng, J. & Sznaier, M. Convex data-driven contraction with riemannian metrics. *IEEE Control Systems Letters* (2025).
94. Papachristodoulou, A., Anderson, J., Valmorbida, G., Prajna, S., Seiler, P., Parrilo, P. A., Peet, M. M. & Jagt, D. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB* Available from <https://github.com/oxfordcontrol/SOSTOOLS> (2021).
95. Papachristodoulou, A. & Prajna, S. *On the construction of Lyapunov functions using the sum of squares decomposition* in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.* **3** (2002), 3482–3487 vol.3. doi:[10.1109/CDC.2002.1184414](https://doi.org/10.1109/CDC.2002.1184414).
96. Parks, P. C. AM Lyapunov’s stability theory—100 years on. *IMA journal of Mathematical Control and Information* **9**, 275–303 (1992).
97. Poincaré, H. *Les méthodes nouvelles de la mécanique céleste* (Gauthier-Villars et fils, imprimeurs-libraires, 1893).
98. Rawlings, J., Mayne, D. & Diehl, M. *Model Predictive Control: Theory, Computation, and Design* (Nob Hill Publishing, 2017).
99. Rouche, N., Habets, P. & Laloy, M. *Stability theory by Liapunov’s direct method* (Springer, 1977).
100. Şahan, G. Stability analysis by a nonlinear upper bound on the derivative of Lyapunov function. *European Journal of Control* **56**, 118–123. doi:<https://doi.org/10.1016/j.ejcon.2020.02.006> (2020).
101. Şahan, G. & Özdemir, D. *Uniform asymptotic stability by indefinite Lyapunov functions* in *2022 13th Asian Control Conference (ASCC)* (2022), 1771–1774. doi:[10.23919/ASCC56756.2022.9828046](https://doi.org/10.23919/ASCC56756.2022.9828046).
102. Sastry, S. Lyapunov stability theory. *Nonlinear Systems: Analysis, Stability, and Control*, 182–234 (1999).
103. Shen, Y., Bichuch, M. & Mallada, E. *Model-free learning of regions of attraction via recurrent sets* in *2022 IEEE 61st Conference on Decision and Control (CDC)* (2022), 4714–4719.

Bibliographic references

104. Shen, Y., Sibai, H. & Mallada, E. *Generalized barrier functions: integral conditions & recurrent relaxations* in *60th Allerton Conference on Communication, Control, and Computing* (2024), 1–8. presented.
105. Sibai, H. & Mallada, E. *Recurrence of nonlinear control systems: entropy and bit rates* in *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control (HSCC)* (Association for Computing Machinery, New York, NY, USA, 2024), 1–9. doi:<https://doi.org/10.1145/3641513.3650121>.
106. Sibai, H. & Mallada, E. *Recurrence of nonlinear control systems: entropy, bit rates, and finite alphabets*, 1–16 (2025). Submitted.
107. Siegelmann, R. *Complementarity: toward better metrics and optimizing data efficiency in llms*. *Transactions on Machine Learning Research*. Accepted with minor revisions (2025).
108. Siegelmann, R., Shen, Y., Paganini, F. & Mallada, E. *A recurrence-based direct method for stability analysis and gpu-based verification of non-monotonic lyapunov functions* in *2023 62nd IEEE Conference on Decision and Control (CDC)* (2023), 6665–6672. doi:[10.1109/CDC49753.2023.10383373](https://doi.org/10.1109/CDC49753.2023.10383373).
109. Siegelmann, R., Shen, Y., Paganini, F. & Mallada, E. *A recurrence-based direct method for stability analysis and gpu-based verification of non-monotonic Lyapunov functions* in *62nd IEEE Conference on Decision and Control (CDC)* (2023), 6665–6672. doi:[10.1109/CDC49753.2023.10383373](https://doi.org/10.1109/CDC49753.2023.10383373).
110. Siegelmann, R., Shen, Y., Paganini, F. & Mallada, E. *Stability analysis and data-driven verification via recurrent Lyapunov functions*. *IEEE Transactions on Automatic Control* (2025). Submitted.
111. Siegelmann, R. & Siegelmann, H. T. *Meta-analytic operation of threshold-independent filtering (motif) reveals sub-threshold genomic robustness in trisomy: the jörmungandr effect*. *Biochemical and Biophysical Research Communications* **737**, 150802. doi:<https://doi.org/10.1016/j.bbrc.2024.150802> (2024).
112. Sontag, E. *Mathematical Control Theory: Deterministic Finite Dimensional Systems* (Springer New York, 2013).

Bibliographic references

113. Spjotvold, J., Rakovic, S. V., Tondel, P. & Johansen, T. A. *Utilizing reachability analysis in point location problems* in *Proceedings of the 45th IEEE Conference on Decision and Control* (2006), 4568–4569. doi:[10.1109/CDC.2006.377580](https://doi.org/10.1109/CDC.2006.377580).
114. Strässer, R., Berberich, J. & Allgöwer, F. *Data-driven control of nonlinear systems: beyond polynomial dynamics* in *2021 60th IEEE Conference on Decision and Control (CDC)* (2021), 4344–4351.
115. Subbaraman, A. & Teel, A. R. A matrosov theorem for strong global recurrence. *Automatica* **49**, 3390–3395. doi:<https://doi.org/10.1016/j.automatica.2013.08.009> (2013).
116. Sui, D., Feng, L. & Hovd, M. Algorithms for online implementations of explicit mpc solutions. *IFAC Proceedings Volumes* **41**. 17th IFAC World Congress, 3619–3622. doi:<https://doi.org/10.3182/20080706-5-KR-1001.00611> (2008).
117. Teel, A. R., Loria, A., Panteley, E. & Popovic, D. *Matrosov's theorem using a family of auxiliary functions: an analysis tool to aid time-varying nonlinear control design* 2003. arXiv: [math/0304255](https://arxiv.org/abs/math/0304255) [[math.OC](#)].
118. Tøndel, P., Johansen, T. & Bemporad, A. Evaluation of piecewise affine control via binary search tree. *Automatica* **39**, 945–950. doi:[https://doi.org/10.1016/S0005-1098\(02\)00308-4](https://doi.org/10.1016/S0005-1098(02)00308-4) (2003).
119. Tsukamoto, H., Chung, S.-J. & Slotine, J.-J. E. Contraction theory for nonlinear stability analysis and learning-based control: a tutorial overview. *Annual Reviews in Control* **52**, 135–169 (2021).
120. Van Waarde, H. J., Camlibel, M. K. & Mesbahi, M. From noisy data to feedback controllers: nonconservative design via a matrix s-lemma. *IEEE Transactions on Automatic Control* **67**, 162–175 (2020).
121. Vannelli, A. & Vidyasagar, M. Maximal Lyapunov functions and domains of attraction for autonomous nonlinear systems. *Automatica* **21**, 69–80 (1985).
122. Vershynin, R. *Introduction to the non-asymptotic analysis of random matrices* 2011. arXiv: [1011.3027](https://arxiv.org/abs/1011.3027) [[math.PR](#)].

Bibliographic references

123. Wabartha, M. & Pineau, J. *Piecewise linear parametrization of policies: towards interpretable deep reinforcement learning* in *The Twelfth International Conference on Learning Representations* (2024).
124. Wada, T. Some results on asymptotic stability by extending matrosov's theorems. *Proceedings of the Japan Academy, Series A, Mathematical Sciences* **65**, 127–130. doi:[10.3792/pjaa.65.127](https://doi.org/10.3792/pjaa.65.127) (1989).
125. Wen, C., Ma, X. & Ydstie, B. E. Analytical expression of explicit mpc solution via lattice piecewise-affine function. *Automatica* **45**, 910–917 (2009).
126. Werner, S. W. & Peherstorfer, B. On the sample complexity of stabilizing linear dynamical systems from data. *Foundations of Computational Mathematics* **24**, 955–987 (2024).
127. Williams, G., Drews, P., Goldfain, B., Rehg, J. M. & Theodorou, E. A. *Aggressive driving with model predictive path integral control* in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (2016), 1433–1440. doi:[10.1109/ICRA.2016.7487277](https://doi.org/10.1109/ICRA.2016.7487277).
128. Williams, G., Drews, P., Goldfain, B., Rehg, J. M. & Theodorou, E. A. Information-theoretic model predictive control: theory and applications to autonomous driving. *IEEE Transactions on Robotics* **34**, 1603–1622. doi:[10.1109/TRO.2018.2865891](https://doi.org/10.1109/TRO.2018.2865891) (2018).
129. Wu, A., Sadraddini, S. & Tedrake, R. *The nearest polytope problem: algorithms and application to controlling hybrid systems* in *2020 American Control Conference (ACC)* (2020), 1815–1822. doi:[10.23919/ACC45564.2020.9147990](https://doi.org/10.23919/ACC45564.2020.9147990).
130. Yorke, J. A. A theorem on Liapunov functions using \ddot{V} . *Mathematical Systems Theory* **4**, 40–45. doi:[10.1007/bf01705884](https://doi.org/10.1007/bf01705884) (1970).
131. Zeilinger, M. N., Jones, C. N. & Morari, M. Real-time suboptimal model predictive control using a combination of explicit mpc and online optimization. *IEEE Transactions on Automatic Control* **56**, 1524–1534. doi:[10.1109/TAC.2011.2108450](https://doi.org/10.1109/TAC.2011.2108450) (2011).
132. Zhou, B. On asymptotic stability of linear time-varying systems. *Automatica* **68**, 266–276. doi:<https://doi.org/10.1016/j.automatica.2015.12.030> (2016).

Bibliographic references

133. Zhou, B. Stability analysis of non-linear time-varying systems by Lyapunov functions with indefinite derivatives. *IET Control Theory & Applications* **11**, 1434–1442. doi:[10.1049/iet-cta.2016.1538](https://doi.org/10.1049/iet-cta.2016.1538) (2017).
134. Zhou, B. & Egorov, A. *Time-varying Razumikhin and Krasovskii stability theorems for time-varying delay systems* in (2016), 1041–1046. doi:[10.1109/CCDC.2016.7531137](https://doi.org/10.1109/CCDC.2016.7531137).
135. Zhou, K., Doyle, J. & Glover, K. Robust and optimal control. *Control Engineering Practice* **4**, 1189–1190 (1996).