LEARNING SAFE REGIONS IN HIGH-DIMENSIONAL DYNAMICAL SYSTEMS VIA RECURRENT SETS

by
Yue Shen

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

Baltimore, Maryland
September, 2024

Abstract

Safety certification in dynamical systems has relied heavily on the identification of an invariant set that strictly requires trajectories to always lie within it. This foundational approach addresses two key safety objectives: stability and avoidance. To this end, techniques such as the Lyapunov method and the barrier function method that are capable of characterizing an invariant set have been pivotal. However, the invariance requirement poses strict constraints on the learning outcome. Therefore, as the system dimension increases, directly characterizing invariant sets or identifying functions that define invariant sets typically demands considerable domain-specific knowledge or extensive computational resources.

This thesis seeks to develop new data-driven methodologies that facilitate the verification of stability and avoidance in dynamical systems, without relying on the identification of invariant sets. A key innovation of this thesis is the application of the concept of recurrence to relax the stringent constraints imposed by invariance. Specifically, a set is recurrent if trajectories originating from it return to it infinitely often. By leveraging recurrence, safety can be characterized with enhanced efficiency and accuracy.

This thesis theoretically establishes necessary and sufficient conditions for using recurrence to characterize safety, offering a deeper understanding of how recurrence can serve as a reliable proxy for invariance. Practically, it introduces practical, data-driven algorithms that utilize only a finite number of finite-length sampled trajectories to determine safe regions within a dynamical system. Optimized for

computational efficiency, these algorithms can be implemented on parallel processing units, making them highly applicable in real-world scenarios where rapid and reliable safety verification is crucial.

Thesis Committee

Dr. Enrique Mallada (Primary Advisor)
      Associate Professor
      Department of Electrical and Computer Engineering
      Johns Hopkins University

Dr. Pablo Iglesias
      Edward J. Schaefer Professor
      Department of Electrical and Computer Engineering
      Johns Hopkins University

Dr. Sijia Geng
      Assistant Professor
      Department of Electrical and Computer Engineering
      Johns Hopkins University

Dedicated to world peace.

Acknowledgements

Over the five years of my Ph.D. journey, I have been incredibly fortunate to receive invaluable support and guidance from many individuals.

I would like to begin by expressing my deepest gratitude to my advisor, Enrique Mallada, who has been an exceptional mentor, collaborator, and friend. I cannot imagine how different my journey would have been without such a great advisor in my life. I could easily dedicate an entire thesis to highlighting Enrique's remarkable research abilities, his unwavering support, and the genuine care he shows to everyone. In short, saying someone is similar to Enrique is the highest compliment I, and all of our group family members, can offer. I deeply admire his diligence, intelligence, and passion, and it has become my life goal to strive toward being the near-perfect individual that he is.

I extend my sincere thanks to Pablo Iglesias and Sijia Geng for serving on my thesis committee and for their valuable comments and suggestions on my dissertation. I also wish to thank Louis Whitcomb and Marin Kobilarov for serving on my graduate board oral exam committee. I am fortunate to be able to collaborate with many genius people, including Maxim Bichuch and Hussein Sibai.

I want to extend my appreciation to my colleagues who have supported and collaborated with me throughout this journey: Yan Jiang, Chengda Ji, Haralampos Avraam, Mustafa Devrim Kaba, Tianqi Zheng, Hancheng Min, Rajni Kant Bansal, Pengcheng You, Jay Guthrie, Eliza Cohn, Eli Pivo, Gary Gao, Dhananjay Anand,

Contents

List of Figures

Chapter 1

Introduction

The contemporary control paradigm is undergoing one of the most fundamental transformations since its inception. Sequential decision-making techniques, such as reinforcement learning enhanced with deep neural network architectures, have shown incredible success by achieving super-human performance in various situations like strategic gameplay, product recommendations, and stock trading. These approaches hold the potential to significantly impact numerous application areas, including autonomous transportation, industrial automation, and power networks.

Despite the significant advances, the deployment of these techniques in safety-critical systems presents a formidable challenge. These environments often involve high-dimensional dynamical systems where the stakes for operational safety and reliability are exceptionally high. As the dimensionality of the system increases, the task of characterizing a safe region becomes computationally intensive and may result in a highly complex representation of the resultant safety set.

Among the prevailing challenges, the fundamental issue lies in the reliance on identifying an invariant set — a set within which system trajectories are continuously contained — to characterize a safe region. This invariance requirement implicitly couples the shape and topology of the learning outcomes with the system trajectories. For instance, any invariant safe set approximation must be connected,

and system trajectories must consistently point inward at every point along its boundary. Therefore, identifying an invariant safe region typically demands considerable domain-specific knowledge or extensive computational resources due to this intricate coupling.

This dissertation research aims to develop rigorous methodologies for efficiently certifying the safety of high-dimensional dynamical systems with minimal system description. It specifically addresses the following two commonly desired objectives:

- What is the set of initial conditions such that a given stable equilibrium is guaranteed to be converged to?

- What is the set of initial conditions such that a given unsafe state space region is guaranteed to be avoided with?

A key innovation of this work is the application of the concept of **recurrence** to relax the stringent constraints imposed by **invariance**. Specifically, a set is recurrent if trajectories originating from it return to it infinitely often. Recurrent sets allow trajectories to leave the set and thus constitute a strict relaxation of invariance—every invariant set is recurrent, but not the other way around. Leveraging this innovative notion of recurrence, the aforementioned dependence can be disentangled, allowing for the characterization of safety with greater efficiency and accuracy across diverse settings. Our parallel work has shown recurrence to be a powerful mechanism for analyzing dynamical systems, including certifying stability via generalized Lyapunov conditions [1]. Note that, from an information theoretical viewpoint, making a set recurrent requires less information than making it invariant [2], thus making it a more beneficial search target. This dissertation will explore these two objectives in detail, illuminating the opportunities and challenges associated with this vital field of safe learning.

1.1 Objectives

1.1.1 Learning the region of attraction of a stable equilibrium

The first goal of this dissertation research is to determine the region of attraction (ROA) of a stable equilibrium point, a fundamental task in the control design process. This challenge has a significant history in the fields of nonlinear control and dynamical systems theory [3]. For an asymptotically stable equilibrium point $x^* \in \mathcal{X}$, we aim to learn the largest set $\mathcal{A}(x^*) \subseteq \mathcal{X}$ for which $x(0) \in \mathcal{A}(x^*) \Rightarrow \lim_{t \to \infty} x(t) = x^*$.

From a theoretical standpoint, there has been a thorough study of conditions that guarantee several topological properties of such set, e.g., being connected, open, dense, smooth [4]. From a practical standpoint, having a representation of such region allows us to test the limits of controller designs, which are usually based on (possibly linear) approximations of nonlinear systems [5], and provides a mechanism for safety verification of different (possibly disturbed) operating conditions [6].

Unfortunately, it is known that finding an analytic form of the region of attraction is difficult and, in general, impossible [3, p. 122]. As a result, most efforts in characterizing the ROA focus on finding inner approximations by means of invariant sets, which, while not exhaustive and efficient when applied to high-dimensional systems, provide valuable insights into the stability boundaries.

1.1.2 Characterizing a safe region of a dynamical system

In addition to stabilizing the system towards a stable equilibrium point, the second goal of this dissertation research is to characterize a safe state space region that is a set of initial conditions such that a given unsafe state space region is guaranteed to be avoided. Specificaly, we consider a continuous-time dynamical system $\dot{x}(t) =$

$f(x(t))$ with the state vector $x(t) \in \mathcal{X}$. We aim to guarantee safety by learning a safe set $\mathcal{X}_s \subseteq \mathcal{X}$ such that the trajectory starting from $\mathcal{X}_s$ can avoid a set of unsafe states $\mathcal{X}_u$. Precisely, given an unsafe set $\mathcal{X}_u \subseteq \mathcal{X}$, we say a set $\mathcal{X}_s$ is safe if starting from any initial states $x(0) \in \mathcal{X}_s$, $x(t) \notin \mathcal{X}_u$ for all $t \geq 0$. This formulation is general enough to capture complex dynamical systems. For example, we can set the individual movements (pitch, yaw, and roll) as well as mutual distances of a group of UAVs as states. Then, the unsafe set contains the regions that lead to a crush (e.g., the pitch is low) or a collision (e.g., the mutual distance is small). Guaranteeing safety is challenging even when the dynamic model is available. Due to the curse of dimensionality, direct sampling methods that check the invariance of a tentatively safe set are usually infeasible [7]. A more tractable approach is to use Barrier Functions (BFs) that render invariant safe sets [8].

1.2 Prior Work

The related works aimed at achieving the aforementioned safety objectives are summarized in Section 1.2.1 and 1.2.2, respectively. It is important to note that these approaches all heavily depend on the identification of an invariant set, which requires trajectories to consistently remain within specified boundaries. Additionally, we invite the reader to refer to Section 1.2.3, where similar but simpler definitions of recurrence are discussed in different contexts such as stability and constrained model predictive control.

1.2.1 Prior work on learning the region of attraction of a stable equilibrium

Several methodologies for computing inner approximations of the ROA have been proposed in the literature. In a broad sense, they can be classified into two groups: **Lyapunov methods** and **non-Lyapunov methods**. We briefly review such methods

in cases when accurate, inaccurate, or no information about the dynamic model is present.

Lyapunov methods: Without calculating the exact solution of a system, Lyapunov methods utilize the fact that Lyapunov functions are certificates of asymptotic stability and then build inner approximations using its sublevel sets.

When an exact description of the dynamics is available, methods for finding such Lyapunov functions are surveyed in, e.g., [9]. In particular, [10] and [11] construct Lyapunov functions that are solutions of Zubov's equation, and [12] searches for piece-wise linear Lyapunov functions that are found via linear programming. Similarly, piece-wise quadratic parameterizations of Lyapunov functions using methods based on linear matrix inequalities are considered in [13]. Finally, recent work [14] leverages the universal approximation property of neural networks to estimate the ROA of an asymptotically stable equilibrium point of general nonlinear dynamical systems.

In the presence of uncertainty in the dynamic model, [15, 16, 17, 18] generalize Lyapunov methods by finding a common Lyapunov function across the entire uncertainty set. Additionally, [19, 20, 21] first acquire a Lyapunov function from the deterministic part of the system dynamics and then use experimental data to bound the uncertainty and expand the Lyapunov function level set.

Finally, when there is no information about the dynamic model, trajectory data can be used to fit Lyapunov functions by leveraging tools such as converse Lyapunov theorems [22] and neural networks [23, 24].

Non-Lyapunov methods: Alternatively, non-Lyapunov methods focus directly on the properties of the ROA. With an exact description of the dynamics, [25] [26] sample reverse trajectories to derive the boundary of ROA directly from the stable manifold of the equilibria on the boundary. In a model-free setting, [27]

utilizes a support vector machine, trained with experimental data gathered through hybrid active learning techniques, to directly delineate the boundary of the ROA. Meanwhile, [28] integrates auto-encoding neural networks with topological tools to represent the underlying nonlinear dynamics as a graph, and then characterize the ROA from it.

1.2.2 Prior work on characterizing the safe region of a dynamical system

We similarly classify methodologies for computing a safe region of a dynamical system into two groups: **Barrier function methods** and **Reachability methods**. Note that **Barrier function methods** characterize the safe region using a Lyapunov-like function instead of computing the exact solution of the dynamical system, which can be considered as a generalization of the **Lyapunov methods**. In contrast, **Reachability methods** are based directly on the trajectory information just like **non-Lyapunov methods**. We will discuss these connections later as we progress.

Barrier function methods: Similar to the Lyapunov function, the barrier certificate function, first proposed by [29, 30], can verify the safety (avoidance) property of an autonomous system. Then, [31] further relaxes the conditions that the barrier certificate function should satisfy to capture the safe region with a tighter underestimation. Given an autonomous system and an unsafe set, these references use SOS programming to compute a barrier certificate function and use its zero-level set as a boundary of the safe state space region. The control barrier function is a generalization of the barrier certificate function in the presence of control inputs. In particular, [32] first obtains feedback controls for affine control systems to avoid unsafe state regions. References [33, 34] unify the control barrier function and the control Lyapunov function (first proposed in [35]) to achieve safety and stability at the same time. Reference [36] considers the disturbed setting and extends the

control barrier function to the robust setting. Finally, when the system is unknown, most of the current corpus accordingly learns a control barrier function from expert trajectories using, for example, a Neural Network [37] or a Support Vector Machine (SVM) [38].

Reachability methods: In contrast to barrier function methods, reachability approaches focus on verifying safety by approximating the reachable set. A common strategy is to use geometric shapes such as ellipsoids [39, 40], polyhedra [41, 42], or geometric programs [43] to over-approximate the unsafe set and propagate it backward over time, particularly under linear dynamics. For nonlinear systems, this approach has been generalized to propagate more complex geometric approximations, as explored by [44]. These methods are collectively referred to as 'geometric' or 'over-approximative methods.'

Alternatively, level set methods represent the reachable set as the zero sub-level set of a signed distance function, which allows for a more accurate representation of the true reachable set compared to simple geometric shapes. The boundary of this reachable set can then be propagated under nonlinear dynamics by solving the Hamilton-Jacobi equation [45, 46] or the Hamilton-Jacobi-Isaacs partial differential equation [47, 48].

1.2.3 Prior work on the notion of recurrence

Recurrence is a fundamental mechanism of dynamical systems. For example, the classic Poincaré recurrent theorem [49] is commonly discussed in statistical physics [50, 51], states that certain classes of dynamical systems will recur to points arbitrarily close to their initial state in finite time. Along the same line, [52] defines the notion of chain recurrence that requires a ε-chain from a point to itself. The concept of chain recurrence, as utilized by [53] in his Fundamental Theorem of Dynamical Systems, demonstrates the existence of a continuous Lyapunov function for con-

tinuous flows in compact metric spaces. This highlights the intimate connection between chain recurrence and Lyapunov functions. Moreover, [54] expands this application to discrete dynamical systems.

In the context of constrained control, [55, 56] share a similar definition of recurrence with our approach. Specifically, [55] revisits the concept of set invariance and introduces the notion of p-invariance, which requires trajectories to satisfy a constraint set within p steps. Building on this, [56] develops a model predictive control scheme that relaxes the stringent terminal requirement, which usually demands trajectories to terminate within an invariant set subject to all constraints. Instead, the scheme assumes a given p-recurrent set inside the invariant set and then requires trajectories to visit this recurrent set within p steps.

1.3 Thesis Contributions

The principal innovation of this dissertation is the application of the concept of recurrence to mitigate the stringent constraints imposed by invariance. Rather than focusing on learning invariant sets that strictly require trajectories to always lie within the set, we advocate learning sets that satisfy this more flexible notion, i.e., recurrence. The application of this innovative approach to the two primary objectives of this research yields multiple contributions:

In the context of learning the region of attraction of a stable equilibrium:

- **Recurrent based approach — Sufficiency:** We propose the notion of recurrence as an alternative property of invariance that guarantees a set to be contained in the region of attraction. This result allows us, in turn, to recast the region of attraction learning problem as a problem of characterizing a recurrent set that contains the stable equilibrium point inside.

- **Recurrent based approach — Necessity:** Conversely, we show that almost

every set contained within the region of attraction of an asymptotically stable equilibrium point, containing such point within, is inherently a recurrent set. This observation is distinctively not applicable to invariant sets in general. Therefore, a much vaster family of sets can be certified as subsets of the region of attraction.

- **Data-driven algorithms:** We provide several practical algorithms designed to characterize the region of attraction through counter-examples of recurrence, utilizing finite-length trajectory samples. These methodologies are termed "model-free" as they do not necessitate an explicit description of the system's dynamics but instead depend solely on a process that generates sample trajectories. Such trajectories can be processed simultaneously, leveraging parallelizable processing units. Moreover, even after the initial completion of our algorithm, it can be reinitiated to enhance the accuracy of the approximation.

- **Sample complexity guarantees:** We demonstrate that after collecting a finite number of samples, the sampling process can be concluded, and it can be asserted that our learning outcome is an inner approximation of the region of attraction with any predetermined error probability.

(This set of contributions is detailed in the following two publications: "Model-free Learning of Regions of Attraction via Recurrent Sets" (Presented at IEEE Conference on Decision and Control 2022) and "Model-free Learning of Regions of Attraction via Recurrent Sets: Sample Complexity and Progressive Improvement" (Submitted to IEEE Transactions on Automatic Control))

In the context of characterizing the safe region of a dynamical system:

- **Recurrent based approach — Sufficiency:** We systematically generalize barrier functions, commonly used to ensure the invariance of their zero superlevel set, to recurrent barrier functions that guarantee the recurrence of these

sets. Subsequently, we demonstrated how a recurrent set is sufficient to guarantee the system's safety (avoidance).

- **Recurrent based approach — Necessity:** Under mild conditions, we show that a simple sign distance function can satisfy our relaxed recurrent barrier function criteria. Again, this observation does not generally apply to the classical barrier functions based on the notion of invariance, highlighting the generality of our recurrent conditions.

- **Data-driven algorithms:** We develop GPU-based algorithms designed to characterize the safe region by utilizing a finite number of finite-length trajectory samples. These algorithms maintain a "model-free" setting, foregoing an explicit description of the system and relying solely on processes that generate sample trajectories

(This set of contributions is detailed in the following two publications: "Generalized Barrier Functions: Integral Conditions & Recurrent Relaxations" (Accepted for presentation at Allerton Conference 2024) and "Model-free Learning of Safe Regions via Recurrent Barrier Functions" (To be submitted to IEEE Transactions on Automatic Control))

We validate our proposed algorithms using both demonstrative and practical examples, which reveal improved accuracy and efficiency. These validations underscore the algorithms' robustness and their potentials in real-world scenarios.

1.4 Additional Publications

Earlier publications by the author on the energy storage optimal control [57], and model predictive control [58] problems are not included in the thesis as they diverge in focus from the current subject. They are referenced here for completeness.

Chapter 2

Model-free Learning of Regions of Attraction via Recurrent Sets

In this chapter, we explore the problem of learning the region of attraction (ROA) of a stable equilibrium point within a model-free framework. The chapter is organized as follows: Section 2.1 outlines the ROA learning problem we aim to solve. In this section, we revisit foundational concepts such as set invariance and Lyapunov stability, which are crucial to our approach. Section 2.2 introduces the pivotal concept of recurrence, which is central to our analysis. This section presents initial results that clarify the relationship between recurrence and containment within the ROA, illustrating both the sufficiency and necessity of our recurrent-based approach. We demonstrate that under mild conditions, recurrent sets invariably form subsets of the ROA, and nearly every subset of the ROA will exhibit recurrence within a finite timeframe. The proposed algorithms and the corresponding guarantees are given in Section 2.3. Section 2.4 offers numerical examples, including the analysis of a demonstrative 2D dynamical system and a real-world 4D power system transient stability problem, showcasing the accuracy and efficiency of our algorithms. Finally, we conclude in Section 2.6.

Notation

Given a set S, we denote its boundary as ∂S, its interior as $\operatorname{int} S$. We use $\operatorname{sd}(x, S)$ to denote the signed distance between a point x and S, i.e.,

$$\operatorname{sd}(x, S) := \begin{cases} \inf_{y \in \partial S} \|y - x\|_2 & \text{if } x \notin S \\ -\inf_{y \in \partial S} \|y - x\|_2 & \text{if } x \in S. \end{cases}$$

We further use $\operatorname{P}_S(x)$ to denote the set of projections of a point x on a non-empty closed set S, i.e.,

$$\operatorname{P}_S(x) := \arg \min_{y \in S} \|y - x\|_2.$$

We use $\mathcal{B}_r(x)$ to denote the closed ball of radius r around a point x. When the point x is at the origin 0, we abbreviate $\mathcal{B}_r(x) = \mathcal{B}_r(0)$ as $\mathcal{B}_r$.

We then consider the following set operations that expand or shrink the set S by a width of $\alpha \in \mathbb{R}$:

$$(S)_\alpha := \begin{cases} S + \mathcal{B}_\alpha & \text{if } \alpha \geq 0 \\ S \backslash \{\partial S + \operatorname{int} \mathcal{B}_\alpha\} & \text{o.w.}, \end{cases}$$

where the $'+'$ stands for the Minkovski sum.

2.1 Problem Formulation

We consider a continuous time dynamical system

$$\dot{x}(t) = f(x(t)), \tag{2.1}$$

where $x(t) \in \mathbb{R}^d$ is the state at time t, and the map $f : \mathbb{R}^d \to \mathbb{R}^d$ is continuously differentiable and (globally) Lipschitz. Given initial condition $x(0) = x_0$, we use $\phi(t, x_0)$ to denote the solution of (2.1). Using this notation, the positive orbit of x_0 is given by $\mathcal{O}_+(x_0) = \{y \in \mathbb{R}^d : y = \phi(t, x_0), t \in \mathbb{R}^+\}$.

Definition 2.1 (ω-limit Set)**.** *Given an initial condition x_0, its ω-limit set $\Omega(x_0)$ is the set of points $y \in \mathbb{R}^d$ for which there exists a sequence t_n indexed by $n \in \mathbb{N}$ satisfying*

$\lim_{n \to \infty} t_n = \infty$ *and* $\lim_{n \to \infty} \phi(t_n, x_0) = y$. *We will further use* $\Omega(f)$ *to denote the ω-limit set of* (2.1), *which is the union of ω-limit sets of all* $x \in \mathbb{R}^d$.

Note that by definition, if x^* is an equilibrium of (2.1), then it follows that $x^* \in \Omega(f)$. We would like then to learn the set of initial conditions that converge to x^*.

Definition 2.2 (Region of Attraction). *Given an invariant set* $S \subseteq \Omega(f)$, *the region of attraction (ROA) of* S *under* (2.1) *is defined as*

$$\mathcal{A}(S) := \left\{ x_0 \in \mathbb{R}^d \,|\, \liminf_{t \to \infty} d(\phi(t, x_0), S) = 0 \right\}, \tag{2.2}$$

where $d(y, S)$ *is the distance from the solution y to the set S, i.e.,* $d(y, S) := \min_{x \in S} \|x - y\|_2$. *When the set S is a singleton that contains exactly one point (say x), we abbreviate* $\mathcal{A}(S) = \mathcal{A}(\{x\})$ *as* $\mathcal{A}(x)$.

Note that without further assumptions, the set (2.2) may be a singleton, have zero measure, or be disconnected, making the problem of characterizing (2.2) from samples almost impossible. We thus make the following assumption.

Assumption 2.1. *The system* (2.1) *has an asymptotically stable equilibrium at* $x^* \in \mathbb{R}^d$.

Remark 1. *It follows from Assumption 2.1 that the ROA* $\mathcal{A}(x^*)$ *is an open contractible set [59], i.e., the identity map of* $\mathcal{A}(x^*)$ *to itself is null-homotopic [60].*

Having set up the necessary assumption for an ROA to be learnable, we now move on to a certain property that helps us to characterize subsets of the region of attraction.

By definition, $\mathcal{A}(S)$ satisfies the invariant property that every trajectory that starts in the set $\mathcal{A}(S)$ remains in the set for all future times, i.e., $\mathcal{A}(S)$ is a positively invariant set [3].

Definition 2.3 (Invariant Set). *A set $\mathcal{I} \subseteq \mathbb{R}^d$ is invariant w.r.t. (2.1) if and only if:*

$$x_0 \in \mathcal{I} \implies \phi(t, x_0) \in \mathcal{I}, \quad \forall t \in \mathbb{R}^+. \tag{2.3}$$

The notion of positive invariance is fundamental for control. It is used to trap trajectories in compact sets and allows the development of the Lyapunov theory. By trapping trajectories on sub-level sets of a function, one can guarantee boundedness of trajectories, stability, and even asymptotic stability via a gradual reduction of the value of the Lyapunov function. A natural approach is therefore to search for Lyapunov functions [3] that render its sublevel sets as invariant inner-approximations of $\mathcal{A}(x^*)$. Such methods are particularly justified after the fundamental result by Vladimir Zubov [61] that guarantees the existence of such a function:

Theorem 2.1 (Zubov's Existence Criterion). *A set $\mathcal{A}$ containing x^* in its interior is the region of attraction of x^* under (2.1) if and only if there exist continuous functions V, h such that the following hold:*

- *$V(x^*) = h(x^*) = 0$, $0 < V(x) < 1$ for $x \in \mathcal{A} \backslash \{x^*\}$, $h(x) > 0$ for $x \in \mathbb{R}^d \backslash \{x^*\}$.*

- *For every $\gamma_2 > 0$, there exists $\gamma_1 > 0$, $\alpha_1 > 0$ such that $V(x) > \gamma_1$, $h(x) > \alpha_1$, whenever $\|x\| \geq \gamma_2$.*

- *$V(x_k) \to 1$ for all sequences $\{x_k\}$ such that $x_k \to \partial \mathcal{A}$ or $\|x_k\| \to \infty$.*

- *V and h satisfy*

$$(\mathcal{L}_f V)(x) = -h(x)(1 - V(x)) \sqrt{1 + \|f(x)\|^2}, \tag{2.4}$$

  where $(\mathcal{L}_f V)(x)$ is the Lie derivative of V under the flow induced by f.

Particularly, when $f(x)$ is continuously differentiable, $h(x)$ can always be selected such that V is differentiable, i.e., $(\mathcal{L}_f V)(x) = \nabla V(x)^T f(x)$.

Corollary 2.1. *Under Assumption 2.1, there exists a Lyapunov function V with domain on $\mathcal{A}(x^*)$ such that for any $c \in (0,1)$ the sublevel set $V_{\leq c} := \{x : V(x) \leq c\}$ is a contractible invariant subset of $\mathcal{A}(x^*)$.*

Proof. See Appendix 2.5 □

The Zubov's function V of Theorem 2.1 provides a parametric family $\{V_{\leq c} : c \in (0,1)\}$ of positively invariant sets inside $\mathcal{A}(x^*)$. Further, while Zubov's result provides a constructive method for $V(x)$, by means of solving a partial differential equation, such a method becomes impractical in the absence of a descriptive model for (2.1). Thus, in the absence of an exact model of the dynamics, it is natural to try to find a set inside $\mathcal{A}(x^*)$ that is positively invariant in a robust sense, in the presence of bounded uncertainty [18], or that is positively invariant with high probability [19].

However, one of the caveats of positively invariant sets is that they need to be specified very carefully, in the sense that even a good approximation of a positively invariant set is not necessarily positively invariant. Particularly, subsets of positively invariant sets need not be positively invariant. This indirectly imposes strict constraints on the complexity of the set that one needs to learn via (2.3). This motivates the alternative proposed in the next section.

2.2 Recurrent and τ-Recurrent Sets

This section delves into the relaxed notion of invariance, termed here as recurrence and τ-recurrence, contrasting it with traditional invariant sets. A visual illustration highlighting the differences between invariant, recurrent, and τ-recurrent sets can be found in Figure 2-1.

We then introduce the recurrent set and show how it constitutes a more flexible and more general class of objects of study.

15
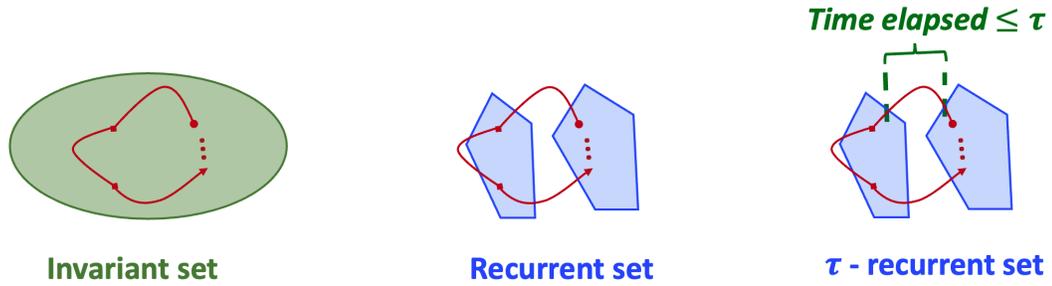
Figure 2-1. Illustration of recurrent and τ-recurrent sets. In particular, trajectories starting from a recurrent set will return to it infinitely often, and a τ-recurrent set further poses a time limit for such a return.

Definition 2.4 (Recurrent Set). *A set $\mathcal{R} \subseteq \mathbb{R}^d$ is recurrent w.r.t. (2.1), if for any point $x_0 \in \mathcal{R}$ and any time $t \geq 0$, there exists a time $t' > t$, such that $\phi(t', x_0) \in \mathcal{R}$.*

Note that while a recurrent set is not inherently invariant, it guarantees that solutions originating from this set will invariably return to it infinitely often. According to Definition 2.3, any positively invariant set, $\mathcal{I}$, is inherently recurrent. This establishes that Definition 2.4 generalizes the notion of positive invariance by allowing the solution $\phi(t, x_0)$ to step outside the set $\mathcal{R}$ for some finite time. Additionally, it is important to highlight that our analysis does not assume connectivity of $\mathcal{R}$, allowing for the possibility that $\mathcal{R}$ could be disconnected. This flexibility can enhance our ability to more accurately approximate the ROA.

However, a notable limitation of the current definition of recurrence is that while it confirms that trajectories will eventually return to the set $\mathcal{R}$, it does not specify the time frame within which this return must occur. The absence of a defined time limit for a trajectory's return to $\mathcal{R}$ after departure can complicate the practical characterization of recurrent sets. This limitation motivates the following stricter notion of recurrence, aimed at providing a more precise temporal framework for these returns.

Definition 2.5 (τ-Recurrent Set). *A set $\mathcal{R} \subseteq \mathbb{R}^d$ is τ-recurrent w.r.t. (2.1), if for any*

16

point $x_0 \in \mathcal{R}$ and any time $t \geq 0$, there exists a $t' \in (t, t + \tau]$, such that $\phi(t', x_0) \in \mathcal{R}$.

Such a τ-recurrent set further guarantees that solutions starting in this set will visit it back within τ-seconds, infinitely often. Thus, τ-recurrent sets are recurrent by definition. Also, according to Definition 2.3, any invariant set is τ-recurrent for any $\tau > 0$. Conversely, a 0-recurrent set is strictly invariant.

Similar to the definition of recurrent, Definition 2.5 generalizes the notion of invariance by permitting the solution $\phi(t, x_0)$ to temporarily exit the set $\mathcal{R}$, but limits this departure to a predetermined time interval, τ. This adaptation introduces a necessary regularity condition to the general concept of recurrence, providing a practical framework to characterize the region of attraction using finite-length trajectories.

2.2.1 Recurrent sets are subsets of the ROA

One concern may be however that by allowing $\phi(t, x_0)$ to leave the set $\mathcal{R}$, this will lead to trajectories that diverge, thus leading to unstable behavior. The following result shows that under mild assumptions, this should not be a source of concern as any trajectory starting in a recurrent set $\mathcal{R}$ will remain within $\mathcal{R}$ after a certain time $T > 0$.

Lemma 2.1. *Let $\mathcal{R} \subset \mathbb{R}^d$ be a compact recurrent set satisfying $\partial \mathcal{R} \cap \Omega(f) = \emptyset$. Then for any $x_0 \in \mathcal{R}$, there exists some time $T > 0$, such that the solution $\phi(t, x_0) \in \mathcal{R}$ for all $t \geq T$.*

Proof. See Appendix 2.5 □

After characterizing regularity conditions for trajectories starting from a recurrent set $\mathcal{R}$, we are ready to show how recurrent sets can be used to characterize subsets of an ROA.

Theorem 2.2. *Let $\mathcal{R} \subset \mathbb{R}^d$ be a compact set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$. Then $\mathcal{R}$ is recurrent if and only if $\Omega(f) \cap \mathcal{R} \neq \emptyset$ and $\mathcal{R} \subset \mathcal{A}(\Omega(f) \cap \mathcal{R})$.*

Proof. See Appendix 2.5 □

Theorem 2.2 illustrates the recurrence of a compact set $\mathcal{R}$, together with the condition $\partial\mathcal{R} \cap \Omega(f) = \emptyset$, necessarily implies its containment within the region of attraction of $\Omega(f) \cap \mathcal{R}$. As a result, by imposing mild conditions on $\Omega(f)$, one leads to the following quite useful result.

Corollary 2.2. *Let assumptions 2.1 hold. Further, let $\mathcal{R}$ be a compact set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \mathcal{R} = \{x^*\}$. Then the set $\mathcal{R}$ is recurrent if and only if $\mathcal{R} \subset \mathcal{A}(x^*)$.*

Proof. See Appendix 2.5 □

Note that Theorem 2.2 and Corollary 2.2 unveil the equivalent relationship between recurrence and containment within the region of attraction. This equivalence is distinctively not applicable to invariant sets in general, emphasizing the broader applicability of recurrent sets.

2.2.2 Almost every ROA subset is finite-time recurrent

Corollary 2.2 implies that one may use recurrence as a mechanism for finding inner approximations for $\mathcal{A}(x^*)$. Specifically, if a compact recurrent set encompasses the equilibrium point x^*, it can be confidently regarded as a subset of the ROA $\mathcal{A}(x^*)$.

However, in practical applications, we often have access only to finite-length trajectory samples. This limitation necessitates focusing on characterizing τ-recurrent sets, which are adapted to the constraints of finite observation windows.

To address this, we introduce the following theorem, which asserts that a compact set $\mathcal{R} \subseteq \mathcal{A}(x^*)$ will satisfy the τ-recurrent requirement as long as τ exceeds a

certain lower bound. This theorem ensures that, even with finite-length trajectory data, we can still reliably characterize a subset of the ROA.

Theorem 2.3. *Let Assumption 2.1 hold, and consider a compact set $\mathcal{R} \subseteq \mathcal{A}(x^*)$ satisfying $x^* \in \mathrm{int}\, \mathcal{R}$ and $\mathcal{R} \cap \partial\mathcal{A}(x^*) = \emptyset$. Then there exists positive constants $\underline{c}$, $\bar{c}$, and a, depending on $\mathcal{R}$, such that for all $\tau \geq \bar{\tau} := \frac{\bar{c}-\underline{c}}{a}$, the set $\mathcal{R}$ is τ-recurrent. Further, starting from any point $x \in \mathcal{R}$, the solution $\phi(t,x) \in \mathcal{R}$ for all $t \geq \bar{\tau}$.*

Proof. See Appendix 2.5 □

Note that the lower bound on τ in Theorem 2.3 implicitly depends on the set $\mathcal{R}$. This makes the process of learning a recurrent set difficult as τ would change, and the set is updated. To eliminate this dependence, one is required to introduce conservativeness. To that end, for given $\delta > 0$, $c \in (0,1)$, and V as in Theorem 2.1, we consider the set

$$\tilde{\mathcal{A}} := V_{\leq c}\backslash\{\mathrm{int}\, \mathcal{B}_\delta + x^*\}, \tag{2.5}$$

where as mentioned before $V_{\leq c} := \{x : V(x) \leq c\}$ is a compact Lyapunov sublevel set contained in $\mathcal{A}(x^*)$. The sign '+' in (2.5) represents the Minkowski sum, and $\mathcal{B}_\delta$ is a closed δ ball centered at the origin, i.e., $\mathcal{B}_\delta = \{x | \|x\|_2 \leq \delta\}$. Note we further choose $\delta > 0$ to be small enough such that $\mathcal{B}_\delta + x^* \subseteq V_{\leq c}$, and the set $V_{\leq c}$ can approximate the ROA $\mathcal{A}(x^*)$ with arbitrary (2-norm) accuracy as $c \to 1$ in the case that $\mathcal{A}(x^*)$ is bounded.

Then, by denoting $\underline{c}(\delta)$ as the minimum Lyapunov function value in $\tilde{\mathcal{A}}$, and $a(\delta)$ as the largest Lie derivative within the set $C_\delta = \{x \in \mathbb{R}^d : \underline{c}(\delta) \leq V(x) \leq c\}$, i.e.,

$$\underline{c}(\delta) := \min_{x \in \tilde{\mathcal{A}}} V(x), \quad \text{and} \quad a(\delta) := \max_{x \in C_\delta} \nabla V(x)^T f(x),$$

we obtain a lower bound on τ that is independent of $\mathcal{R}$.

Theorem 2.4. *Let Assumption 2.1 hold, and consider $\delta > 0$, $c \in (0,1)$ and a compact set $\mathcal{R}$ satisfying: $\mathcal{B}_\delta + x^* \subseteq \mathcal{R} \subseteq V_{\leq c}$. Then $\mathcal{R}$ is τ-recurrent for $\tau \geq \bar{\tau}(\delta) := (\underline{c}(\delta) - c)/a(\delta)$. Moreover, when $t \geq \bar{\tau}(\delta)$, $\phi(t,x) \in \mathcal{R}$ for any point $x \in \mathcal{R}$.*

Proof. Let us first construct a contradiction to show $V_{\leq \underline{c}(\delta)} \subseteq \mathcal{B}_\delta + x^*$. Particularly, if $V_{\leq \underline{c}(\delta)} \not\subseteq \mathcal{B}_\delta + x^*$, then for any point $\tilde{x} \in V_{\leq \underline{c}(\delta)} \backslash \{\mathcal{B}_\delta + x^*\}$, $\lim_{t \to \infty} \phi(t, \tilde{x}) = x^*$ and $V(\phi(t, \tilde{x})) < \underline{c}(\delta)$ for all $t > 0$. Therefore, there exists a $\tilde{t} > 0$ such that $V(\phi(\tilde{t}, \tilde{x})) < \underline{c}(\delta)$ and $\phi(\tilde{t}, \tilde{x}) \in \partial\{\mathcal{B}_\delta + x^*\} \subset \tilde{\mathcal{A}}$, which contradicts with the definition of $\underline{c}(\delta)$.

Now, since $V_{\leq \underline{c}(\delta)} \subseteq \mathcal{B}_\delta + x^* \subseteq \mathcal{R} \subseteq V_{\leq c}$, any point $x \in \mathcal{R}$ must have $V(x) \leq c$. Then, it follows from the definition of $a(\delta)$ that after $t \geq \bar{\tau}(\delta)$, the Lyapunov value $V(\phi(t,x)) \leq \underline{c}(\delta)$, and thus $\phi(t,x) \in \mathcal{R}$. $\qquad\square$

2.3 Learning the ROA as a recurrent set

Having laid down the basic theory underlying recurrent sets, we now propose a method to compute inner approximations of the region of attraction $\mathcal{A}(x^*)$ based on checking the recurrence property on finite-length trajectory samples. For concreteness, we consider the following type of sampled trajectories for system (2.1):

$$x_n = \phi(n\tau_s, x_0), \qquad x_0 \in \mathbb{R}^d, \quad n \in \mathbb{N}, \tag{2.6}$$

where $\tau_s > 0$ is the sampling period. In this setting, we will define two notions of discrete-time recurrence w.r.t. a length k trajectory. An illustration that visually highlights these differences can be found in Figure 2-2.

Definition 2.6 (k-Recurrent Set).

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is k-steps recurrent (k-recurrent for short) w.r.t. (2.6), if for any point $x_0 \in \mathcal{R}$ and any step index $n \geq 0$,

$$\exists\, n' \in \{n+1, ..., n+k\}, \ s.t. \ x_{n'} \in \mathcal{R}.$$
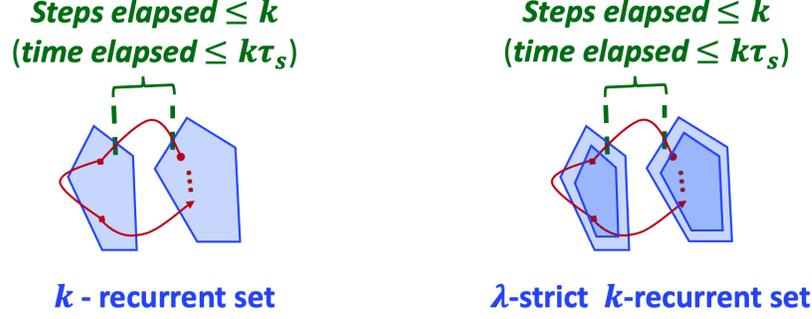
Figure 2-2. Illustration of k-recurrent and λ-strict k-recurrent sets. Note that a k-recurrent is simply a discrete version of a τ-recurrent set. Conversely, a λ-strict k-recurrent set imposes a more stringent requirement for trajectories to reach.

Remark 2. *Note that a set $\mathcal{R}$ being k-recurrent implies that $\mathcal{R}$ is τ-recurrent with $\tau = k\tau_s$. One can then conclude that $\mathcal{R} \subset \mathcal{A}(x^*)$ under the assumptions of Corollary 2.2.*

In our previous work [62], we leverage this k-recurrent property to develop several algorithms that can learn inner approximations of the ROA almost surely, potentially requiring an infinite number of finite-length trajectory samples. To further propose algorithms that require only a finite number of finite-length trajectory samples, we define a more conservative notion referred to here as strict recurrence.

Definition 2.7 (λ-strict k-recurrent Set)**.**
Given any strictness parameter $\lambda \in \mathbb{R}^+$, the set $\mathcal{R}$ is further λ-strict k-recurrent w.r.t. (2.6), if for any $x_0 \in \mathcal{R}$ and any step index $n \geq 0$,

$$\exists\, n' \in \{n+1, ..., n+k\},\ \ s.t.\ x_{n'} \in (\mathcal{R})_{-\lambda}.$$

Remark 3. *In the strict recurrence case, trajectory starting from $p \in \mathcal{R}$ is required to visit not only $\mathcal{R}$, but a more conservative approximation $(\mathcal{R})_{-\lambda}$. Note that if a set $\mathcal{R}$ is λ-strict k-recurrent, then $\mathcal{R}$ is also a k-recurrent set, thus recurrent.*

To ensure one can find such a k-recurrent set or a λ-strict k-recurrent set, we consider again the specific set $\tilde{\mathcal{A}}$ defined in (2.5) that gives the following sufficient conditions for a set $\mathcal{R}$ to be k-recurrent or λ-strict k-recurrent, respectively.

Theorem 2.5. *Let Assumption 2.1 hold, and consider $\delta > 0$, $c \in (0, 1)$ and a compact set $\mathcal{R}$ satisfying: $\mathcal{B}_\delta + x^* \subseteq \mathcal{R} \subseteq V_{\leq c}$. Then $\mathcal{R}$ is τ-recurrent for $k > \bar{k}(\delta) := \bar{\tau}(\delta)/\tau_s$, where $\bar{\tau}(\delta)$ is defined as in Theorem 2.4.*

Proof. Given Theorem 2.4, this result follows directly from $\phi(t, x) \in \mathcal{R}$ for all $x \in \mathcal{R}$ when $t \geq \bar{\tau}(\delta)$. □

Theorem 2.6. *Let Assumption 2.1 hold, and consider $\delta > 0$, $c \in (0, 1)$ and a compact set $\mathcal{R}$ satisfying: $\mathcal{B}_{\delta+\lambda} + x^* \subseteq \mathcal{R} \subseteq V_{\leq c}$. Then $\mathcal{R}$ is λ-strict k-recurrent for any $\lambda \in (0, r - \delta]$ and $k > \bar{k} := \bar{\tau}(\delta)/\tau_s$, where $\bar{\tau}(\delta)$ is defined in Theorem 2.4 and r is the smallest distance between the origin (equilibrium) and the boundary $\partial V_{\leq c}$.*

Proof. See Appendix 2.5 □

In the rest of this section, we assume w.l.o.g. that the asymptotically stable equilibrium is at the origin, i.e., $x^* = 0$. Next, we will outline the mechanism that will be employed to characterize the region of attraction by learning strictly recurrent sets.

2.3.1 Algorithm Summary

We will first initialize a compact ROA approximation $\mathcal{S} \subset \mathbb{R}^d$ satisfying $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$, and then sequentially update it towards a subset of the ROA by computing λ-strict k-recurrent sets $\mathcal{R}$ that seek to satisfy the properties of Theorem 2.6. Precisely, we sample points uniformly from an augmented sampling set $(\mathcal{S})_\eta$ parameterized by a constant $\eta > 0$. Each sample point will be classified, and we update the approximation $\mathcal{S}$ and restart the sampling process if a counter-example is found. This method is illustrated in Algorithm 1.

The rest of this section provides a detailed explanation of each step of the algorithm, as well as a rigorous justification of the proposed methodology.

Algorithm 1: Learning a λ-strict k-recurrent set

Initialize $\mathcal{S}$ according to (2.8)
while $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$ **do**
$\quad$ Generate enough random samples within $(\mathcal{S})_\eta$ according to (2.11)
$\quad$ **if** *observe a counter-example w.r.t* $(\mathcal{S})_\eta$ **then**
$\quad\quad$ | Update $\mathcal{S}$ according to (2.9)
$\quad$ **else**
$\quad\quad$ └ **return** $\mathcal{S}$

Classification of sample points

To classify each point p sampled from $(\mathcal{S})_\eta$, we correspondingly define the set $\mathcal{S}^p$ as

$$\mathcal{S}^p := (\mathcal{S})_d,$$

where d is the signed distance between sample point p and the set $\mathcal{S}$, i.e.,

$$d := \mathrm{sd}(p, \mathcal{S}) = \begin{cases} -d(p, \partial\mathcal{S}) & \textbf{if } p \in \mathcal{S} \\ d(p, \partial\mathcal{S}) & \textbf{o.w.,} \end{cases}$$

Then, we say p is a λ-strict k-steps point if starting from $x_0 = p$,

$$\exists\, n \in \{1, ..., k\},\ \ s.t.\ x_n \in \mathcal{B}_\delta \cup (\mathcal{S}^p)_{-\lambda}. \tag{2.7}$$

That is, the trajectory starting from p either visits the ball $\mathcal{B}_\delta \subset \mathcal{A}(0)$, or certifies the set $\mathcal{S}^p$ is a λ-strict k-recurrent set.

If (2.7) does not hold, we say p is a *counter-example*. We will use such counter-examples to update our current set approximation $\mathcal{S}$.

Construction of set approximations

In order to gradually update the sets $\mathcal{S}$, we consider a parametric family of set approximations known as the sphere approximation. To construct such a sphere approximation, we start by choosing a radius $b > 0$ large enough such that the set

$$\mathcal{S} := \{x \mid \|x\|_2 \le b\} \supseteq \mathcal{B}_{\delta+\lambda}. \tag{2.8}$$

23

Then, given a sample point $p \in (\mathcal{S})_\eta$, we update $\mathcal{S}$ based on the following criterion:

$$p \text{ \textbf{is a counter-example}} \implies b = \|p\|_2 - \varepsilon, \tag{2.9}$$

where $\varepsilon > \eta$ is an algorithm parameter expressing the level of conservativeness in our update.

If the process reaches a value of $b < \delta + \lambda$, we declare the search a failure. At such point, one may choose to either reduce the value of ε or increase the length of the trajectories sampled.

2.3.2 Bound on the number of updates

As mentioned before, the aforementioned search for approximations will fail if $b < \delta + \lambda$ after an update. We will show next that, provided that k and ε are chosen appropriately, there will be no failure. In other words, there will be no counter-examples after a finite number of set updates.

Given $k \geq \bar{k}$, and an arbitrary approximation $\mathcal{S}$ satisfying $\mathcal{B}_{\delta+\lambda} \subseteq (\mathcal{S})_\eta \subseteq V_{\leq c}$, then Theorem 2.6 guarantees that any sample $p \in (\mathcal{S})_\eta$ will lead to a λ-strict k-recurrent trajectory. As a result, the algorithm will stop updating at this point since we cannot find further counter-examples within $(\mathcal{S})_\eta$.

This means that, if it is possible for $(\mathcal{S})_\eta$ to become a subset of $V_{\leq c}$, without violating the condition $\mathcal{B}_{\delta+\lambda} \subseteq \mathcal{S}$, then the algorithm will stop updating and will never fail. The following theorem shows that this is indeed the case, whenever ε and k are properly chosen.

Theorem 2.7. *Let the approximation $\mathcal{S}$ satisfy $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$ and trajectory length $k > \bar{k}$, for $\bar{k}$ as defined in Theorem 2.6. Then, given any counter-example p, the resulting updated approximation will always satisfy $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$ whenever*

$$\varepsilon \leq r - \delta - \lambda, \tag{2.10}$$

24

where r is the smallest distance between the origin (equilibrium) and the boundary $\partial V_{\leq c}$.

Proof. See Appendix 2.5 □

Theorem 2.7 establishes that one can choose parameters k and ε so that the updates of approximation $\mathcal{S}$ never leads to $b < \delta + \lambda$, i.e., the algorithm never fails. However, this requires prior knowledge of $\bar{k}$, r, and δ. We argue that local information on the dynamics can be sufficient to find conservative bounds for r and δ, and thus ε. However, $\bar{k}$ depends in a highly non-trivial way on δ. We solve this issue by, doubling the side of k, i.e. $k^+ = 2k$, every time the failure conditions are met, and re-initializing the approximation $\mathcal{S}$.

In what follows, we use $\mathcal{F}_{\bar{b}}$ to denote the parametric family of closed balls defined by $\{x : ||x||_2 \leq b\}$, for $b \in [0, \bar{b}]$. This leads to the following total bound on the number of iterations.

Theorem 2.8. *Given the initial approximation $\mathcal{S} \in \mathcal{F}_{\bar{b}}$ and initial constant $\bar{b}$ defined in (2.8), the total number of counter-examples encountered in Algorithm 1, with k-doubling after each failure and $\varepsilon > \eta$, is bounded by $\frac{\bar{b}}{\varepsilon - \eta} \log_2 \bar{k}$.*

Proof. See Appendix 2.5 □

Our results provide an upper bound on the number of updates the set approximation may experience by ensuring that $\mathcal{S}$ always contains an $(\delta + \lambda)$-ball around the equilibrium point. However, this is not sufficient to guarantee that $\mathcal{S}$ is k-recurrent, which is required to guarantee that $\mathcal{S} \subseteq V_{\leq c}$. This issue is addressed next.

2.3.3 Sample complexity

By Definition 2.6, a set $\mathcal{S}$ is λ-strict k-recurrent if every point $p \in \mathcal{S}$ satisfies (2.7). As shown before, certifying this property will enable us to guarantee that $\mathcal{S} \subset \mathcal{A}(0)$. However, it is infeasible to enforce condition (2.7) for every point in $\mathcal{S}$. Instead, we

25

will demonstrate that after processing a finite number of samples, our algorithm can identify a set $\mathcal{S}$ satisfying $\text{int } \mathcal{S} \subseteq \mathcal{A}(0)$ with the error probability bounded within a specified error bound $\beta \in (0,1)$.

To achieve this, we first define $\mathcal{S}_{sc}$ as the set that contains all λ-strict k-steps counter-examples within $(\mathcal{S})_\eta$, i.e.,

$$\mathcal{S}_{sc} := \{p \in (\mathcal{S})_\eta \,|\, p \text{ is a counter-example}\}.$$

This implies that for any starting point $p \in \mathcal{S}_{sc} \subseteq (\mathcal{S})_\eta$, the trajectory x_n does not enter $\mathcal{B}_\delta$ nor stay within $(\mathcal{S}^p)_{-\lambda}$ for all $n \in \{1, ..., k\}$.

The following result confirms that there must exist at least one point $p \in \partial \mathcal{S}$ belongs to $\mathcal{S}_{sc}$ whenever $\mathcal{S} \not\subseteq \mathcal{A}(0)$. Consequently, in such instances, $\mathcal{S}_{sc}$ is guaranteed to be non-empty.

Lemma 2.2. *Given any compact set $\mathcal{S}$ satisfying $\partial \mathcal{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \mathcal{S} = \{0\}$, if $\mathcal{S} \not\subseteq \mathcal{A}(0)$, there then exists a point $p \in \partial \mathcal{S}$ such that the solution $\phi(t, p) \notin \text{int}\{\mathcal{S}\}$ for all $t \geq 0$.*

Proof. Since $\mathcal{S} \not\subseteq \mathcal{A}(0)$, Corollary 2.2 implies that $\mathcal{S}$ is not recurrent, i.e., there exists a point $q \in \mathcal{S}$ and a $t' > 0$ such that $\phi(t, q) \notin \mathcal{S}$, $\forall t > t'$. Then, by the continuity of solutions, there exists a point $p \in \partial \mathcal{S}$ such that $\phi(t, p) \notin \text{int}\{\mathcal{S}\} \subseteq \mathcal{S}$ for all $t \geq 0$. Therefore, the result follows. $\square$

Rather than merely establishing that $\mathcal{S}_{sc}$ is non-empty, we will demonstrate that by appropriately selecting η, the volume of $\mathcal{S}_{sc}$ can be lower bounded whenever $\mathcal{S}$ is convex and $\mathcal{S} \not\subseteq \mathcal{A}(0)$.

Theorem 2.9. *let $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$ be a convex compact approximation satisfying $\partial \mathcal{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \mathcal{S} = \{0\}$. If $\mathcal{S} \not\subseteq \mathcal{A}(0)$, then $(\mathcal{S})_\eta$ has λ-strict k-steps counter-example set volume $vol(\mathcal{S}_{sc}) \geq vol(\mathcal{B}_\eta)$ for all $\lambda \in (0, r-\delta)$ and $k \in \{1, 2, ...\}$ whenever*

$$\eta < \min\{\lambda(1 + \exp(Lk\tau_s))^{-1}, r - \delta - \lambda\},$$

where r is defined in Theorem 2.7 and L is the globally Lipschitz constant of the system (2.1).

Proof. See Appendix 2.5 □

We now leverage the findings from Lemma 2.9 to establish a sample complexity bound. Specifically, since the volume of $\mathcal{S}_{sc}$ is lower bounded, the probability of not encountering a counter-example and consequently returning an incorrect result is also bounded. Therefore, by collecting a sufficient number of samples uniformly, we can meet any pre-specified error bound $\beta \in (0, 1)$.

Theorem 2.10. *Given any convex compact approximation $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$ satisfying $\partial \mathcal{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \mathcal{S} = \{0\}$, and consider parameters $\lambda \in (0, r - \delta)$, $\beta \in (0, 1)$, $k > \bar{k}$, and $\eta < \min\{\lambda(1 + \exp(Lk\tau_s))^{-1}, r - \delta - \lambda\}$. Then, whenever we fail to find a λ-strict k-steps counter-example within*

$$m \geq \frac{vol((\mathcal{S})_\eta)}{vol(\mathcal{B}_\eta)} \log \frac{1}{\beta} \tag{2.11}$$

points uniformly sampled from $\mathcal{S}_\eta$, we can terminate sampling and conclude current approximation $\mathcal{S} \subseteq \mathcal{A}(0)$ with error probability bounded by β.

Proof. We generate samples p_j uniformly within $(\mathcal{S})_\eta$, i.e., $p_j \overset{iid}{\sim} U((\mathcal{S})_\eta)$ for all j. Thus, we define the indicator random variable:

$$X_j = \begin{cases} 0 & \textbf{if } p_j \in \mathcal{S}_{sc} \\ 1 & \textbf{o.w.}, \end{cases}$$

as well as a hypothesis test:

$$\begin{cases} H_0 : & \rho < \bar{\rho} \\ H_1 : & \rho \geq \bar{\rho}, \end{cases} \tag{2.12}$$

where $\bar{\rho} \in (0, 1)$ is a preset error threshold for the potentially unknown error ratio $\rho := vol(\mathcal{S}_{sc})/vol((\mathcal{S})_\eta)$. If $m \in \mathbb{N}^+$ points sampled uniformly from $\mathcal{S}$ are all λ-strict,k-recurrent, then hypothesis set (2.12) can be terminated with the probability

27

of type II error (accept H_0 when H_1 is true) bounded by:

$$\mathbb{P}\left(X_1 = \ldots = X_m = 1 \big| H_1\right) \leq \sup_{\rho \geq \bar{\rho}} (1 - \rho)^m = (1 - \bar{\rho})^m.$$

Therefore, for any tolerance $\beta \in (0,1)$, we have the error

$$\mathbb{P}\left(X_1 = \ldots = X_m = 1 \big| H_1\right) \leq \beta, \text{ if } (1 - \bar{\rho})^m \leq \beta,$$

which can be further simplified as the following lower bound:

$$\log \frac{1}{\beta} \leq m \log \frac{1}{1 - \bar{\rho}} \leq m \frac{\bar{\rho}}{1 - \bar{\rho}} \implies m \geq \frac{1}{\bar{\rho}} \log \frac{1}{\beta}.$$

Finally, given parameters $\lambda \in (0, r - \delta)$ and $\eta < \min\{\lambda(1 + \exp(Lk\tau_s))^{-1}, r - \delta - \lambda\}$, Theorem 2.9 implies $\mathcal{S} \subseteq \mathcal{A}(0)$ whenever $\text{vol}(\mathcal{S}_{sc}) < \text{vol}(\mathcal{B}_\eta)$. Thus, we let $\bar{\rho} = \text{vol}(\mathcal{B}_\eta)/\text{vol}((\mathcal{S})_\eta)$ and the result follows. $\qquad\square$

2.3.4 Multiple center point approximation

When the ROA $\mathcal{A}(0)$ is distorted or nonconvex, Algorithm 1 may significantly underestimate $\mathcal{A}(0)$, meaning that the volume of the resulting approximation $\text{vol}(\mathcal{S}) \ll \text{vol}(\mathcal{A}(0))$. To address this problem, we can refine Algorithm 1 by generating additional approximations similar to $\mathcal{S}$ but centered at points different from the equilibrium $x^* = 0$.

In particular, we consider $h \in \mathbb{N}^+$ center points x_i indexed by $i \in \{1, 2, \ldots, h\}$, where the first center point is $x_1 = x^* = 0$. Then other centers, i.e., $x_2, \ldots, x_h$, can be chosen uniformly within some region of interest or selected to be in some preferred place. At each center point x_i the sphere approximation is defined by $\mathcal{S}_i := \{x | \|x - x_i\|_2 \leq b_i\}$, where b_i represents the radius to be updated in the presence of counter-examples. As before we initialize $b_i = \bar{b}$.

Then, the multi-center ROA approximation $\mathcal{S}_{\text{multi}}$ is the union of all approximations, i.e., $\mathcal{S}_{\text{multi}} := \cup_{i=1}^h \mathcal{S}_i$. Note that $\mathcal{S}_1$ is equivalent to the original approximation $\mathcal{S}$ of the previous sections, and $\mathcal{S}_2$ to $\mathcal{S}_h$ are additional enhancements.

Similar to Algorithm 1, sample points are generated uniformly within $\mathcal{S}_{\text{multi},\eta} := \cup_{i=1}^h (\mathcal{S}_i)_\eta$. In this multi-center case, a sample point p is classified as a counter-example if starting from $x_0 = p$,

$$x_n \notin \{\cup_{i=1}^h (\mathcal{S}_i^p)_{-\lambda}\} \cup \mathcal{B}_\delta, \forall n \in \{1, ..., k\},$$

where $\mathcal{S}_i^p := (\mathcal{S}_i)_{-d_i}$ and

$$d_i := \begin{cases} -d(p, \partial \mathcal{S}_i) & \text{if } p \in \mathcal{S}_i \\ d(p, \partial \mathcal{S}_i) & \text{if } p \in (\mathcal{S}_i)_\eta \backslash \mathcal{S}_i \\ 0 & \text{o.w.}, \end{cases}$$

i.e., the signed distances between the point p and the set $\mathcal{S}_i$ whenever $p \in (\mathcal{S}_i)_\eta$ and 0 otherwise.

Once encounter a counter-example, we update $\mathcal{S}_{\text{multi}}$ and restart the sampling process. In particular, given a counter-example $p \in \mathcal{S}_{\text{multi}}$, every approximations $\mathcal{S}_i$ satisfying $p \in \mathcal{S}_i$ are subjected to update respectively via the following criterion:

$$b_i = \|p - x_i\|_2 - \varepsilon.$$

Then, those approximations not containing p will not be updated. Note that the parameter ε is strictly positive. Thus, for all center points $x_i \notin \mathcal{A}(0)$, the corresponding constraint parameters b_i could decrease to negative values and result in $\mathcal{S}_i = \emptyset$ without affecting our results.

In this multi-center setting, we use $\mathcal{F}_{\bar{b}}^h$ to denote the parametric family of h closed balls defined by $\cup_{q=1}^h \mathcal{S}_q$, where $\mathcal{S}_i = \{x : \|x - x_i\|_2 \leq b_i\}$, for $b_i \in [0, \bar{b}]$ and $x_i \in \mathbb{R}^d$ indexed by $i = \{1, ..., h\}$. We aim to demonstrate that the established bounds on the number of updates and the sample complexity continue to apply within this framework.

Theorem 2.11. *The approximation $\mathcal{S}_{\text{multi}}$ is non-vanishing, i.e., $\mathcal{S}_{\text{multi}} \supseteq \mathcal{B}_{\delta+\lambda}$, whenever $k > \bar{k}$ and condition (2.10) is satisfied. The total number of counter-examples encountered, with k-doubling after each failure, is bounded by $h \frac{\bar{b}}{\varepsilon} \log_2 \bar{k}$.*

Proof. See Appendix 2.5 □

Theorem 2.12. *Given any compact multi-center approximation $\mathcal{S}_{\mathrm{multi}} \in \mathcal{F}_{\tilde{b}}^{h}$ satisfying $\partial\mathcal{S}_{\mathrm{multi}} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \mathcal{S}_{\mathrm{multi}} = \{0\}$, and consider parameters $\lambda \in (0, r - \delta)$, $\beta \in (0,1)$, $k > \bar{k}$, and $\eta < \min\{\lambda(1 + \exp(Lk\tau_s))^{-1}, r - \delta - \lambda\}$. If we fail to find a λ-strict k-steps counter-example within*

$$m \geq \underline{m} := \frac{vol(\mathcal{S}_{multi,\eta})}{vol(\mathcal{B}_\eta)} \log \frac{1}{\beta}$$

points uniformly sampled from $\mathcal{S}_{multi,\eta}$, we can terminate sampling and conclude current approximation $\mathcal{S}_{multi} \subseteq \mathcal{A}(0)$ with error probability bounded by β.

Proof. See Appendix 2.5 □

2.4 Numerical Methods

In this section, we will first illustrate the details of our numerical methods based on Theorem 2.12. Precisely, we will leverage the highly parallelizable processing units to process samples simultaneously, approximate the Lipschitz constant locally to relax our sample complexity, and run our algorithm episodically to further increase the learning accuracy. Finally, we validate the accuracy and efficiency of our proposed algorithms through two autonomous dynamic system examples.

Parallel processing sample points

Recall from Theorem 2.12 that we need $\underline{m}$ λ-strict k-steps recurrent samples uniformly taken from $\mathcal{S}_{\mathrm{multi},\eta}$ to bound the error probability by β.

However, it is very computationally costly to sample uniformly from sphere approximations that are potentially interacting with each other. Therefore, we instead sample

$$m_i := \frac{\mathrm{vol}((\mathcal{S}_i)_\eta)}{\mathrm{vol}(\mathcal{B}_\eta)} \log \frac{1}{\beta}$$

points uniformly from each approximation $i \in \{1, .., h\}$. If all of them are λ-strict k-steps recurrent samples, we can terminate sampling and return the resulting approximation since:

$$\sum_{i=1}^{h} m_i = \sum_{i=1}^{h} \frac{\text{vol}((\mathcal{S}_i)_\eta)}{\text{vol}(\mathcal{B}_\eta)} \log \frac{1}{\beta} \geq \underline{m}.$$

When running our algorithm, we use the Torchode toolbox [63] to compute many trajectories in parallel. In particular, we divide samples into N_b batches and sample m_i/N_b points uniformly from each approximation $i \in \{1, .., h\}$ to from a batch. If a batch contains λ-strict k-steps counter-examples, we update $\mathcal{S}_{\text{multi}}$ and restart the sampling process. Otherwise, we can terminate sampling and return the resulting approximation after N_b batches.

Local estimation of the Lipschitz constants:

In order to apply Theorem 2.12, we need to choose

$$\eta < \lambda(1 + \exp(Lk\tau_s))^{-1} \tag{2.13}$$

w.r.t a global consistent Lipschitz constant L. Since the vector field is assumed to be locally Lipschitz, it is also locally one-sided Lipschitz[64, Page 70], i.e., for any point $z \in D$, there exists a neighborhood U_z around z and a constant L_z such that $\forall x, y \in U_z$:

$$(y - x)^T (f(y) - f(x)) \leq L_z \|y - x\|^2.$$

We note that a uniform one-sided Lipschitz constant can be defined under these conditions over any (bounded) subset.

Therefore, we build a uniform grid G that covers our region of interest with a maximum separation $l > 0$ small enough. For each grid cell $g \subseteq G$, we will approximate a local Lipschitz constant L_g for each sample points $p \in g$. Precisely, we simulate the trajectories starting from all sample points $p \in g$ for k steps. Then,

we let L_g be the largest log-norm (defined in [64, p. 76]) evaluated along those trajectory points, i.e.,

$$L_g := \max \text{log-norm}(\phi(p, n\tau_s)),$$

for all sample points $p \in g$ and $n \in \{1, 2, ..., k\}$.

In our algorithm, we fix the choice of parameters η, λ, τ_s, and simulate trajectories starting from all sample points for k steps. We then leverage the trajectory information of all sample points $p \in g$ to solve L_g and further define a step limit:

$$k_g := \frac{1}{L_g \tau_s} \ln \left(\frac{\lambda}{\eta} - 1 \right)$$

within each cell g in our grid. Finally, if a sample point p belongs to a cell g with $k_g \geq k$, we check if it is a λ-strict k-steps recurrent point. If p belongs to a cell g with $k_g < k$, we instead check if p is a k_g-steps recurrent point. Such that the requirement (2.13) will always be satisfied.

Grow the ROA approximation:

When dealing with complex dynamical systems, simulating them over a large time-step horizon k becomes inefficient, and fully covering the region of interest by merely increasing the number of approximations h is not feasible in higher dimensions. To address this, we exploit the recurrent property to add new approximations episodically.

Specifically, after completing our algorithm with h initial approximations, we introduce an additional h distinct initial approximations and rerun our algorithm. This iterative process can be repeated until the accumulated volume of our multi-center point ROA approximation meets our satisfaction. In subsequent runs, we only require sample points from these newly added approximations to confirm their inclusion within the ROA. Nonetheless, we update our approximations whenever we detect a counter-example w.r.t. the union of all existing approximations.

2.4.1 Learning the ROA of a 2D system:

We initiate our analysis by approximating the region of attraction for a 2D autonomous dynamical system, which is governed by the following equations:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}.$$

In Figure 2-3, the black dashed line delineates the boundary of the ROA for the origin, providing a visual benchmark for our approximations.

In our algorithm, we set parameters $\eta = 0.0125$, $\lambda = 0.1$, $\varepsilon = 0.2$, $\tau_s = 0.1s$, and $\delta = 0.3$. The outcomes of our approximation are marked in green. Specificaly, Figure 2-3 compares the results of our multiple center points algorithm for different choices of h (the number of sets used for approximation) and k (the number of time steps considered for recurrence).

It is important to note that the accuracy of our approximation critically depends on the choice of h and k. A smaller h and k limit the likelihood of certifying additional approximations as strict recurrent due to reduced opportunity for trajectories to return within the designated time frame. Specifically, the green line in Figure 2-3 illustrates the boundary of the recurrent area. This boundary is determined by systematically testing a mesh grid of points; any point from which the trajectory fails to visit any $\mathcal{S}_i$ for $i = 1, ..., h$ within k steps is considered to be outside of this region.

To extend the ROA approximation, we strategically increase the number of approximations by including $h = 50$ additional approximations in each iteration and repeat this process for 200 episodes. This episodic enhancement allows us to progressively refine our ROA approximation. The impacts of different time-step horizons k on the accuracy and efficiency of our algorithm are depicted in Figure 2-4.
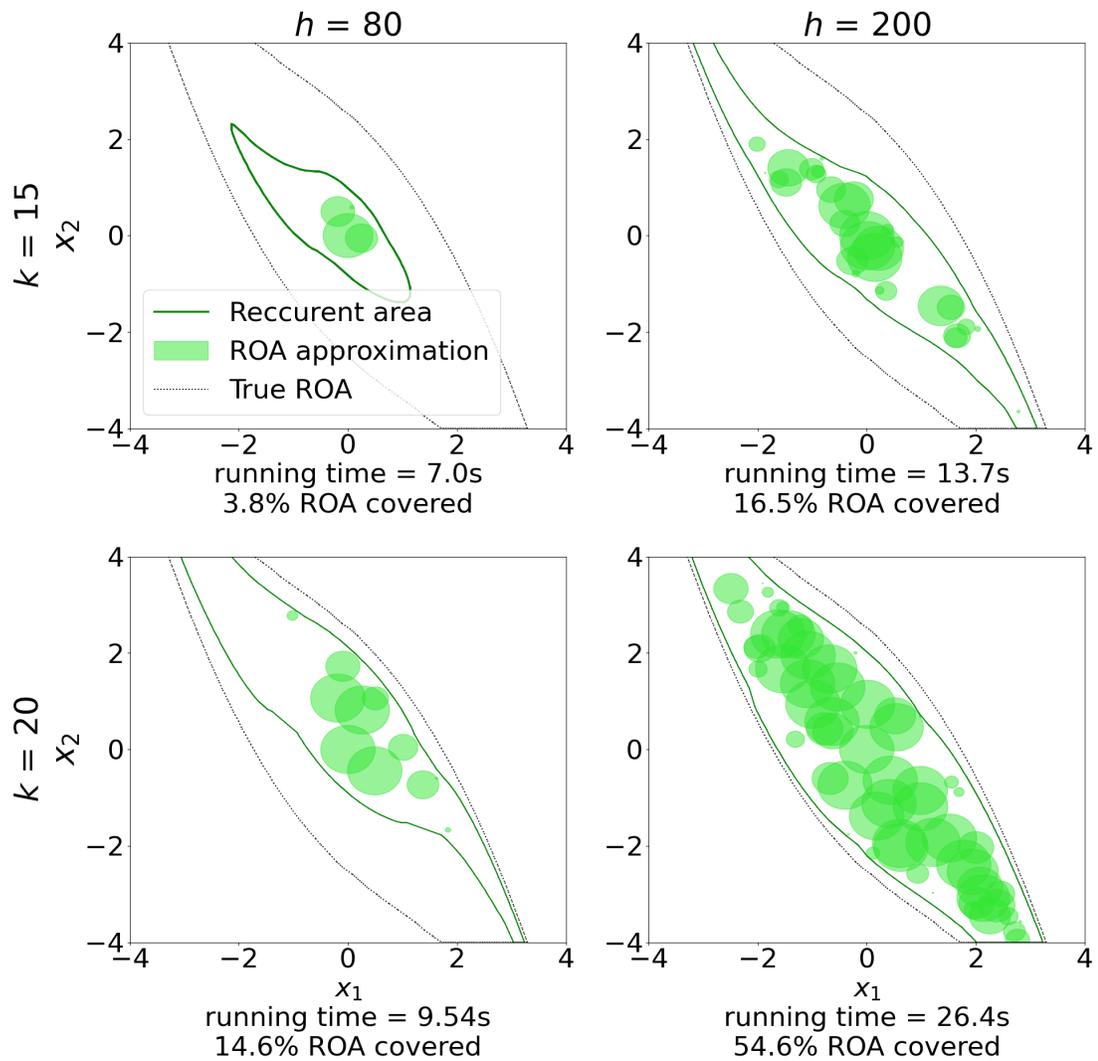
Figure 2-3. Our ROA approximation outcomes under different settings. The black dashed line represents the true boundary of the ROA. The union of all green circles illustrates our computed approximation of the ROA. The solid green line delineates the boundary of the recurrent area, showing the limits within which trajectories are guaranteed to return.
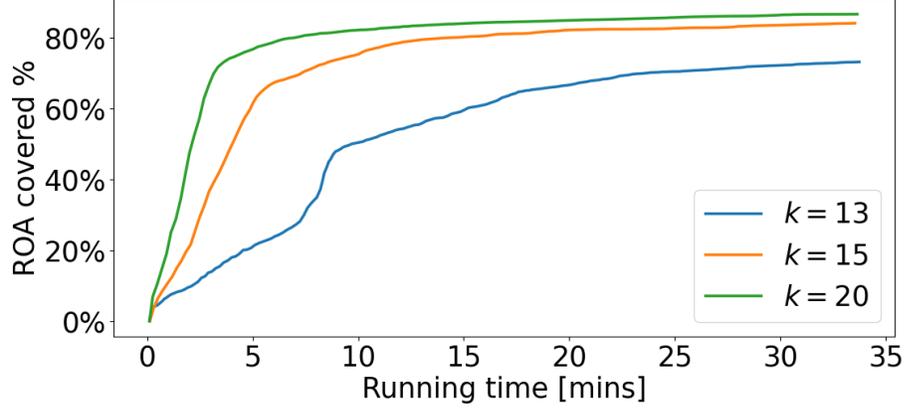
Figure 2-4. The growth of the ROA approximation in 200 episodes w.r.t. different time-step horizon k.

2.4.2 Learning the ROA of a 4D transient stability problem:

We then study this 4D power system transient stability problem introduced in [65]:

$$
\begin{bmatrix} \dot{\delta}_1 \\ \dot{\omega}_1 \\ \dot{\delta}_2 \\ \dot{\omega}_2 \end{bmatrix} = \begin{bmatrix} \omega_1 \\ -\sin\left(\delta_1\right) - 0.5\sin\left(\delta_1 - \delta_2\right) - 0.4\omega_1 \\ \omega_2 \\ -0.5\sin\left(\delta_2\right) - 0.5\sin\left(\delta_2 - \delta_1\right) - 0.5\omega_2 + 0.05 \end{bmatrix}.
$$

In this problem, we set parameters $\eta = 0.025$, $\lambda = 0.2$, $\varepsilon = 0.15$, $\tau_s = 0.5s$, $\delta = 2$, and $k = 10$. Figure 2-6 shows the resulting ROA approximation and recurrent area volume by episodically including $h = 20$ additional approximations and growing our ROA approximation for 80 episodes.

The green line in Figure 2-6 demonstrates the accuracy of the Sum-of-Squares (SOS) method as introduced in [65]. Notably, our method surpasses the SOS-based results in accuracy after just ten episodes. Furthermore, as depicted in Figure 2-5, the precision of our approximation method continues to improve, ultimately reaching an accuracy of 58%.
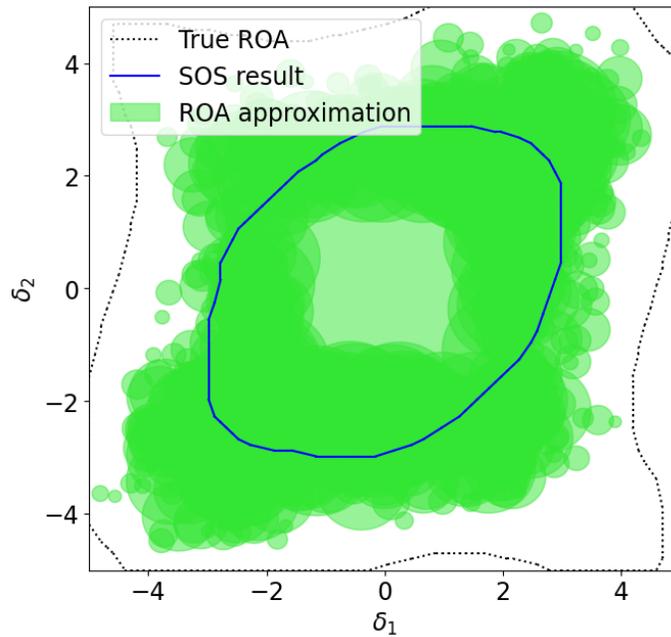
Figure 2-5. Our ROA approximation results after 80 episodes. The black dashed line depicts the true ROA boundary. The union of all green circles is the projection of our ROA approximation.
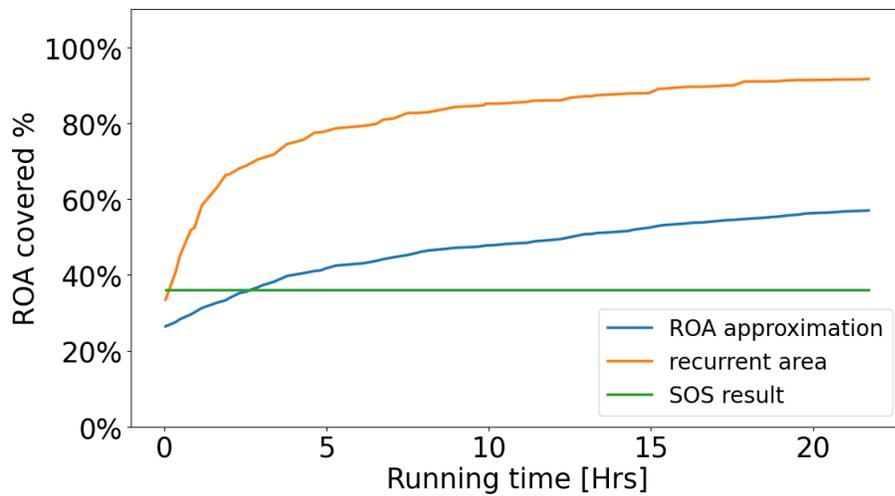
Figure 2-6. The growth of our ROA approximation and the associated recurrent area in 80 episodes.

2.5 Appendix

Proof of Corollary 2.1

Let V be the Zubov's function whose existence is guaranteed by Theorem 2.1. Thus by the definition of V, for $c \in (0,1)$, $V_{\leq c} \subseteq \mathcal{A}(x^*)$. Further from (2.4), it follows that $(\mathcal{L}_f V)(x) \leq 0$, for $x \in V_{\leq c} \subset \mathcal{A}(x^*)$. Thus, $V_{\leq c}$ is positively invariant.

To prove the $V_{\leq c}$ is contractible, we need to provide a continuous mapping $H : [0,1] \times V_{\leq c} \to V_{\leq c}$ such that $H(0,x) = x$ and $H(1,x) = x^*$ for all $x \in V_{\leq c}$. Similar to [59], we define $H(s,x) := \phi(\frac{s}{1-s}, x)$ for $s < 1$, and $H(1,x) \equiv x^*$. Note that H is continuous in s and x for $s < 1$, as in [3]. We are thus left to prove continuity at each $(1,x)$. To do so, we take any such x and pick any open neighborhood $\mathcal{V}$ of $H(1,x) = x^*$. By Assumption 2.1 as well as the definition of asymptotic stability, it follows that there exists another open neighborhood $\mathcal{W} \subseteq \mathcal{A}(x^*)$ of x^* for which all trajectories starting in $\mathcal{W}$ remain in $\mathcal{V}$, i.e., $\phi(t, x_0) \in \mathcal{V}$ for all $x_0 \in \mathcal{W}$ and $t > 0$. Given $V_{\leq c} \subseteq \mathcal{A}(x^*)$, any point $x \in V_{\leq c}$ satisfies $\phi(T, x) \in \mathcal{W}$ for some $T > 0$. This, together with the continuity of $\phi(T, \cdot)$, implies that there is a neighborhood $\mathcal{V}' \subseteq V_{\leq c}$ of x such that $\phi(T, y) \in \mathcal{W}$ for all $y \in \mathcal{V}'$, which let us conclude:

$$H(s,y) \in \mathcal{V} \quad \text{whenever } y \in \mathcal{V}' \text{ and } s > 1 - \frac{1}{T+1}$$

and continuity follows since $\mathcal{V}$ could be made arbitrarily small.

Proof of Lemma 2.1

We will prove this statement by contradiction. Assume the result does not hold, i.e., there exists $x_0 \in \mathcal{R}$ s.t. for any $t > 0$ there exists a $t' \geq t$ such that $\phi(t', x_0) \notin \mathcal{R}$. This, together with the definition of the recurrent set (Definition 2.4) and the continuity of the solution, implies there exists a $t'' \geq t$ such that $\phi(t'', x_0) \in \partial \mathcal{R}$ for any $t > 0$. Therefore, we can construct an infinite sequence $\{x_n\}_{n=0}^{\infty}$ that lies within $\partial \mathcal{R}$, i.e., $\{x_n\}_{n=0}^{\infty} \subset \partial \mathcal{R}$.

Precisely, let $t_0 \geq 0$ be a time such that $\phi(t_0, x_0) \in \partial\mathcal{R}$. Then, given $x_n := \phi(t_n, x_0) \in \partial\mathcal{R}$ and some fixed time interval $\tau > 0$, we defined t_{n+1} as the first time since $t_n + \tau$ that the solution $x_{n+1} := \phi(t_{n+1}, x_0)$ lies within $\partial\mathcal{R}$, i.e., $\phi(t_{n+1}, x_0) \in \partial\mathcal{R}$ and $\phi(t, x_0) \notin \partial\mathcal{R}$ for all $t \in [t_n + \tau, t_{n+1})$.

Then, since $\partial\mathcal{R}$ is compact, by Bolzano-Weierstrass theorem, $\{x_n\}_{n=0}^{\infty}$ must have a sub-sequence $\{x_{n_i}\}_{i=1}^{\infty}$ that converges to an accumulation point $\bar{x} \in \partial\mathcal{R}$. It follows then from the definition of ω-limit sets (Definition 2.1) that $\bar{x} = \lim_{i\to\infty} x_{n_i} \in \Omega(f) \cap \partial\mathcal{R}$, which contradicts with the assumption that $\partial\mathcal{R} \cap \Omega(f) = \emptyset$.

Proof of Theorem 2.2

($\implies$): If $\mathcal{R}$ is a compact recurrent set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$, Lemma 2.1 implies that for any point $x_0 \in \mathcal{R}$, there exists a time $T > 0$ such that $\phi(t, x_0) \in \mathcal{R}, \forall t \geq T$, i.e., the solution is bounded in the compact set $\mathcal{R}$ for all $t \geq T$. It then follows from [3, p. 127] that the limit set $\Omega(x_0) \neq \emptyset$ and $\lim_{t\to\infty} d(\phi(t, x_0), \Omega(x_0)) = 0$. Therefore, we conclude $\Omega(f) \cap \mathcal{R} \supseteq \Omega(x_0) \neq \emptyset$ and $x_0 \in \mathcal{A}(\Omega(f) \cap \mathcal{R})$. Finally, since x_0 was chosen arbitrarily within $\mathcal{R}$, it follows that $\mathcal{R} \subset \mathcal{A}(\Omega(f) \cap \mathcal{R})$.

($\impliedby$): By assumption $\Omega(f) \cap \mathcal{R} \subset \text{int } \mathcal{R}$. Therefore, we can always construct an open ζ-neighborhood $\Omega_\zeta^{\mathcal{R}} := \{x \in \mathbb{R}^d | d(x, \Omega(f) \cap \mathcal{R}) < \zeta\}$ of $\Omega(f) \cap \mathcal{R}$ for some $\zeta > 0$ small enough such that $\Omega_\zeta^{\mathcal{R}} \subset \text{int } \mathcal{R}$.

Then for any point $x_0 \in \mathcal{R}$, by the assumption that $\mathcal{R} \subset \mathcal{A}(\Omega(f) \cap \mathcal{R})$, the solution $\phi(t, x_0)$ converges to $\Omega(f) \cap \mathcal{R}$, i.e., $\liminf_{t\to\infty} d(\phi(t, x_0), \Omega(f) \cap \mathcal{R}) = 0$. It follows then that for any $\zeta > 0$ and time $t > 0$, there always exists some time $t' \geq t$ such that $d(\phi(t', x_0), \Omega(f) \cap \mathcal{R}) < \zeta$, and thus $\phi(t', x_0) \in \Omega_\zeta^{\mathcal{R}} \subset \mathcal{R}$. Therefore, $\mathcal{R}$ is recurrent.

Proof of Theorem 2.3

The proof of the theorem relies on Zubov's existence criterion stated in Theorem 2.1. Given $\mathcal{R}$, let us now define

$$\underline{c} := \min_{x \in \partial \mathcal{R}} V(x), \quad \bar{c} := \max_{x \in \partial \mathcal{R}} V(x),$$
$$\text{and} \quad a := \max_{x \in C} \nabla V(x)^T f(x),$$

where $C = \{x \in \mathbb{R}^d : \underline{c} \le V(x) \le \bar{c}\}$ is compact.

We first argue that $V_{\le \underline{c}} := \{x : V(x) \le \underline{c}\} \subseteq \mathcal{R}$. Let $\underline{x}$ be the point in $\partial \mathcal{R}$ that achieves the minimum, i.e, $V(\underline{x}) = \underline{c}$. Since $\mathcal{R}$ is not necessarily connected, we use $\mathcal{R}'$ to denote the connected component of $\mathcal{R}$ containing $\underline{x}$. Note that $x^* \in \text{int}\,\mathcal{R}$ must be contained in $\mathcal{R}'$, since otherwise, the trajectory $\phi(t, \underline{x})$, which strictly decreases V must eventually find a point $x' \in \partial R$ with $V(x') < \underline{c}$; which contradicts the definition of $\underline{c}$, see Fig 2-7. Thus, $x^* \in \mathcal{R}' \subseteq \mathcal{R}$.
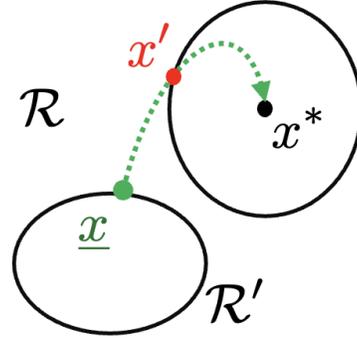
Suppose then that $V_{\le \underline{c}} \not\subseteq \mathcal{R}' \subseteq \mathcal{R}$, for any point $\tilde{x} \in V_{\le \underline{c}} \backslash \mathcal{R}'$, $V(\phi(t, \tilde{x})) < \underline{c}$, for $t > 0$, and $\lim_{t \to \infty} \phi(t, \tilde{x}) = x^*$. Thus there exists $\tilde{t} > 0$ s.t. $V(\phi(\tilde{t}, \tilde{x})) < \underline{c}$ and $\phi(\tilde{t}, \tilde{x}) \in \partial \mathcal{R}$; which contradicts again with the definition of $\underline{c}$. It follow then that $V_{\le \underline{c}} \subseteq \mathcal{R}' \subseteq \mathcal{R}$.

Figure 2-7. An visualization of the proof of Theorem 2.3. Particularly, when $\mathcal{R}$ is disconnected, the equilibrium x^* must be contained in $\mathcal{R}'$, since otherwise, one could find a point $x' \in \partial \mathcal{R}$ with $V(x') < V(\underline{x})$ along the trajectory $\phi(t, \underline{x})$ that strictly decreases V.

Similarly, since the contradictable set $V_{\le \bar{c}}$ contains every point in the boundary of $\mathcal{R}$, there cannot be any point in $x \in \mathcal{R}$ with $V(x) > \bar{c}$. We therefore get that the

following inclusions must hold:

$$V_{\leq \underline{c}} \subseteq \mathcal{R} \subseteq V_{\leq \bar{c}}. \tag{2.14}$$

Finally, by (2.14), for any point $x \in \mathcal{R}$ we must have $V(x) \leq \bar{c}$. Since the time derivative of $V(x)$ is at most $a < 0$, it follows that after $t \geq \bar{\tau} := \frac{\underline{c} - \bar{c}}{a}$ the Lyapunov value $V(\phi(t, x)) \leq \underline{c}$, which implies that $\phi(t, x) \in \mathcal{R}$ and result follows.

Proof of Theorem 2.6

Note first that r is the smallest distance between the origin (equilibrium) and the boundary $\partial V_{\leq c}$, hence if $\lambda \leq r - \delta$, then $\mathcal{B}_{\delta+\lambda} \subseteq V_{\leq c}$ and such a $\mathcal{R}$ always exists. Then, we further conclude $(\mathcal{R})_{-\lambda} \supseteq \mathcal{B}_\delta$ since $\mathcal{R} \supseteq \mathcal{B}_{\delta+\lambda}$. Finally, it follows from Theorem 2.4 that after $k > \bar{k} = \bar{\tau}(\delta)/\tau_s$ steps, trajectory $\phi(t, x) \in \mathcal{B}_\delta \subseteq (\mathcal{R})_{-\lambda}$ for any point $x \in \mathcal{R}$, and thus results follow.

Proof of Theorem 2.7

Given any counter-example p w.r.t $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$, one can conclude $p \notin V_{\leq c}$ from Theorem 2.6; since otherwise, p would generate a λ-strict k-recurrent trajectory. It then follows from the definition of r that $\|p\|_2 \geq r$.

Now, since $\|p\|_2 \geq r$, the update leads to $b = \|p\|_2 - \varepsilon \geq r - \varepsilon \geq \delta + \lambda$ whenever $\varepsilon \leq r - \delta - \lambda$. This, together with the fact that $\mathcal{S} \supseteq \mathcal{B}_{\delta+\lambda}$, implies the result.

Proof of Theorem 2.8

Note that once a counter-example is encountered, we decrease the radius constraint by at least $\varepsilon - \eta$. Therefore, our approximation $\mathcal{S} \in \mathcal{F}_{\bar{b}}$ after every updates. Then, for any fixed k, our method can find at most $\bar{b}/(\varepsilon - \eta)$ counter-examples without failing. Since it takes at most $\log_2 \bar{k}$ updates on k to find some $k \geq \bar{k}$ using the doubling method, the result follows.

Proof of Theorem 2.9

The proof of Theorem 2.9 relies on the following properties of the set $(\mathcal{S})_\eta$.

Lemma 2.3. *Given any compact convex set $\mathcal{S}$, then for all points $q \in \partial \mathcal{S}$, the distance $d(q, \partial(\mathcal{S})_\eta) = \eta$. Moreover, $d(p, \partial(\mathcal{S})_\eta) < \eta$ whenever $p \in (\mathcal{S})_\eta \backslash \mathcal{S}$.*

Lemma 2.4. *Given any compact convex set $\mathcal{S}$ and any sample point $p \in (\mathcal{S})_\eta$, then $(\mathcal{S}^p)_{-\lambda}$ is a subset of $\mathcal{S}$ whenever $\lambda > \eta$.*

Corollary 2.3. *Given any compact convex set $\mathcal{S}$, any sample point $p \in (\mathcal{S})_\eta$, and any parameter $\zeta > 0$, then $(\mathcal{S}^p)_{-\lambda}$ is a subset of $(\mathcal{S})_{-\zeta}$ whenever $\lambda > \eta + \zeta$.*

Note that Lemma 2.2 guarantees the existence of a point $p \in \partial \mathcal{S}$ with solution $\phi(t, p) \notin \text{int}\{\mathcal{S}\}$ for all $t \geq 0$. We then use $\mathcal{B}(p, \eta)$ to denote a closed ball with radius η centered at this point p, i.e., $\mathcal{B}(p, \eta) = \{x | \|x - p\|_2 \leq \eta\}$. Then, for all points $q \in \mathcal{B}(p, \eta) \subseteq (\mathcal{S})_\eta$, we will show that by choosing η appropriately, q will be classified as a λ-strict counter-example, i.e., $\phi(q, n\tau_s) \notin (\mathcal{S}^q)_{-\lambda}$ and $\phi(q, n\tau_s) \notin \mathcal{B}_\delta$ for all $n \in \{1, 2, ...\}$.

Towards this goal, we first argue that it is sufficient for us to just show $\phi(q, n\tau_s) \notin (\mathcal{S}^q)_{-\lambda}$ for all $n \in \{1, 2, ...\}$ since $\eta < r - \delta - \lambda$ implies the distance $d(\partial(\mathcal{S}^q)_{-\lambda}, 0) \geq r - \eta - \lambda > \delta$, and thus $(\mathcal{S}^q)_{-\lambda} \supseteq \mathcal{B}_\delta$.

Now recall the assumption that the dynamical system (2.1) is globally L-Lipschitz, it follows from [3, p. 96] that the distance between solutions $\|\phi(t, p) - \phi(t, q)\| \leq \|p - q\| \exp(Lt) \leq \eta \exp(Lt)$ for all $t \geq 0$. This, together with the fact that $\phi(t, p) \notin \text{int}\{\mathcal{S}\}$, implies $\phi(t, q) \notin (\mathcal{S})_{-\zeta}$ for all $t \geq 0$ if $\zeta > \eta \exp(Lt) = \eta \exp(Lk\tau_s)$.

Finally, Corollary 2.3 further implies $(\mathcal{S}^q)_{-\lambda} \subseteq (\mathcal{S})_{-\zeta}$ whenever $\lambda > \eta + \zeta > \eta(1 + \exp(Lk\tau_s))$. Therefore, by choosing $\eta < \lambda(1 + \exp(Lk\tau_s))^{-1}$, every point $q \in \mathcal{B}(p, \eta)$ will be classified as a λ-strict counter-example, and thus the λ-strict k-steps counter-example set volume $\text{vol}(\mathcal{S}_{sc}) \geq \text{vol}(\mathcal{B}_\eta)$.

Proof of Theorem 2.11

By definition $\mathcal{S}_{\text{multi}} \supseteq \mathcal{S}_1$ for all $i \in \mathbb{N}^+$, Theorem 2.7 therefore implies $\mathcal{S}_{\text{multi}} \supseteq \mathcal{S}_1 \supseteq \mathcal{B}_\delta$ under (2.10). The bound on the total number of counter-examples follows as in Theorem 2.8, since every additional approximation $\mathcal{S}_i \in \mathcal{F}_{\bar{b}}$ for all $i \in \{1, ..., h\}$.

Proof of Theorem 2.12

Since $\mathcal{S}_{multi} \in \mathcal{F}_{\bar{b}}^h$ is the union of convex compact sets with $\mathcal{S}_1 \supseteq \mathcal{B}_{\delta+\lambda}$, we can generalize the proof of Theorem 2.9-2.10 to this case.

Similar to the proof of Theorem 2.9, it follows from Lemma 2.2 that there exists a point $p \in \partial \mathcal{S}_{multi}$ with solution $\phi(t, p) \notin \text{int} \, \mathcal{S}_{multi}$ for all $t \geq 0$. There further exists a ball $\mathcal{B}(p, \eta)$ such that $\phi(t, q) \notin (\mathcal{S}_{multi})_{-\zeta}$ for all $q \in \mathcal{B}(p, \eta)$ and $t \geq 0$, whenever $\zeta > \eta \exp\left(Lk\tau_s\right)$.

Then, by choosing a $\lambda > \eta + \zeta > \eta(1 + \exp\left(Lk\tau_s\right))$, Corollary 2.3 implies $(\mathcal{S}_i^q)_{-\lambda} \subseteq (\mathcal{S}_i)_{-\zeta}$ for all $i = 1, ..., h$. This, together with the fact that $\{\cup_{i=0}^h (\mathcal{S}_i)_{-\zeta}\} \subseteq (\mathcal{S}_{multi})_{-\zeta}$, implies

$$\{\cup_{i=1}^h (\mathcal{S}_i^q)_{-\lambda}\} \subseteq \{\cup_{i=1}^h (\mathcal{S}_i)_{-\zeta}\} \subseteq (\mathcal{S}_{multi})_{-\zeta}.$$

Now, since the trajectory $\phi(t, q) \notin (\mathcal{S}_{multi})_{-\zeta}$, we conclude $\phi(t, q) \notin \{\cup_{i=1}^h (\mathcal{S}_i^q)_{-\lambda}\}$.

Since the other condition $\eta < r - \delta - \lambda$ implies $\mathcal{B}_\delta \subseteq \{\cup_{i=1}^h (\mathcal{S}_i^q)_{-\lambda}\}$, we further conclude $\phi(t, q) \notin \mathcal{B}_\delta$. Thus, $\mathcal{S}_{\text{multi}}$ has λ-strict k-steps counter-example set volume $\text{vol}(\mathcal{S}_{sc}) \geq \text{vol}(\mathcal{B}_\eta)$.

Finally, by constructing the same hypothesis test as defined in the proof of Theorem 2.10 with parameter $\bar{\rho} = \text{vol}(\mathcal{B}_\eta)/\text{vol}((\mathcal{S}_{\text{multi}})_\eta)$, result follows.

2.6 Conclusion

In this chapter, we introduce a novel methodology for learning the region of attraction of an asymptotically stable equilibrium point within a model-free context, utilizing only sampled system trajectories.

We begin by advocating for a more flexible notion of invariance, termed recurrence. Utilizing this concept, we demonstrate that under mild conditions, a compact set containing an asymptotically stable equilibrium point is a subset of the ROA if and only if it is recurrent. To facilitate the practical learning of an ROA inner approximation, we establish that there exists a time horizon k within which almost every compact ROA subset exhibits recurrence. This finding not only highlights a regularity condition necessary for ROA inner approximations but also enables their characterization using only finite-length trajectory samples.

Building on this theoretical foundation, we propose a practical algorithm designed to characterize the region of attraction for a given asymptotically stable equilibrium point. This algorithm is optimized for efficiency, requiring only a finite number of finite-length trajectory samples and updates. It also takes full advantage of highly parallelizable processing units, allowing for simultaneous processing of trajectory samples.

Moreover, the adaptability of our approach is demonstrated by its capacity for incremental improvement: even after the algorithm has terminated, it can be warm-started with additional sets to further enhance the accuracy of the approximation.

The chapter concludes with the presentation of two numerical examples that illustrate the effectiveness of our method.

Chapter 3

Model-free Learning of Safe Regions via Recurrent Barrier Functions

Having established a practical method to characterize the region of attraction for a stable equilibrium point, we now shift our focus to identifying the set of initial conditions such that a given unsafe state space region is guaranteed to be avoided.

Specifically, in Section 3.1 we systematically relax the classic differential barrier conditions into integral conditions and further into recurrent conditions. These adaptations simplify the searching process of a barrier function that can be utilized to confirm safety. We further extended our analysis in Section 3.2 to a practical safety objective and developed in Section 3.3 an algorithm that uses only a finite number of finite-length trajectory samples to characterize safety regions and analyze safety levels.

Notation

Given a set S and an arbitrary norm (denoted by $\| \cdot \|$), we use $\mathrm{sd}(x, S)$ to denote the signed distance between a point x and S, i.e.,

$$\mathrm{sd}(x, S) := \begin{cases} \inf_{y \in \partial S} \|y - x\| & \text{if } x \notin S \\ -\inf_{y \in \partial S} \|y - x\| & \text{if } x \in S. \end{cases}$$

We further use $\mathrm{P}_S(x)$ to denote the set of projections of a point x on a non-empty

closed set S, i.e.,

$$\mathrm{P}_S(x) := \arg\min_{y \in S} \|y - x\|.$$

We respectively use $[x]_+$ and $[x]_-$ to denote the projection of a number $x \in \mathbb{R}$ on the set of all non-negative real numbers and the set of all non-positive real numbers, respectively. For a continuous function h, we use $h_{<0}$ ($h_{\leq 0}$), $h_{>0}$ ($h_{\geq 0}$), and $h_{=0}$ to denote the open (closed) sub-level set, open (closed) super-level set and closed level set, respectively.

For a vector field that is assumed to be locally Lipschitz, it is also locally one-sided Lipschitz [64, Page 70], i.e., for any point $z \in D$, there exists a neighborhood U_z around z and a constant L_z such that $\forall x, y \in U_z$:

$$(y - x)^T (f(y) - f(x)) \leq L_z \|y - x\|^2$$

We note that a uniform one-sided Lipschitz constant can be defined under these conditions over any (bounded) subset of D. In what follows, we will use:

$$F(S) := \sup_{z \in S} \|f(z)\|; \quad L(S) := \sup_{z \in S} L_z,$$

for a set $S \subseteq D$.

3.1 Generalized Barrier Functions

Barrier functions constitute an effective tool for assessing and enforcing safety-critical constraints on dynamical systems. To this end, one is required to find a function h that satisfies a Lyapunov-like differential condition, thereby ensuring the invariance of its zero super-level set $h_{\geq 0}$. This methodology, however, does not prescribe a general method for finding the function h that satisfies such differential conditions, which, in general, can be a daunting task. In this section, we seek to overcome this limitation by developing a generalized barrier condition that makes the search for h easier.

The section is structured as follows: Section 3.1.1 introduces preliminary definitions pertinent to dynamical systems and safety and revisits classical barrier conditions that will be utilized throughout this work. In Section 3.1.2, we detail the development of integral-based barrier functions. This discussion is expanded in Section 3.1.3, where we introduce and elaborate on the recurrence conditions. Under these conditions, one can certify that a super-level set of h is τ-recurrent, imposing regularity conditions on the trajectories. In Section 3.1.4, we demonstrate the generality of the recurrence conditions, specifically showing that almost every safe set would satisfy our recurrent barrier conditions with a bounded time horizon τ. In Section 3.1.5, we argue that recurrent sets are functionally equivalent to invariant sets and, thus, can be indicative of safety.

3.1.1 Preliminary results

Throughout the chapter, we consider a continuous-time dynamical system described using the following ordinary differential equation:

$$\dot{x} = f(x), \tag{3.1}$$

where $x \in D \subset \mathbb{R}^n$ is the state, and the map $f : D \to \mathbb{R}^n$ is a continuous and locally Lipschitz function defined over a domain D. Given an initial state x_0, we use $\phi(t, x_0)$ to denote the solution of (3.1). We assume system 3.1 is forward complete, which is specified in the following assumption.

Assumption 3.1 (Forward Completeness). *For any $x \in D$, the trajectory $\phi(\cdot, x)$ is defined for all $t \in [0, \infty)$.*

In the following, we formally define the notions of safety and invariance.

Definition 3.1 (Safety). *Let $\mathcal{X}_u \subseteq D$ be a set of unsafe states, a trajectory $\phi(t, x_0)$ of (3.1) is unsafe if there exists a time $t \geq 0$ such that $\phi(t, x_0) \cap \mathcal{X}_u \neq \emptyset$.*

We say that a set $\mathcal{X}_s \subseteq D$ is a safe state space region if there are no unsafe trajectories starting from $\mathcal{X}_s$.

Note that the notion of invariance defined in Definition 2.3 is closely related to the barrier function methods that characterize safe state space regions. By trapping trajectories on level sets of a function, one can ensure the system's safety whenever its initial state belongs to an invariant set S that does not intersect with $\mathcal{X}_u$.

Next, we review some classic formulations of barrier functions that certify the invariance of their super-level sets. These formulations require the barrier functions to be differentiable, which we will relax in the following sections. We start with the most basic formulation: Nagumo's. It only requires the time derivative of the barrier function h, which is also its Lie derivative along f, to be non-negative at any state of its zero level set $h_{=0}$. That implies that whenever a trajectory reaches the boundary of that set $h_{\geq 0}$ from its interior, h must not decrease. Thus, the trajectory must remain in the set, ensuring the latter's invariance.

Theorem 3.1 (Nagumo's Barrier Functions). *[66, Th 3.1]*

Consider a dynamical system (2.1) and a differentiable function $h : D \subset \mathbb{R}^n \to \mathbb{R}$, then h is a Nagumo's Barrier Function (NBF) satisfying:

$$L_f h(x) := \lim_{t \to 0} \frac{h(\phi(t, x)) - h(x)}{t} \geq 0, \quad \forall x \in h_{=0}, \tag{3.2}$$

if and only if the super-level set $h_{\geq 0}$ is invariant.

The following theorem adds an additional constraint to Nagumo's definition that further lower-bounds the rates at which h can at most decrease along the trajectories starting from the interior of its super-level set and the least at which it should increase along the trajectories starting from the exterior of that set. That ensures safe trajectories approach the boundary slowly, if at all, and possibly unsafe ones converge to the safe set fast enough, ensuring its stability. The latter is not guaranteed by Nagumo's version.

Definition 3.2 (Extended class $\mathcal{K}$ function). *[34, Def 2]*

A continuous function $\zeta := (-b, a) \to (-\infty, \infty)$ is said to belong to extended class $\mathcal{K}$ for some $a, b > 0$ if it is strictly increasing and $\zeta(0) = 0$.

Theorem 3.2 (Zeroing Barrier Functions). *[34, Prop 1]*

Consider a differentiable function $h : D \to \mathbb{R}$ and an extended class $\mathcal{K}$ function ζ. Assume there exists a super-level set $D_0 := h_{\geq -c} \subseteq D$ for some $c \geq 0$ such that:

$$L_f h(x) \geq -\zeta(h(x)), \quad \forall x \in D_0, \tag{3.3}$$

then:

(i) *h is called a Zeroing Barrier Functions (ZBF), and the super-level set $h_{\geq 0}$ is invariant.*

(ii) *whenever $x \in h_{<0} \cap D_0$, then as long as $h(\phi(t, x)) < 0$, $h(\phi(t, x))$ must monotonically increase to zero, at least, with a positive rate of $-\zeta(h(\phi(t, x)))$, and*

(iii) *whenever $h(\phi(t, x)) > 0$, then $h(\phi(t, x))$ may decrease to zero, at most, with a negative rate of $-\zeta(h(\phi(t, x)))$.*

The bounds on the boundary-approaching rates of Zeroing barrier functions in Theorem 3.2 are non-uniform and state-dependent. In the following remark and the rest of the section, we describe the special case when ζ is a (piece-wise) linear function of $h(x)$, which results in uniform exponential bounds on the evolution of $h(\phi(t, x))$.

Remark 4. *[34, Remark 6] A special case of (3.3) is:*

$$L_f h(x) \geq -\alpha h(x), \ \forall x \in D_0, \tag{3.4}$$

for some $\alpha > 0$. Since $\zeta(s) = \alpha s$ is an extended class $\mathcal{K}$ function, the super-level set $h_{\geq 0}$ is invariant as stated in Theorem 3.2. This formulation is commonly used since it leads to a convex problem that can be efficiently solved using techniques like SoS programming [31, 67].

The following lemma provides the formal statement on the exponential bounds on the convergence rates of the function value $h(\phi(t, x))$ under condition (3.4).The proof is analogous to [31, Thm 1], with an extension to consider all states $x \in h_{\leq -c}$ instead of just $h_{\leq 0}$.

Lemma 3.1. *Consider a continuous function $h : D \to \mathbb{R}$, an $\alpha > 0$, and a super-level set $D_0 := h_{\geq -c}$ for some $c \geq 0$, then condition (3.4) implies the following exponential convergence result:*

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x), \quad \forall t \geq 0, \, x \in D_0.$$

Proof. See Appendix 3.4 □

We end this section by generalizing the case in Remark 4 to one where the bounds on the different sides of the boundary differ, calling the resulting functions exponential barrier functions (EBFs). The need for this generalization will become clear in Section 3.1.4, where we show the generality of our newly developed conditions.

Theorem 3.3 (Exponential Barrier Functions). *Consider a differentiable function $h : D \subset \mathbb{R}^n \to \mathbb{R}$, and parameters $\alpha, \beta > 0$. Assume there exists a super-level set $D_0 := h_{\geq -c}$ for some $c \geq 0$ such that:*

$$L_f h(x) \geq -(\beta[h(x)]_- + \alpha[h(x)]_+), \quad \forall x \in D_0, \tag{3.5}$$

then:

   (i) we call h an Exponential Barrier Function (EBF), the super-level set $h_{\geq 0}$ is positively invariant,

   (ii) whenever $x \in h_{<0} \cap D_0$, then as long as $h(\phi(t, x)) < 0$, $h(\phi(t, x))$ must monotonically increase to zero, at least, with a positive linear rate of $-\beta h(\phi(t, x))$, which

implies,

$$h(\phi(t, x)) \geq e^{-\beta t} h(x), \ \forall t \geq 0, \ x \in h_{<0} \cap D_0, \tag{3.6}$$

and

(iii) *whenever* $x \in h_{>0}$, *then* $h(\phi(t, x))$ *may decrease to zero, at most, with a negative linear rate of* $-\alpha h(\phi(t, x))$, *which implies,*

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x), \quad \forall t \geq 0, \ x \in h_{\geq 0}. \tag{3.7}$$

Proof. See Appendix 3.4 □

We finalize by reminding that all the results of this section require the differentiability of h. In the next section, we relax this requirement.

3.1.2 Integral Barrier Functions

This section examines integral-based versions of the classic barrier function formulations previously discussed. We will explore how these integral forms relate to earlier formulations and their implications for the invariance of a set. Specifically, in Theorems 3.4, 3.5, and 3.6, we present the integral-based versions of the barrier function formulations of Theorems 3.1, 3.2, and 3.3, respectively. Additionally, we discuss in Remarks 5, 6, and 7 the equivalence between the integral formulations and classical ones when the barrier function is differentiable.

Theorem 3.4 (Integral Nagumo's Barrier Functions)**.** *Consider a continuous function* $h : D \rightarrow \mathbb{R}$. *The super-level set* $h_{\geq 0}$ *is invariant if and only if* h *satisfies:*

$$h(\phi(t, x)) \geq 0, \ \forall \, x \in h_{=0}, \ t \geq 0, \tag{3.8}$$

in which case we call h *an Integral Nagumo's Barrier Function (INBF).*

Proof. See Appendix 3.4 □

50

Remark 5. *Theorem 3.4 only requires h to be continuous. If h is also differentiable, then (3.8) is equivalent to (3.2). Specifically, starting from an arbitrary point $x \in h_{=0}$, condition (3.8) implies that $h(\phi(t, x)) \geq 0$ for all $t \geq 0$. Therefore, $h(\phi(t, x)) - h(x) \geq -h(x) = 0, \ \forall x \in h_{=0}, t \geq 0$. Then, by evaluating the Lie derivative:*

$$L_f h(x) = \lim_{t \to 0} \frac{h(\phi(t, x)) - h(x)}{t} \geq 0,$$

condition (3.2) follows.

For the other direction, note that condition (3.2) implies $h_{\geq 0}$ is invariant, which is equivalent to condition (3.8).

In the following theorem, we relax the differentiability conditions on zeroing barrier functions recalled in Theorem 3.2 and define the integral zeroing barrier functions.

Theorem 3.5 (Integral Zeroing Barrier Functions)**.** *Consider a continuous function $h : D \subset \mathbb{R}^n \to \mathbb{R}$, and an extended class $\mathcal{K}$ function ζ. Assume there exists a super-level set $D_0 := h_{\geq -c}$ for some $c \geq 0$ such that:*

$$h(\phi(t, x)) + \int_0^t \zeta(h(\phi(s, x)))ds \geq h(x), \tag{3.9}$$

for all $t \geq 0$ and $x \in D_0$, then:

   (i) *we call h an Integral Zeroing Barrier Function (IZBF), the super-level set $h_{\geq 0}$ is invariant, and*

  (ii) *the conditions (ii)-(iii) stated in Theorem 3.2 are satisfied.*

Proof. See Appendix 3.4 □

Remark 6. *Theorem 3.5 only requires h to be continuous. If h is also differentiable, (3.9) is equivalent to (3.3), and thus the asymptotic convergence results (ii)-(iii) stated in Theorem 3.2 are satisfied.*

Precisely, condition (3.3) follows by evaluating the following Lie derivative everywhere

under condition (3.9):

$$
\begin{aligned}
L_f h(x) &= \lim_{t \to 0} \frac{h(\phi(t,x)) - h(x)}{t} \\
&\geq \lim_{t \to 0} \frac{-\int_0^t \zeta(h(\phi(s,x)))ds}{t} \\
&= \lim_{t \to 0} \frac{0 - \zeta(h(\phi(0,x)))t - o(t)}{t} \\
&= -\zeta(h(x)).
\end{aligned}
$$

For the other direction, condition (3.9) follows directly from (3.3) by applying the 'funda-

mental theorem of calculus II'[68].

In the last theorem of this section, we relax the differentiability conditions on
exponential barrier functions introduced in Theorem 3.3 and define the integral
exponential barrier functions.

Theorem 3.6 (Integral Exponential Barrier Functions). *Consider a continuous function*

$h : D \subset \mathbb{R}^n \to \mathbb{R}$ *and parameters* $\alpha, \beta > 0$. *Assume there exists a super-level set*

$D_0 := h_{\geq -c}$ *for some* $c \geq 0$ *such that:*

$$
h(\phi(t,x)) \geq e^{-\beta t}[h(x)]_- + e^{-\alpha t}[h(x)]_+, \tag{3.10}
$$

for all $x \in D_0$ *and* $t \geq 0$, *then:*

(i) *we call* h *an Integral Exponential Barrier Function (IEBF), the super-level set* $h_{\geq 0}$ *is*
invariant, and

(ii) *the conditions (ii)-(iii) stated in Theorem 3.3 are satisfied.*

Proof. See Appendix 3.4 □

Remark 7. *Theorem 3.6 only requires* h *to be continuous. If* h *is also differentiable,*

then (3.10) is equivalent to (3.5). In particular, condition (3.5) follows by evaluating the

following Lie derivative everywhere under condition (3.10), i.e.,

Whenever $x \in h_{\geq 0}$:

$$L_f h(x) = \lim_{t \to 0} \frac{h(\phi(t, x)) - h(x)}{t}$$

$$\geq \lim_{t \to 0} \frac{h(x)e^{-\alpha t} - h(x)}{t}$$

$$= h(x) \lim_{t \to 0} \frac{e^{-\alpha t} - 1}{t}$$

$$= -\alpha h(x).$$

Similarily, whenever $x \in D_0 \cap h_{<0}$:

$$L_f h(x) = \lim_{t \to 0} \frac{h(\phi(t, x)) - h(x)}{t} \geq -\beta h(x).$$

For the other direction, we have shown condition (3.5) implies (3.10) in Theorem 3.3 part (ii-iii).
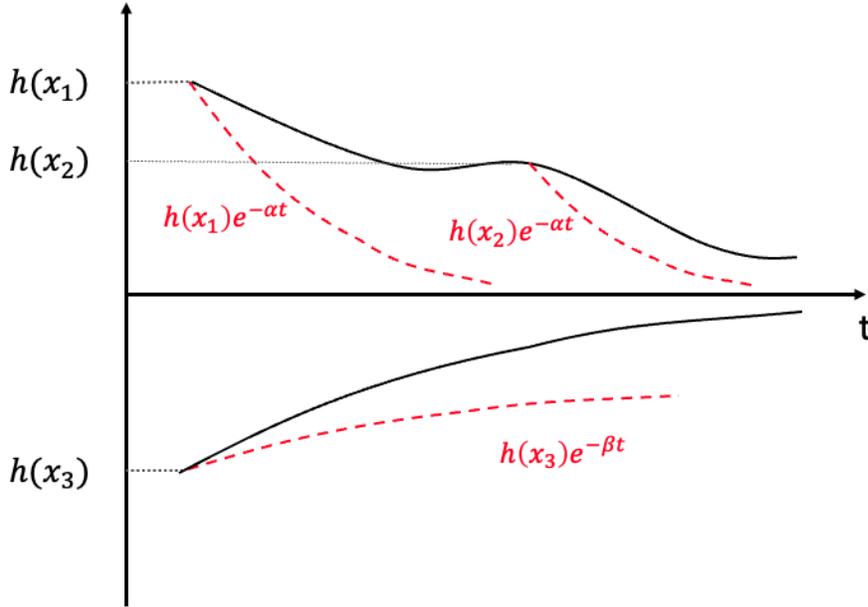
Figure 3-1. This illustration depicts the integral exponential barrier function. The red dashed lines represent $e^{-\beta t}[h(x)]_- + e^{-\alpha t}[h(x)]_+$, i.e., the right-hand side of equation 3.10, evaluated at three different states x_1, x_2, and x_3 over time $t \geq 0$. The two black lines, representing $h(\phi(t, x_1))$ and $h(\phi(t, x_3))$, demonstrate that these function values do not fall below their respective red dashed thresholds.

3.1.3 Recurrent Barrier Functions

We are now ready to provide a relaxation to the barrier conditions that lead to the invariance of the super-level set $h_{\geq 0}$. To relax the notion of invariance, one must allow trajectories to temporarily leave a set. Therefore, we recall the definitions of recurrence and τ-recurrence defined in Definition 2.4 and 2.5, respectively.
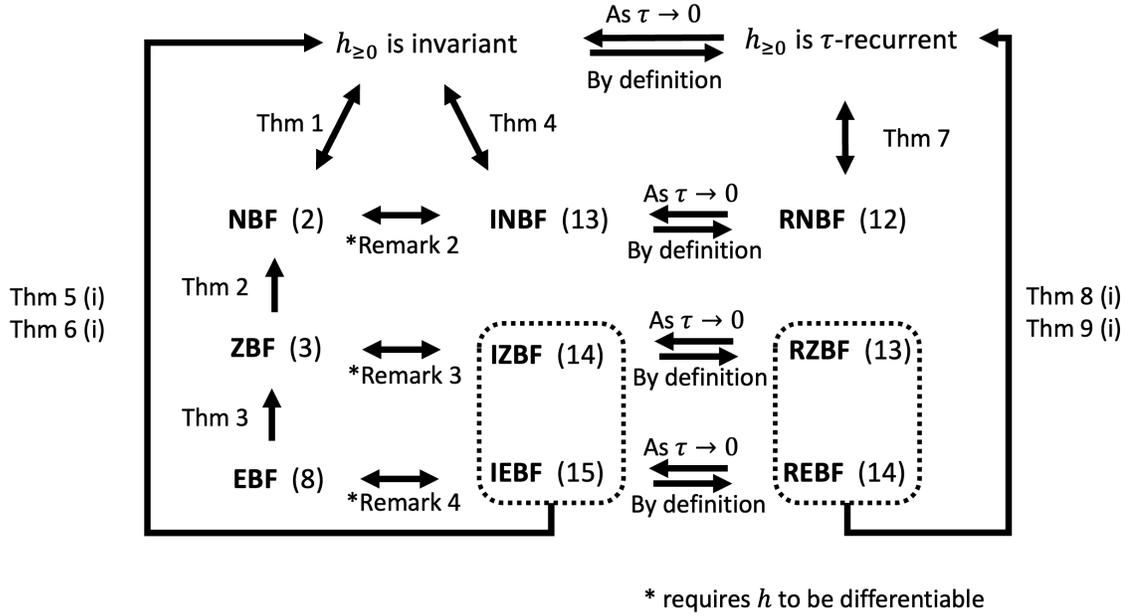
Figure 3-2. Relationships among classic barrier functions (Theorems 3.1-3.3), integral barrier functions (Theorem 3.4-3.6), and recurrent barrier functions (Theorem 3.7-3.9).

In the rest of this section, we generalize the aforementioned barrier function formulations into recurrent-based versions that certify the τ-recurrence of the super-level set $h_{\geq 0}$. Specifically, in theorems 3.7, 3.8, and 3.9, we present the recurrence-based versions of the integral barrier function formulations of theorems 3.4, 3.5, and 3.6, respectively. The relationships between all barrier functions are summarized in Fig 3-2. As usual, we start with Nagumo's version.

Theorem 3.7 (Recurrent Nagumo's Barrier Functions)**.** *Consider a continuous function*

$h : D \to \mathbb{R}$, then the super-level $h_{\geq 0}$ is τ-recurrent if and only if h satisfies:

$$\max_{t \in (0,\tau]} h(\phi(t,x)) \geq 0, \quad \forall x \in h_{=0}, \tag{3.11}$$

in which case we call it a Recurrent Nagumo's Barrier Function (RNBF).

Proof. See Appendix 3.4 □

As in the differential and integral formulations, Nagumo's version of recurrence does not restrict the rate at which the barrier function changes along the trajectories. For that, we will need to introduce the Zeroing formulation.

Definition 3.3 (Recurrent Zeroing Barrier Functions).

A continuous function $h : D \subset \mathbb{R}^n \to \mathbb{R}$ is a Recurrent Zeroing Barrier Function (RZBF) if there exists an extended class $\mathcal{K}$ function ζ and a super-level set $D_0 := h_{\geq -c}$, $c \geq 0$, such that:

$$\max_{t \in (0,\tau]} h(\phi(t,x)) + \int_0^t \zeta(h(\phi(s,x)))ds \geq h(x), \forall x \in D_0. \tag{3.12}$$

In contrast with the differential and integral formulations of zeroing barrier functions, the recurrent one does not constrain the rates at which the barrier function changes at all time instants of a trajectory, but only on countably infinite many times that are most separated by τ in consecutive steps.

Lemma 3.2. *Consider a Recurrent Zeroing Barrier Function h defined over $D_0 := h_{\geq -c}$ for some $c \geq 0$. Then, for any $x \in D_0$, there exists a sequence of times $\{t_n\}_{n \in \mathbb{N}}$, with $t_0 = 0$,*

$$t_{n+1} = \max\{\arg\max_{t \in (t_n, t_n+\tau]} h(\phi(t,x)) + \int_{t_n}^t \zeta(h(\phi(s,x)))ds\},$$

$$\lim_{n \to \infty} t_n = \infty, \text{ and } t_{n+1} - t_n \in (0,\tau], \quad \forall n \in \mathbb{N}, \tag{3.13}$$

such that for each $x_n := \phi(t_n, x)$, $x_n \in D_0$, and

$$h(x_{n+1}) \geq \max\{h(x_n) + \delta_n, [h(x_n)]_-\}, \forall n \in \mathbb{N}, \tag{3.14}$$

with $\delta_n := -\int_0^{t_{n+1}-t_n} \zeta(h(\phi(s, x_n)))ds$, satisfying $\delta_n > 0$ whenever $h(x_n)$ and $h(x_{n+1}) <$ 0.

Proof. See Appendix (Section 3.4). □

The following theorem gives a detailed explanation of the implications of the Recurrent Zeroing Barrier Function.

Theorem 3.8. *Consider a Recurrent Zeroing Barrier Function h defined over $D_0 := h_{\geq -c}$ for some $c \geq 0$ and let $x_n := \phi(t_n, x)$ be the states along the sequence of times $\{t_n\}_{n \in \mathbb{N}}$ specified in Lemma 3.2. Then:*

(i) the super-level set $h_{\geq 0}$ is τ-recurrent,

(ii) whenever $x \in h_{<0} \cap D_0$, then as long as $h(x_n) < 0$, $h(x_{n+1})$ must monotonically increase, at least, by a step size of $\delta_n > 0$, if $h(x_{n+1}) < 0$, or $-h(x_n) > 0$, if $h(x_{n+1}) \geq 0$, and

(iii) whenever $h(x_n) > 0$, then $h(x_{n+1})$ may decrease to zero, at most, by a negative step size of $\max\{\delta_n, -h(x_n)\}$.

Proof. See Appendix (Section 3.4). □

As before, we end with the exponential formulation which constrains the rates using exponential functions of time. As in the zeroing version, it only constrains the rates at countably infinite time instants.

Definition 3.4 (Recurrent Exponential Barrier Functions).

A continuous function $h : D \subset \mathbb{R}^n \to \mathbb{R}$ is a Recurrent Exponential Barrier Function (REBF) if there exists parameters $\alpha, \beta > 0$ and a super-level set $D_0 := h_{\geq -c}$, for some $c \geq 0$, such that:

$$\max_{t \in (0, \tau]} e^{\beta t}[h(\phi(t, x))]_- + e^{\alpha t}[h(\phi(t, x))]_+ \geq h(x) \qquad (3.15)$$

for all $x \in D_0$.
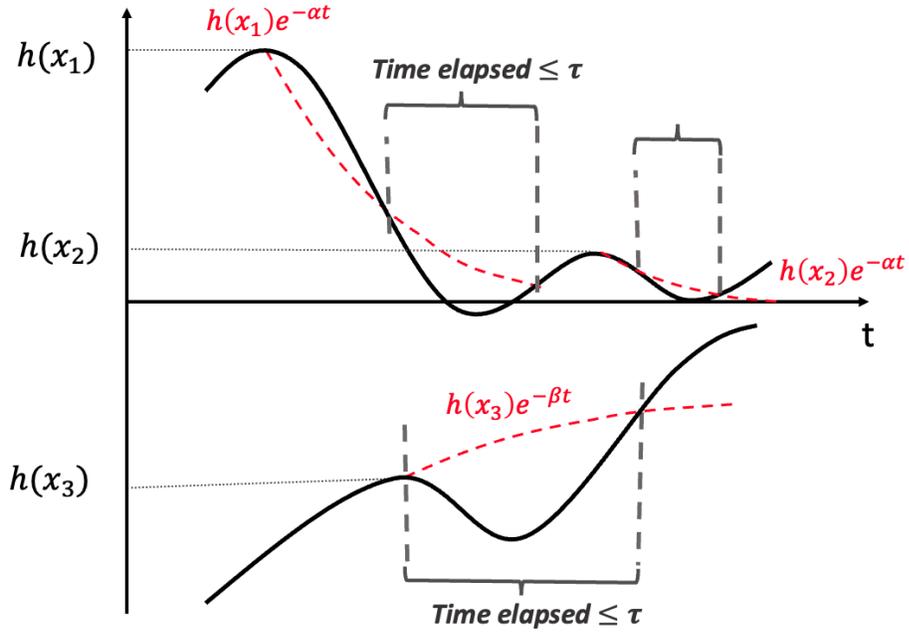
Figure 3-3. This illustration demonstrates the recurrent exponential barrier function. The red dashed lines represent $e^{-\beta t}[h(x)]_- + e^{-\alpha t}[h(x)]_+$, i.e., the rearranged right-hand side of equation 3.15, evaluated at three different states x_1, x_2, and x_3 over time $t \geq 0$. The two black lines, representing $h(\phi(t, x_1))$ and $h(\phi(t, x_3))$, demonstrate that these function values are permitted to temporarily drop below their respective red dashed thresholds but only for periods shorter than τ.

Lemma 3.3. *Consider a Recurrent Exponential Barrier Function h defined over $D_0 := h_{\geq -c}$ for some $c \geq 0$ with parameters $\alpha, \beta > 0$. Then, for any $x \in D_0$, there exists a sequence of times $\{t_n\}_{n \in \mathbb{N}}$, with $t_0 = 0$,*

$$t_{n+1} = \max\{\arg\max_{t \in (t_n, t_n + \tau]} e^{\beta t}[h(\phi(t, x))]_- + e^{\alpha t}[h(\phi(t, x))]_+\}$$

$$\lim_{n \to \infty} t_n = \infty \ and \ t_{n+1} - t_n \in (0, \tau], \ \ \forall n \in \mathbb{N}, \tag{3.16}$$

such that for each state $x_n := \phi(t_n, x)$, we have, $x_n \in D_0$, and

$$h(x_{n+1}) \geq e^{-\beta \Delta t_n}[h(x_n)]_- + e^{-\alpha \Delta t_n}[h(x_n)]_+ , \tag{3.17}$$

for all $n \in \mathbb{N}$, with $\Delta t_n := t_{n+1} - t_n$.

Proof. See Appendix (Section 3.4). $\qquad\square$

We then summarize the implications of the Recurrent Exponential Barrier Function in the next theorem.

Theorem 3.9. *Consider a Recurrent Exponential Barrier Function h defined over $D_0 :=$ $h_{\geq -c}$ for some $c \geq 0$ with parameters $\alpha, \beta > 0$ and let $x_n := \phi(t_n, x)$ be the states along the sequence of times $\{t_n\}_{n\in\mathbb{N}}$ specified in Lemma 3.3. Then:*

(i) *the super-level set $h_{\geq 0}$ is τ-recurrent,*

(ii) *whenever $x \in h_{<0} \cap D_0$, then as long as $h(x_n) < 0$, $h(x_{n+1})$ must monotonically increase, at least, by a positive step size of $\delta_n = (e^{-\beta \Delta t_n} - 1)h(x_n)$, which implies,*

$$h(x_n) \geq h(x)e^{-\beta t_n}, \ \ \forall n \in \mathbb{N}, \ x \in h_{<0} \cap D_0, \tag{3.18}$$

and

(iii) *whenever $h(x_n) > 0$, then $h(x_{n+1})$ may decrease to zero, at most, by a negative step size of $\delta_n = (e^{-\beta \Delta t_n} - 1)h(x_n)$, which implies,*

$$h(x_n) \geq h(x)e^{-\alpha t_n}, \ \ \forall n \in \mathbb{N}, \ x \in h_{\geq 0}. \tag{3.19}$$

Proof. See Appendix (Section 3.4). $\qquad\square$

3.1.4 The generality of recurrence conditions

In the previous section, we introduced a set of novel barrier conditions that relaxed the invariant requirement on the zero super level set of h. We will now show that this relaxation widely decouples the geometry of the vector field with the geometry of the level sets of h. This allows us to characterize a vast family of sets and functions that can be used to certify safety. Our prior work in Chapter 2 inspires our results, wherein we show that under mild conditions, every set contained within the region of attraction of an equilibrium point is recurrent see, e.g., Corollary 2.2.

In this section, we generalize this idea in the context of certifying safety. We start by introducing mild regularity constraints on h, which we will need later in Theorem 3.10. This requires us to introduce the notion of sector-bounded functions.

Definition 3.5 (Sector Containment). *Let $h : D \subset \mathbb{R}^n \to \mathbb{R}$ be continuous. If $\exists \alpha_1, \alpha_2 > 0$ such that*

$$(h(x) - \alpha_1 \mathrm{sd}(x, h_{\leq 0}))(h(x) - \alpha_2 \mathrm{sd}(x, h_{\leq 0})) \leq 0, \tag{3.20}$$

*for all $x \in D$, we say that h is **sector-contained**.*

Remark 8. *Given parameters $\alpha_2 \geq \alpha_1 > 0$, (3.20) is satisfied if and only if for all $x \in D$:*

$$\alpha_2 \mathrm{sd}(x, h_{\leq 0}) \geq h(x) \geq \alpha_1 \mathrm{sd}(x, h_{\leq 0}) \geq 0 \;\; \textbf{if } h(x) \geq 0, \tag{3.21}$$
$$0 \geq \alpha_1 \mathrm{sd}(x, h_{\leq 0}) \geq h(x) \geq \alpha_2 \mathrm{sd}(x, h_{\leq 0}) \;\; \textbf{if } h(x) < 0. \tag{3.22}$$

In the following theorem, we show that the existence of a sector-contained IEBF h is sufficient to make the signed distance to the zero sub-level set of h a REBF. As such, this theorem illustrates the generality of our recurrent condition.

Theorem 3.10. *Let h be an Integral Exponential Barrier Function with exponential rates α and β, defined over $D_0 := h_{\geq -c}$ for some $c \geq 0$. Then, if h is sector-contained with parameters α_1 and α_2, the function $\hat{h}(\cdot) = \mathrm{sd}(\cdot, h_{\leq 0})$ is a Recurrent Exponential Barrier Function, i.e., the following conditions hold:*

$$\max_{t \in (0, \hat{\tau}]} e^{\hat{\beta} t} [\hat{h}(\phi(t, x))]_- + e^{\hat{\alpha} t} [\hat{h}(\phi(t, x)]_+ \geq \hat{h}(x), \tag{3.23}$$

for all $x \in D_0$ and any $\hat{\alpha}, \hat{\beta}, \hat{\tau} > 0$ satisfying $\hat{\alpha} > \alpha$, $\hat{\beta} < \beta$, and

$$\hat{\tau} \geq \max\{\frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha} - \alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta - \hat{\beta}}\}. \tag{3.24}$$

Proof. See Appendix (Section 3.4). □

While the REBF in Theorem 3.10 has a simple definition, it still requires the knowledge of the sub-level set of the IEBF that is assumed to exist. In the following

theorem, we relax this assumption and only require the knowledge of any set that contains the super-level set of the IEBF and is contained in the domain that satisfies the IEBF conditions.

Theorem 3.11. *Let h be an Integral Exponential Barrier Function with exponential rates α and β defined over $D_0 := h_{\geq -c}$ for some $c \geq 0$. If h is sector-contained with parameters α_1 and α_2, then, for any closed set S satisfying $h_{\geq 0} \subset S \subseteq D_0 = h_{\geq -c}$ and $\partial S \cap h_{=0} = \emptyset$, the function*

$$\hat{h}(x) := -\mathrm{sd}(x, S)$$

is a Recurrent Exponential Barrier, i.e., the following conditions hold:

$$\max_{t \in (0, \hat{\tau}]} \{ e^{\hat{\beta} t}[\hat{h}(\phi(t, x))]_- + e^{\hat{\alpha} t}[\hat{h}(\phi(t, x))]_+ \} \geq \hat{h}(x) \tag{3.25}$$

for all $x \in \hat{h}_{\geq -\hat{c}}$ with $\hat{c} \geq 0$ such that $\hat{h}_{\geq -\hat{c}} \subseteq D_0$, any $\hat{\alpha}, \hat{\beta}, \hat{\tau} > 0$ satisfying $\hat{\alpha} > \alpha$, $\hat{\beta} < \beta$, $\hat{\beta} \leq \hat{\alpha}$, and

$$\hat{\tau} \geq \max\{ \frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha} - \alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta - \hat{\beta}} \} + \frac{\log(\bar{\delta}/\underline{\delta})}{\min\{\hat{\alpha}, \hat{\beta}\}},$$

with

$$\bar{\delta} := \sup_{x \in D_0} (\mathrm{sd}(x, h_{\geq 0}) - \mathrm{sd}(x, S)), \tag{3.26a}$$

$$\underline{\delta} := \inf_{x \in D_0} (\mathrm{sd}(x, h_{\geq 0}) - \mathrm{sd}(x, S)). \tag{3.26b}$$

Proof. See Appendix (Section 3.4). □

3.1.5 Safety assessment

A τ-recurrent set S outside of the known unsafe region does not immediately imply safety, as is the case with an invariant one. To practically employ the notion of recurrence for ensuring safety, the following result is pivotal: it demonstrates that a τ-recurrent set, along with the states visited by the finite-time trajectories starting from it, i.e., the finite-time reachable set, constitute an invariant set. Consequently,

this process certifies an invariant set in a manner akin to classical methodologies, which can be used to verify safety.

Theorem 3.12. *Consider a closed set S that is τ-recurrent. Then the finite time reachable set*

$$\mathcal{R}_\tau(S) := \bigcup_{x \in S, t \in [0,\tau]} \phi(t, x) \tag{3.27}$$

is invariant.

Proof. See Appendix (Section 3.4). □

We have identified sufficient conditions that guarantee a super-level set of a recurrent-based barrier function is τ-recurrent, which in turn implies its bounded-time reachable set with bound τ is invariant.

Now, one needs to additionally make sure its τ-seconds reachable set, defined in (3.27), does not intersect with the known unsafe. However, characterizing such a finite-time reachable set is non-trivial, as it may require adaptations to accommodate the distinct trajectories of usually undecidable nonlinear systems.

Yet, under mild conditions, trajectories originating from a τ-recurrent set S are restricted from straying too far from S, as they can only leave S for at most τ-seconds. Consequently, it is possible to over-approximate the τ-seconds reachable set and effectively certify safety if the resulting over-approximation does not intersect the unsafe set.

Therefore, we provide the following lemma that bounds how far the τ-recurrent trajectories can stray from the recurrent set in τ seconds. The lemma is an extension of Lemma 1 in [2].

Lemma 3.4 (Containment Lemma). *Let $S \subseteq D$ be a closed set that is τ-recurrent and define:*

$$c_1 = F(\mathcal{R}_\tau(S))\tau, \quad c_2 = F(\partial S)\tau e^{L(\mathcal{R}_\tau(S))\tau},$$
$$c = \min\{c_1, c_2\}. \tag{3.28}$$

Then, starting from any $x \in S$, the trajectory satisfies:

$$\mathrm{sd}(\phi(t, x), S) \le c, \quad \forall t \ge 0. \tag{3.29}$$

Proof. See Appendix (Section 3.4). □

Note that Lemma 3.4 provides necessary regularity conditions for trajectories starting of a τ-recurrent set $h_{\ge 0}$. Building on this foundation, we present the following theorem, which practically leverages the concept of τ-recurrence to characterize a safe state-space region of the system.

Theorem 3.13. *Consider a continuous function $h : D \to \mathbb{R}$ and a set $\mathcal{X}_u$ of unsafe states. If the super-level set $h_{\ge 0}$ is τ-recurrent and $\{h_{\ge 0} + \mathcal{B}_c\} \cap \mathcal{X}_u = \emptyset$, where the '+' stands for the Minkovski sum, $\mathcal{B}_c$ is a closed ball of radius r around the origin, and the constant c is defined in (3.28), then $h_{\ge 0}$ is a safe state space region.*

Proof. Given the closed τ-recurrent set $h_{\ge 0}$, note first that Theorem 3.12 implies its τ-seconds reachable set $\mathcal{R}_\tau(h_{\ge 0})$ is invariant. Then, Lemmma 3.4 further ensures the invariant set $\mathcal{R}_\tau(h_{\ge 0}) \subseteq \{h_{\ge 0} + \mathcal{B}_c\}$ since trajecties starting from $h_{\ge 0}$ cannot depart from it more then c. Finally, we have

$$\phi(t, x) \in \mathcal{R}_\tau(h_{\ge 0}) \subseteq \{h_{\ge 0} + \mathcal{B}_c\},$$

for all $x \in h_{\ge 0}$ and $t \ge 0$. This, together with the fact that $\{h_{\ge 0} + \mathcal{B}_c\} \cap \mathcal{X}_u = \emptyset$, further implies $\phi(t, x) \notin \mathcal{X}_u$. Therefore, result follows. □

3.2 Learning a local safe region

In pursuit of providing practical data-driven methods to delineate the safe state space region, we have developed recurrent conditions that are computationally more tractable to characterize. However, like the differential and integral conditions, these recurrent conditions must be satisfied within a domain $D_0 := h_{\ge -c}$ for some

$c \geq 0$. In practice, such a super-level set of h is often unbounded, rendering it impractical to sample sufficiently.

Therefore, in the remainder of this chapter, we adopt the following practical safety objective, which is commonly employed in the literature, see, e.g., [67].

Definition 3.6 (Local Safety). *Let $\mathcal{X} \subseteq D$ be a region of interest and $\mathcal{X}_u \subseteq \mathcal{X}$ be a set of unsafe states within the region of interest. Then, a trajectory $\phi(t, x)$ of (2.1) is locally unsafe if there exists a time $T \geq 0$ such that $\phi(T, x) \cap \mathcal{X}_u \neq \emptyset$ and $\phi(t, x) \in \mathcal{X}$ for all $t \in [0, T]$.*

We say that a set $\mathcal{X}_s \subseteq \mathcal{S}$ is a local safe state space region if there are no local unsafe trajectories starting from $\mathcal{X}_s$.

Note that this local safety requirement prevents trajectories starting in $\mathcal{X}_s$ from entering $\mathcal{X}_u$ without leaving $\mathcal{X}$ first. According to Definition 3.1, if $\mathcal{X}$ is an invariant set, then local safety is equivalent to safety. However, if $\mathcal{X}$ is not invariant, the safety criteria are not enforced once trajectories depart from $\mathcal{X}$.

From Section 3.2.1 through Section 3.2.4, we adapt the foundational results introduced in Section 3.1 into this practical safety objective. Then, in Section 3.2.5, we demonstrate the method for certifying the recurrent barrier function and other safety requirements for a local neighborhood around a point x, utilizing only one finite-length trajectory that originates from x.

3.2.1 The local differential and integral barrier conditions

We note that by altering the domain within which the differential and integral barrier conditions must be satisfied, the corresponding localized barrier functions can similarly certify these local safety requirements. Furthermore, starting from a point inside $\mathcal{X}$, the bounds on the convergence rates of the function value $h(\phi(t, x))$ remain valid until trajectories exit $\mathcal{X}$. In the following two theorems, we introduce the localized versions of the Exponential Barrier Function and the Integral Exponential

Barrier Function as examples.

Theorem 3.14 (Local EBF)**.** *Consider a differentiable function $h : D \subset \mathbb{R}^n \to \mathbb{R}$, and parameters $\alpha, \beta > 0$. Assume that for all $x \in \mathcal{X}$:*

$$L_f h(x) \geq -(\beta[h(x)]_- + \alpha[h(x)]_+).$$

Then,

   (i) we call h a Local Exponential Barrier Function (LEBF),

  (ii) the set $h_{\geq 0} \cap \mathcal{X}$ is locally safe if $h(x) < 0, \forall x \in \mathcal{X}_u$,

 (iii) whenever $x \in h_{<0} \cap \mathcal{X}$, then as long as there exists a $T > 0$ such that $h(\phi(T, x)) < 0$ and $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T]$, $h(\phi(t, x))$ must monotonically increase for all $t \in [0, T]$, at least, with a positive linear rate of $-\beta h(\phi(t, x))$, which implies,

$$h(\phi(t, x)) \geq e^{-\beta t} h(x),$$

for all $x \in h_{<0} \cap \mathcal{X}, t \in [0, T]$, and all $T \geq 0$ such that $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T]$, and

 (iv) whenever $x \in h_{>0} \cap \mathcal{X}$ and there exists a $T > 0$ such that $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T]$, then $h(\phi(t, x))$ may decrease to zero, at most, with a negative linear rate of $-\alpha h(\phi(t, x))$, which implies,

$$h(\phi(t, x)) \geq e^{-\alpha t} h(x),$$

for all $x \in h_{>0} \cap \mathcal{X}, t \in [0, T]$, and all $T \geq 0$ such that $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T]$.

Theorem 3.15 (Local IEBF)**.** *Consider a continuous function $h : D \subset \mathbb{R}^n \to \mathbb{R}$ and parameters $\alpha, \beta > 0$. If:*

$$h(\phi(t, x)) \geq e^{-\beta t}[h(x)]_- + e^{-\alpha t}[h(x)]_+, \tag{3.30}$$

for all $x \in \mathcal{X}, t \in [0, T]$, and all $T \geq 0$ such that $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T]$, then:

(i) we call h a Local Integral Exponential Barrier Function (LIEBF),

(ii) the set $h_{\geq 0} \cap \mathcal{X}$ is locally safe if $h(x) < 0, \forall x \in \mathcal{X}_u$, and

(iii) the conditions (iii)-(iv) in Theorem 3.14 are satisfied.

3.2.2 The local recurrent barrier condition

To further generalize the localized version of the IEBF into the recurrent form, we again focus only on the trajectory prior to its exit from $\mathcal{X}$. This requires us to first introduce this localized version of τ-recurrent sets.

Definition 1 (Locally τ-Recurrent Sets).
A set $S \subseteq \mathcal{X}$ is locally τ-recurrent w.r.t. $\mathcal{X}$, if for any $x \in S$, and any $T \geq 0$ such that $\phi(t, x) \in \mathcal{X}, \forall t \in [0, T + \tau]$,

$$\exists\, t' > T, \;\; with \;\; t' - T \in (0, \tau] \;\; s.t. \;\; \phi(t', x) \in S.$$

Recall from Definition 2.5 that a τ-recurrent set ensures solutions starting in this set will visit it back within τ seconds, infinitely often. However, in this localized version of τ-recurrent, the requirement is only applicable before trajectories exit $\mathcal{X}$. Leveraging the same idea, we now introduce the localized version of the REBF defined in Definition 3.4.

Definition 2 (Local REBF). *A continuous function $h : D \subset \mathbb{R}^n \to \mathbb{R}$ is a Local Recurrent Exponential Barrier Function (LREBF) if there exist parameters $\alpha, \beta > 0$, such that*

$$\max_{t \in (0, \tau]} e^{\beta t}[h(\phi(t, x))]_- + e^{\alpha t}[h(\phi(t, x))]_+ \geq h(x), \tag{3.31}$$

for all $x \in \mathcal{X}_\tau := \{x \in \mathcal{X} | \phi(t, x) \in \mathcal{X}, \forall t \in [0, \tau]\}$.

Lemma 3.5. *Consider a Local Recurrent Exponential Barrier Function h defined over $\mathcal{X}_\tau$ with parameters $\alpha, \beta > 0$. Then, for each $x \in \mathcal{X}_\tau$ and for all $T \geq 0$ such that $\phi(t, x) \in \mathcal{X}$,*

$\forall t \in [0, T + \tau]$, *there exists a sequence of times* $\{t_n\}_{n \in \{0,...,N\}}$, *with* $t_0 = 0$,

$$t_{n+1} = \max\{\arg\max_{t \in (t_n, t_n + \tau]} e^{\beta t}[h(\phi(t, x))]_- + e^{\alpha t}[h(\phi(t, x))]_+\},$$

$$t_N > T, \text{ and } t_{n+1} - t_n \in (0, \tau], \ \forall n \in \{0, \dots, N-1\}, \tag{3.32}$$

such that for each state $x_n := \phi(t_n, x)$, *we have*

$$h(x_{n+1}) \geq e^{-\beta \Delta t_n}[h(x_n)]_- + e^{-\alpha \Delta t_n}[h(x_n)]_+, \tag{3.33}$$

for all $n \in \{0, \dots, N-1\}$, *with* $\Delta t_n := t_{n+1} - t_n$.

Proof. See Appendix (Section 3.4). □

Theorem 3.16. *Consider a Local Recurrent Exponential Barrier Function* h *defined over* $\mathcal{X}_\tau$ *with parameters* $\alpha, \beta > 0$ *and let* $x_n := \phi(t_n, x)$ *be the states along the sequence of times* $\{t_n\}_{n \in \{0,...,N\}}$ *specified in Lemma* 3.5. *Then,*

(i) *the set* $h_{\geq 0} \cap \mathcal{X}_\tau$ *is locally* τ-*recurrent w.r.t.* $\mathcal{X}$,

(ii) *whenever* $x \in h_{<0} \cap \mathcal{X}_\tau$, *then if* $h(x_n) < 0$ *and* $\phi(t, x) \in \mathcal{X}, \forall t \in [0, t_n + \tau]$, $h(x_{n+1})$ *must be larger than* $h(x_n)$, *at least, by a positive step size of* $\delta_n = (e^{-\beta \Delta t_n} - 1)h(x_n)$ *, which implies that*

$$h(x_n) \geq h(x)e^{-\beta t_n}, \tag{3.34}$$

for all $x \in h_{<0} \cap \mathcal{X}_\tau$, $n \in \{0, \dots, N\}$,

and

(iii) *whenever* $x \in h_{>0} \cap \mathcal{X}_\tau$ *and* $\phi(t, x) \in \mathcal{X}, \forall t \in [0, t_n + \tau]$, *then* $h(x_{n+1})$ *must be larger than zero but may smaller than* $h(x_n)$, *at most, by a negative step size of* $\delta_n = (e^{-\beta \Delta t_n} - 1)h(x_n)$, *which implies that*

$$h(x_n) \geq h(x)e^{-\alpha t_n}, \tag{3.35}$$

for all $x \in h_{\geq 0} \cap \mathcal{X}_\tau$, $n \in \{0, \dots, N\}$.

Proof. See Appendix (Section 3.4). □

3.2.3 The generality of local the recurrent condition

In this subsection, we will again demonstrate that a simple signed distance function can satisfy our local recurrence exponential barrier conditions. Similar to the arguments made in Section 3.1.4, we will assume there exists a Local Integral Exponential Barrier Function h satisfying the following regularity constraints.

Given a set S satisfies $S \cap \mathcal{X} \neq \emptyset$ and an arbitrary norm (denoted by $\| \cdot \|$), we define the constrained signed distance:

$$\mathrm{sd}|_{\mathcal{X}}(x, S) := \begin{cases} \inf_{y \in \partial S \cap \mathrm{int}\mathcal{X}} \|y - x\| & \text{if } x \notin S \\ -\inf_{y \in \partial S \cap \mathrm{int}\mathcal{X}} \|y - x\| & \text{if } x \in S. \end{cases}$$

Definition 3.7. *Let $h : D \subset \mathbb{R}^n \to \mathbb{R}$ be continuous. If $\exists \alpha_1, \alpha_2 > 0$ such that*

$$(h(x) - \alpha_1 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}))(h(x) - \alpha_2 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0})) \leq 0, \tag{3.36}$$

*for all $x \in \mathcal{X}$, we say that h is **locally sector-contained**.*

Given parameters $\alpha_2 \geq \alpha_1 > 0$, (3.36) is satisfied if and only if for all $x \in D$:

$$\alpha_2 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \geq h(x) \geq \alpha_1 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \geq 0 \text{ if } h(x) \geq 0, \tag{3.37}$$
$$0 \geq \alpha_1 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \geq h(x) \geq \alpha_2 \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \text{ if } h(x) < 0. \tag{3.38}$$

We then use the following theorem to demonstrate the generality of our recurrent condition. Precisely, if there exists a locally sector-contained LIEBF h, then the constrained signed distance to the zero sub-level set of h is an LREBF.

Theorem 3.17. *Let h be a Local Integral Exponential Barrier Function defined over $\mathcal{X}$ with exponential rates α and β. Then, if h is locally sector-contained with parameters α_1 and α_2, the function $\hat{h}(\cdot) = \mathrm{sd}|_{\mathcal{X}}(\cdot, h_{\leq 0})$ is a Local Recurrent Exponential Barrier Function, i.e., the following conditions hold:*

$$\max_{t \in (0, \hat{\tau}]} e^{\hat{\beta} t}[\hat{h}(\phi(t, x))]_- + e^{\hat{\alpha} t}[\hat{h}(\phi(t, x))]_+ \geq \hat{h}(x), \tag{3.39}$$

for all $x \in \mathcal{X}_{\hat{\tau}}$ with any $\hat{\alpha}, \hat{\beta}, \hat{\tau} > 0$ satisfying $\hat{\alpha} > \alpha$, $\hat{\beta} < \beta$, and

$$\hat{\tau} \geq \max\{\frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha} - \alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta - \hat{\beta}}\}. \tag{3.40}$$

Proof. See Appendix (Section 3.4). □

Again, the LREBF described in Theorem 3.17 necessitates the knowledge of the sub-level set $h_{\leq 0}$ of the presumed LIEBF. Conversely, the following theorem lessens this dependence. It specifies that knowledge of any set that contains the intersection $h_{\geq 0} \cap \mathcal{X}$ is sufficient to establish a constrained signed distance that fulfills the LREBF conditions.

Theorem 3.18. *Let h be a Local Integral Exponential Barrier Function defined over $\mathcal{X}$ with exponential rates α and β. If h is locally sector-contained with parameters α_1 and α_2, then, for any closed set S satisfying $h_{\geq 0} \cap \mathcal{X} \subset S$ and $\partial S \cap h_{=0} = \emptyset$, the function*

$$\hat{h}(x) := -\mathrm{sd}|_{\mathcal{X}}(x, S)$$

is a Local Recurrent Exponential Barrier function, i.e., the following conditions hold:

$$\max_{t \in (0, \hat{\tau}]} \{ e^{\hat{\beta} t} [\hat{h}(\phi(t, x))]_- + e^{\hat{\alpha} t} [\hat{h}(\phi(t, x))]_+ \} \geq \hat{h}(x) \tag{3.41}$$

for all $x \in \mathcal{X}_{\hat{\tau}}$ and any $\hat{\alpha}, \hat{\beta}, \hat{\tau} > 0$ satisfying $\hat{\alpha} > \alpha$, $\hat{\beta} < \beta$, $\hat{\beta} \leq \hat{\alpha}$, and

$$\hat{\tau} \geq \max\{ \frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha} - \alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta - \hat{\beta}} \} + \frac{\log(\bar{\delta}/\underline{\delta})}{\min\{\hat{\alpha}, \hat{\beta}\}},$$

with

$$\bar{\delta} := \sup_{x \in \mathcal{X}} (\mathrm{sd}|_{\mathcal{X}}(x, h_{\geq 0}) - \mathrm{sd}|_{\mathcal{X}}(x, S)), \tag{3.42a}$$

$$\underline{\delta} := \inf_{x \in \mathcal{X}} (\mathrm{sd}|_{\mathcal{X}}(x, h_{\geq 0}) - \mathrm{sd}|_{\mathcal{X}}(x, S)). \tag{3.42b}$$

Proof. See Appendix (Section 3.4). □

3.2.4 Local safety assessment

Similar to the arguments made Section 3.1.5, a locally τ-recurrent set S w.r.t. $\mathcal{X}$ satisfying $S \cap \mathcal{X}_u = \emptyset$ does not immediately imply local safety. To assess local safety within $\mathcal{X}$, we will utilize a backward finite time reachable set, rather than relying on the invariance of the finite time reachable set $\mathcal{R}_\tau(S)$. This approach will be detailed in the following theorem.

Theorem 3.19. *Let $S \subseteq \mathcal{X}$ be a closed locally τ-recurrent set S w.r.t. $\mathcal{X}$ and define the finite time backward reachable set of $\mathcal{X}_u$ as:*

$$\mathcal{R}_\tau^{-1}(\mathcal{X}_u) := \{x \in \mathcal{X} \mid \exists t \in [0, \tau] \, s.t. \, \phi(t, x) \in \mathcal{X}_u\} \supseteq \mathcal{X}_u.$$

Then S is locally safe whenever $S \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$.

Proof. Suppose $S \subseteq \mathcal{X}$ is not locally safe, there is then an initial state $x \in S$ and a time $T \geq 0$ such that $\phi(T, x) \in \mathcal{X}_u$ and $\phi(t, x) \in \mathcal{X}$ for all $t \in [0, T]$. Since $S \cap \mathcal{X}_u = \emptyset$, the unsafe point $\phi(T, x) \notin S$. We therefore use T' to denote the the trajectory $\phi(t, x)$ stays within the closed set S before T, i.e.,

$$x' := \phi(T', x) \in S \quad \text{and} \quad \phi(t, x) \notin S, \ \forall t \in (T', T].$$

Additionally, since $S \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$, $T' - T$ must be great then τ. Therefore, starting from $x' \in S$, we have $\phi(t, x') \in \mathcal{X}$ and $\phi(t, x') \notin \mathcal{S}$ for all $t \in (0, T]$, which contradicts with the assumption that S is locally τ-recurrent. $\qquad\square$

3.2.5 Sampling-based safety verification

To leverage Theorem 3.16 for ensuring that a set is locally τ-recurrent within $\mathcal{X}$, we need a LREBF h that satisfies condition (3.31) within $\mathcal{X}_\tau$. Subsequently, Theorem 3.19 shows that when $h_{\geq 0} \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$, local safety is guaranteed. Therefore, applying these theorems generally involves four steps, characterizing $\mathcal{X}_\tau$, defining h, verifying condition (3.31) within $\mathcal{X}_\tau$, and verifying that $h_{\geq 0} \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$. It should be noted, however, that numerically testing these conditions is impossible, as they require testing an uncountable set of states. Instead, we will seek to check whether a point x satisfies a set of stricter conditions, which allows certifying its neighborhood. of x.

In this section, we discuss how to certify such a neighborhood: a ball $\mathcal{B}_r(x)$ centered at $x \in \mathcal{X}$ with radius $r > 0$. This consists of verifying the set inclusion

relationships of $\mathcal{B}_r(x)$ w.r.t. $\mathcal{X}_\tau$ and $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$, and then verifying condition (3.31) within $\mathcal{X}_\tau$. We will only need the trajectory originating from x to check these conditions.

We will use $L := L(D)$ to denote the local Lipschitz constant over D in the subsequent lemma and corollary to bound the differences in signed distances between points that are sufficiently close to each other.

Lemma 3.6. *Consider two points x and y. If the distance between x and y is bounded by r, i.e., $\|x - y\| \leq r$, then the signed distances from points x, y to a set S satisfies:*

$$\mathrm{sd}(y, S) \geq \mathrm{sd}(x, S) - r. \tag{3.43}$$

Proof. See Appendix (Section 3.4). □

Corollary 3.1. *Starting from any two points x, y such that $\|x - y\| \leq r$, the trajectory satisfies:*

$$\mathrm{sd}(\phi(t, y), S) \geq \mathrm{sd}(\phi(t, x), S) - re^{Lt}, \ \forall t \geq 0.$$

Proof. Since the vector field is locally one-sided Lipschitz, Theorem 3.9 in [64] implies that

$$\|\phi(t, x) - \phi(t, y)\| \leq e^{Lt}\|x - y\| \leq re^{Lt}, \ \forall t \geq 0. \tag{3.44}$$

Therefore, the result follows from Lemma 3.6. □

By applying Corollary 3.1, one can verify whether $\mathcal{B}_r(x)$ is a subset of, or intersects with, $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$, using only the trajectory originating from x.

Theorem 3.20. *Consider an initial state $x \in \mathcal{X}$, then*

(i) if x satisfies

$$\mathrm{sd}(\phi(t, x), \mathcal{X}_u) > re^{Lt}, \quad \forall t \in [0, \tau], \tag{3.45}$$

for some $r > 0$, then $\mathcal{B}_r(x) \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$.

(ii) Conversely, if x satisfies

$$\exists t \in [0, \tau] \ \ s.t. \ \ \mathrm{sd}(\phi(t, x), \mathcal{X}_u) < -re^{Lt}, \tag{3.46}$$

for some $r > 0$, then $\mathcal{B}_r(x) \subseteq \mathcal{R}_\tau^{-1}(\mathcal{X}_u)$.

Proof. See Appendix (Section 3.4). □

Similarly, one can verify that the trajectories starting from $\mathcal{B}_r(x)$ remain inside, or intersect with, $\mathcal{X}_\tau$, using the trajectory starting from x.

Theorem 3.21. *Consider an initial state $x \in \mathcal{X}$, then*

(i) if x satisfies

$$\exists t \in [0, \tau] \ \ s.t. \ \ \mathrm{sd}(\phi(t, x), \mathcal{X}) > re^{Lt}, \tag{3.47}$$

for some $r > 0$, then $\mathcal{B}_r(x) \cap \mathcal{X}_\tau = \emptyset$.

(ii) Conversely, if x satisfies

$$\mathrm{sd}(\phi(t, x), \mathcal{X}) < -re^{Lt}, \quad \forall t \in [0, \tau], \tag{3.48}$$

for some $r > 0$, then $\mathcal{B}_r(x) \subseteq \mathcal{X}_\tau$.

Once ensured that $\mathcal{B}_r(x) \subseteq \mathcal{X}_\tau$ and $\mathcal{B}_r(x) \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$, we proceed to define a LREBF $h(\cdot) := -\mathrm{sd}|_{\mathcal{X}}(\cdot, S)$, as outlined in Theorem 3.18. Our next goal is to verify that h satisfies condition (3.31) within $\mathcal{B}_r(x)$. To this end, we will leverage the following lemma that is analogous to Lemma 3.6 and Corollary 3.1.

Lemma 3.7. *Consider a convex set $\mathcal{X}$ and two states $x, y \in \mathcal{X}$. If the distance between x and y is bounded by r, i.e., $\|x - y\| \leq r$, then the constrained signed distances to a set S from x and y satisfy:*

$$\mathrm{sd}|_{\mathcal{X}}(y, S) \geq \mathrm{sd}|_{\mathcal{X}}(x, S) - r. \tag{3.49}$$

Additionally, if $x, y \in \mathcal{X}_\tau$, then the trajectories starting from them satisfy:

$$\mathrm{sd}|_{\mathcal{X}}(\phi(t, y), S) \geq \mathrm{sd}|_{\mathcal{X}}(\phi(t, x), S) - re^{Lt}, \ \forall t \in [0, \tau]. \tag{3.50}$$

Proof. See Appendix (Section 3.4). □

By applying Lemma 3.7, we finially verify whether h satisfies condition (3.31) within $\mathcal{B}_r(x)$, using only the trajectory originating from x.

Theorem 3.22. *Given a closed set S, a convex region of interest $\mathcal{X}$ and let $h(\cdot) := -\mathrm{sd}|_{\mathcal{X}}(\cdot, S)$.*

(i) Assume that there is a point x that satisfies

$$\hat{h}_r^-(t, x) := h(\phi(t, x)) - re^{Lt}$$

and

$$\max_{t \in (0, \tau]} e^{\alpha t} [\hat{h}_r^-(t, x)]_+ + e^{\beta t} [\hat{h}_r^-(t, x)]_- \geq h(x) + r, \tag{3.51}$$

for some $r > 0$. Then, for all $y \in B_r(x)$, the LREBF condition is satisfied, i.e.,

$$\max_{t \in (0, \tau]} e^{\alpha t} [h(\phi(t, y))]_+ + e^{\beta t} [h(\phi(t, y))]_- \geq h(y).$$

(ii) Assume that there is a point x that satisfies

$$\hat{h}_r^+(t, x) := h(\phi(t, x)) + re^{Lt}$$

and

$$\max_{t \in (0, \tau]} e^{\alpha t} [\hat{h}_r^+(t, x)]_+ + e^{\beta t} [\hat{h}_r^+(t, x)]_- < h(x) - r, \tag{3.52}$$

for some $r > 0$. Then, for all $y \in B_r(x)$, the LREBF condition is not satisfied, i.e.,

$$\max_{t \in (0, \tau]} e^{\alpha t} [h(\phi(t, y))]_+ + e^{\beta t} [h(\phi(t, y))]_- < h(y).$$

Proof. See Appendix (Section 3.4). □

3.3 Numerical Methods

In this section, inspired by the Theorem 3.16,3.18, and 3.19 we will develop an algorithm that, given region of interest $\mathcal{X}$, characterize a region $h_{\geq 0} \cap \mathcal{X}_\tau$ that is locally safe within $\mathcal{X}$. The algorithm covers $\mathcal{X}$ with a grid G, where each grid cell $g \in G$ is represented as an infinity norm ball $\mathcal{B}_r(x)$ centered at the grid cell center x, aiming to utilize the highly parallelizable processing units to certify many points simultaneously.

Precisely, in Section 3.3.1, we leverage the results in Section 3.2.5 to develop an algorithm to verify the neighborhood of a point x, i.e., $\mathcal{B}_r(x)$. Then, Section 3.3.2 proposes a parallelizable algorithm that allows us to learn the local safe region and the safe boundary approaching rates α, β within the entire region $\mathcal{X}$. Finally, an example of characterizing the safety of a 2D system is provided in Section 3.3.3.

3.3.1 Algorithmic Verification of a Ball

In Algorithm 2, we describe the procedure for verifying a neighborhood $\mathcal{B}_r(x) \subseteq \mathcal{X}$ around a given point x, assuming that a candidate LREBF function h and a specified radius $r > 0$ are provided.

By leveraging conditions (3.45), (3.48), and (3.51) respectively, our goal is to ascertain whether three critical objectives are met for every $y \in \mathcal{B}_r$. These objectives include verifying that $y \in \mathcal{X}_\tau$, $y \notin \mathcal{R}_\tau^{-1}(\mathcal{X}_u)$, and confirming that the LREBF condition (3.31) is satisfied. If all these conditions are met, we will successfully verify the safety of the ball and return true.

Conversely, we return false if we can confirm any of the following contrary conditions: $\mathcal{B}_r \cap \mathcal{X}_\tau = \emptyset$, $\mathcal{B}_r \subseteq \mathcal{R}_\tau^{-1}(\mathcal{X}_u)$, or the LREBF condition (3.31) is not satisfied by any $y \in \mathcal{B}_r$. These conditions are encoded in (3.46), (3.47), and (3.52). In instances where a definitive conclusion cannot be reached regarding any of these conditions,

we classify the result as unknown and seek to refine this neighborhood later in Algorithm 4.

Algorithm 2: VerifyBall(g, τ, L, h)

Input Grid cell $g = \mathcal{B}_r(x) \in G$, $\tau > 0$, $L > 0$
Verify $\mathcal{B}_r(x) \subseteq \mathcal{X}_\tau$ by checking exceptions
if (3.47), *i.e.,* $\mathcal{B}_r(x) \cap \mathcal{X}_\tau = \emptyset$ **then**
$\quad$ **return** (False, Not in $\mathcal{X}_\tau$)
if not (3.48), *i.e.,* $\mathcal{B}_r(x) \not\subseteq \mathcal{X}_\tau$ **then**
$\quad$ **return** (Unknown, $\mathcal{X}_\tau$ unknown)

Verify $\mathcal{B}_r(x) \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$ by checking exceptions
if (3.46), *i.e.,* $\mathcal{B}_r(x) \subseteq \mathcal{R}_\tau^{-1}(\mathcal{X}_u)$ **then**
$\quad$ **return** (False, Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$)
if not (3.45), *i.e.,* $\mathcal{B}_r(x) \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) \neq \emptyset$ **then**
$\quad$ **return** (Unknown, $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$ unknown)

Verify LREBF condition (3.31) *by checking exceptions*
if (3.52), *i.e., verified* $h_{\geq 0}$ *NOT recurrent* **then**
$\quad$ **return** (False, Not recurrent)
if not (3.51), *i.e., unable to verify recurrence* **then**
$\quad$ **return** (Unknown, Recurrence unknown)

Definite success
return (True, Verified a ball)

3.3.2 Verification of a Region

Having developed an algorithm to verify a ball $B_r(x)$, c.f. Algorithm 2, we are now ready to integrate it into an algorithm that can characterize a region S that is locally safe w.r.t. $\mathcal{X}$. Our algorithm divides G into G^+ and G^-, and defines $h(\cdot) := -\text{sd}|_{\mathcal{X}}(\cdot, G^-)$. Consequently, all points $x \in G^+$ (resp. G^-) will have $h(x) \geq 0$ (resp. $h(x) < 0$).

Splitting cells for more accurate verification

When Algorithm 2 returns 'Unknown.', we will refine the local section of the grid by splitting the ball $B_r(x)$ into 3^d balls. We refer the reader to Algorithm 3 for details

and to Figure 3-4 for an illustration for $d = 2$.

Algorithm 3: Splitting a Ball
***
Input Grid cell $g = \mathcal{B}_r(x) \in G$
$(x_1, ..., x_n) = x$
$P = \{(x), (x_1 \pm (2/3)r, ..., x_n), ..., (x_1, ..., x_n \pm (2/3)r)\}$
return Set of cells $\{\mathcal{B}_{2r/3}(p) \mid p \in P\}$
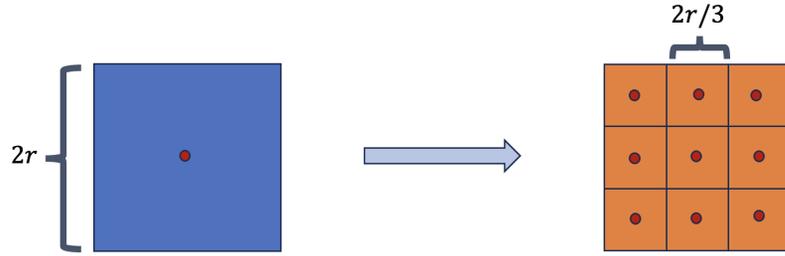
Figure 3-4. Splitting a Ball according to Algorithm 3

Algorithm Summary

The steps outlined above culminate in an algorithm that enables us to characterize a region $h_{\geq 0} \cap \mathcal{X}_\tau$ as locally safe within $\mathcal{X}$. This proposed algorithm, detailed in Algorithm 4, utilizes parallel processing capabilities of GPUs to efficiently verify multiple points simultaneously. In short, Algorithm 4 establishes an initial grid G and proceeds to verify each cell $g \in G$.

In this process, it removes g from G in the case that 'Not in $\mathcal{X}_\tau$,' i.e., condition (3.47) is satisfied. If a $g \in G^+$ is 'Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$,' or 'Not recurrent', i.e., condition (3.46) or (3.52), it moves g to G^-. It then recursively splits the 'unknown' grid cells until each cell is either verified or reaches the minimal radius. To elaborate, cells classified as '$\mathcal{X}_\tau$ unknown' (resp. '$\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$ unknown', 'Recurrence unknown') are approximately reclassified as 'Inside $\mathcal{X}_\tau$' (resp. 'Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$', 'Not recurrent') if their radius is smaller than the minimal threshold. This approximation ensures the robustness of the learning outcomes.

Algorithm 4: Verification of a Region

Input $L, \varepsilon > 0$
G = Grid of balls (x, r) covering $\mathcal{X}$
$G^- = \{g \in G \,|\, g \subseteq \mathcal{X}_u\}$, $G^+ = G \backslash G^-$
while *true* **do**

    (verified, info) = VerifyBall(G, τ, L)
    Unknown_cells = G [verified == Unknown]

    Remove G [info == 'Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$'] from G^+
    Add G [info == 'Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$'] to G^-
    Select small '$\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$ unknown' cells
    $G_\varepsilon = G$ [info == '$\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$ unknown', $r \leq \varepsilon$]
    Treat them as 'Inside $\mathcal{R}_\tau^{-1}(\mathcal{X}_u)$' cells
    Remove G_ε from G^+; Add G_ε to G^-

    Remove G [info == 'Not recurrent'] from G^+
    Add G [info == 'Not recurrent'] to G^-
    Select small 'Recurrence unknown' cells
    $G_\varepsilon = G$ [info == 'Recurrence unknown', $r \leq \varepsilon$]
    Treat them as 'Not recurrent' cells
    Remove G_ε from G^+; Add G_ε to G^-

    Remove G [info == 'False $\mathcal{X}_\tau$'] from G, G^+, G^-
    Select small '$\mathcal{X}_\tau$ unknown' cells
    $G_\varepsilon = G$ [info == '$\mathcal{X}_\tau$ unknown', $r \leq \varepsilon$]
    Treat them as 'Inside $\mathcal{X}_\tau$' cells
    Remove G_ε from Unknown_cells

    if Unknown_cells *is empty* **then**
      └ **return** True
    Remove Unknown_cells[$r > \varepsilon$] from G
    Add Split(Unknown_cells[$r > \varepsilon$]) to G

Convergence rates bounds

Recall that α and β represent the upper and lower bounds on the boundary-approaching rates from the sets $h_{\geq 0}$ and $h_{<0}$, respectively, serving as measures of safety levels. Therefore, in verifying the LREBF condition, we seek to identify the smallest possible $\alpha > 0$ and the largest possible $\beta > 0$ such that the condition (3.51) is satisfied for each grid cell g in G^+ and G^-, respectively. Therefore, we define

$$\alpha(g) = \inf\{\alpha > 0 \mid \text{Cell-wise condition 3.51 is true}\}$$
$$\beta(g) = \sup\{\beta > 0 \mid \text{Cell-wise condition 3.51 is true}\}.$$

It is important to note that the left-hand side of (3.51) is monotonically increasing as α increases for grid cells in G^+ (resp. decreasing as β decreases for grid cells in G^-). Therefore, employing a simple line search algorithm allows for the optimal determination of α or β for each cell.

Algorithm 5: Find best α, β for all cells

Input $L, \varepsilon, \delta_e > 0$; max_err $\in (0,1)$
G = Grid of balls (x, r) covering $\mathcal{X}$
G^+, G^-
for *each g in G^+ with $r > \varepsilon$* **do**
 $\quad$ err$(g) = (\alpha(g) - \underline{\alpha}(g))/\underline{\alpha}(g)$
 $\quad$ **if** *err$(g) >$ max_err **or** $\alpha(g) > \underline{\alpha}(g) + \delta_e$* **then**
 $\quad\quad$ Remove g from G^+; Add Split(g) to G^+

for *each g in G^- with $r > \varepsilon$* **do**
 $\quad$ err$(g) = (\bar{\beta}(g) - \beta(g))/\bar{\beta}(g)$
 $\quad$ **if** *err$(g) >$ max_err **or** $\beta(g) < 0$* **then**
 $\quad\quad$ Remove g from G^-; Add Split(g) to G^-

Find the best α, β

In addition to finding α, β such that the cell-wise LREBF condition (3.51) is satisfied, we can also determine the optimal bounds, $\underline{\alpha}$ and $\bar{\beta}$, specifically for the point x by applying the original LREBF condition (3.31). Therefore, we further define

$$\underline{\alpha}(g) = \inf\{\alpha > 0 \mid \text{Point-wise condition (3.31) is true}\}$$
$$\bar{\beta}(g) = \sup\{\beta > 0 \mid \text{Point-wise condition (3.31) is true}\}.$$

Algorithm 6: Find best α, β for the region

Input $L, \varepsilon > 0$; max_err $\in (0, 1)$

G = Grid of balls (x, r) covering $\mathcal{X}$

G^+, G^-

while *true* **do**
    worst-$\alpha = \max\{\alpha(g) \mid g \in G^+\}$
    worst-$\underline{\alpha} = \max\{\underline{\alpha}(g) \mid g \in G^+\}$
    err = (worst-α − worst-$\underline{\alpha}$)/worst-$\underline{\alpha}$
    $G' := k$ cells in G^+ with largest α
    $r_0 :=$ radius of the cell with largest α
    if *err > max_err **and** $r_0 > \varepsilon$* **then**
        Remove each $g \in G'$ from G^+;
        Add each Split(g), $g \in G'$ to G^+
    else
        break

while *true* **do**
    worst-$\beta = \min\{\beta(g) \mid g \in G^-\}$
    worst-$\bar{\beta} = \min\{\bar{\beta}(g) \mid g \in G^-\}$
    worst_err = (worst-$\bar{\beta}$ − worst-β)/worst-$\bar{\beta}$
    $G' := k$ cells in G^+ with smallest β
    $r_0 :=$ radius of the cell with smallest β
    if *err > max_err **and** $r_0 \geq \varepsilon$* **then**
        Remove each $g \in G'$ from G^+;
        Add each Split(g), $g \in G'$ to G^+
    else
        Success

It is important to note that, by definition, $\bar{\beta} \geq \beta$ and $\underline{\alpha} \leq \alpha$, with both $\alpha \to \underline{\alpha}$ and $\beta \to \bar{\beta}$ as the radius $r \to 0$. This convergence behavior motivates a further subdivision of the ball $B_r(x)$ to refine our grid. This refined grid allows for a more precise determination of α or β, enhancing the granularity and accuracy of our safety analysis. To algorithmically achieve this refinement and optimization, we introduce Algorithms 5 and 6. Algorithm 5 is designed to determine the optimal α and β for each individual cell, ensuring localized precision. Conversely, Algorithm 6 targets the identification of the region-wise optimal α and β.

3.3.3 Characterizing safety of a 2D system

We end this section by providing a preliminary validation of the proposed algorithm. To investigate the efficiency of our proposed method, we consider the following system:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}. \tag{3.53}$$

We fix the set of unsafe states to be $\mathcal{X}_u = \{x \in \mathbb{R}^2 \mid (x_1 + 1)^2 + (x_2 + 1)^2 \leq 0.16\}$. We seek to characterize a local safe region within the region of interest $\mathcal{X} := [-3.5, 3.5] \times [-3.5, 3.5]$. In our experiments, we use a 7×7 initial grid to cover $\mathcal{X}$ and simulate the trajectory for a finite length $\tau = 0.6s$.

The outcome of Algorithm 4 as well as sample trajectories for the system (3.53) can be seen in Figure 3-5. Specifically, the set of unsafe states is visually demarcated in red within the figure, clearly indicating areas that must be avoided. The union of grid cells marked in green, referred to as G^+, illustrates our algorithmically determined local safe region.

Figure 3-6 displays the verified boundary-approaching rates for trajectories originating from both the safe ($\mathcal{X}_\tau \cap h_{\geq 0}$) and the unsafe ($\mathcal{X}_\tau \cap h_{<0}$) regions as they approach the local safety boundary defined by $h_{=0} \cap \text{int}\mathcal{X}$. In the upper panel of Figure 3-6, the intensity of the green color is used to indicate the level of safety for
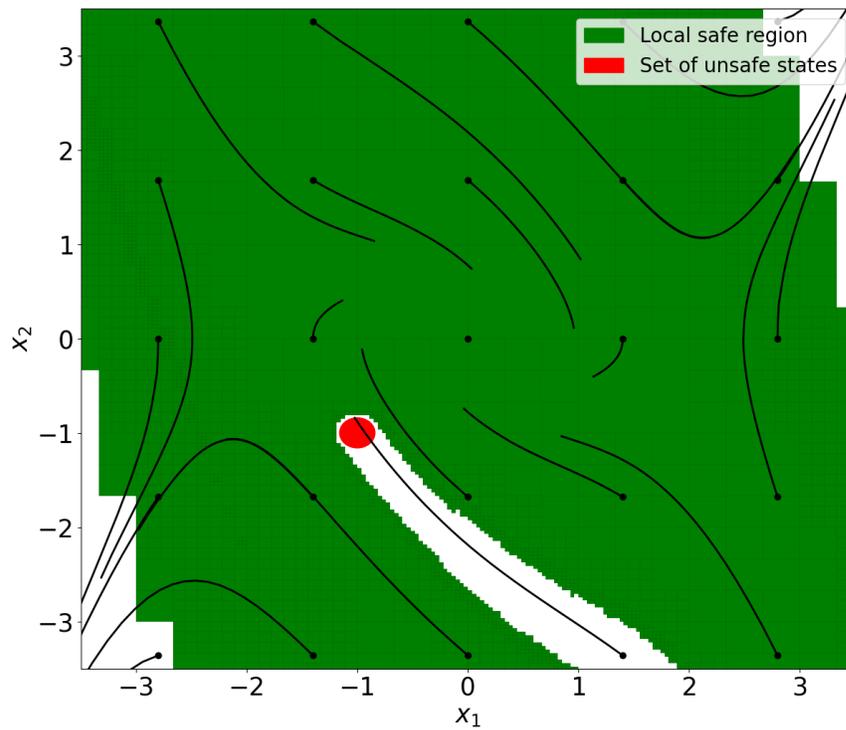
Figure 3-5. The learning outcome of Algorithm 4. Unsafe regions are highlighted in red. The green cells, collectively forming G^+, represent the areas identified as safe through our experiment.

different cells within the local safe region $\mathcal{X}_\tau \cap h_{\geq 0}$. Darker green signifies a higher safety level, attributed to a smaller required α for satisfying the cell-wise condition (3.51). This implies a stricter (smaller) upper bound on the boundary-approaching rates can be posed on trajectories starting from this safe cell. Conversely, lighter green cells demonstrate larger α, indicating that trajectories from these cells may approach the boundary more rapidly.

Similarly, in the lower panel of Figure 3-6, the safety levels within the region $\mathcal{X}_\tau \cap h_{<0}$ are represented by shades of brown. Here, darker brown denotes a higher safety level, i.e., the cells satisfy condition (3.51) with a larger β. This indicates that trajectories from these cells approach the boundary, i.e., return back to the safe set, at a faster rate, allowing for a stricter (larger) upper bound on these rates. Lighter brown, however, points to a lower safety level with a smaller β, signifying that some trajectories may approach the boundary more slowly.

Note that in both panels of Figure 3-6, cells with a larger radius generally exhibit a lower safety level. This trend arises because verifying tighter bounds on the boundary-approaching rates becomes more challenging as the cell size increases. This issue is particularly pronounced in the region of $\mathcal{X}_\tau \cap h_{<0}$, where some cells even record a $\beta \leq 0$. Such values suggest that trajectories originating from these cells may not be approaching the safe boundary as expected, indicating potential safety concerns. To address this, we consider the application of Algorithm 5 or 6. These algorithms aim to optimize the grid structure and accurately determine the best α and β values for each cell, thereby enhancing reliability and precision.

In the rest of our experiments, we set the max_err $= 30\%$, and we stop splitting if the radius is smaller than $\varepsilon = 0.01$. The outcomes of Algorithms 5 and 6 are presented in Figures 3-7 and 3-8, respectively, showcasing their distinct approaches and effectiveness in approximating the boundary-approaching rates. Figure 3-7 highlights the high precision resolution achieved by Algorithm 5, which diligently
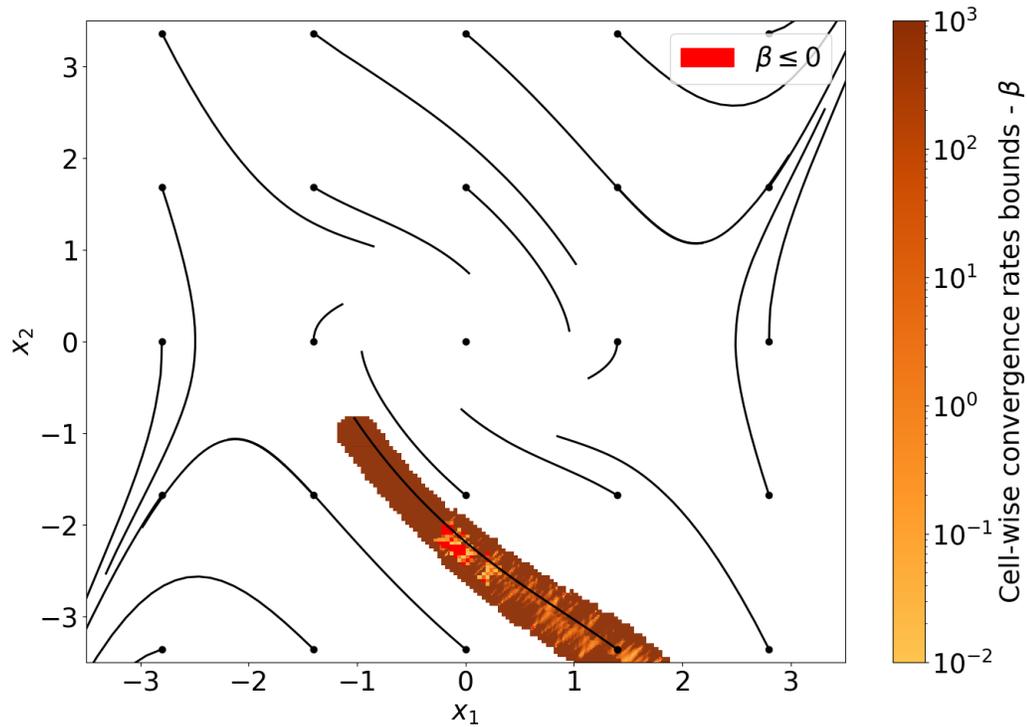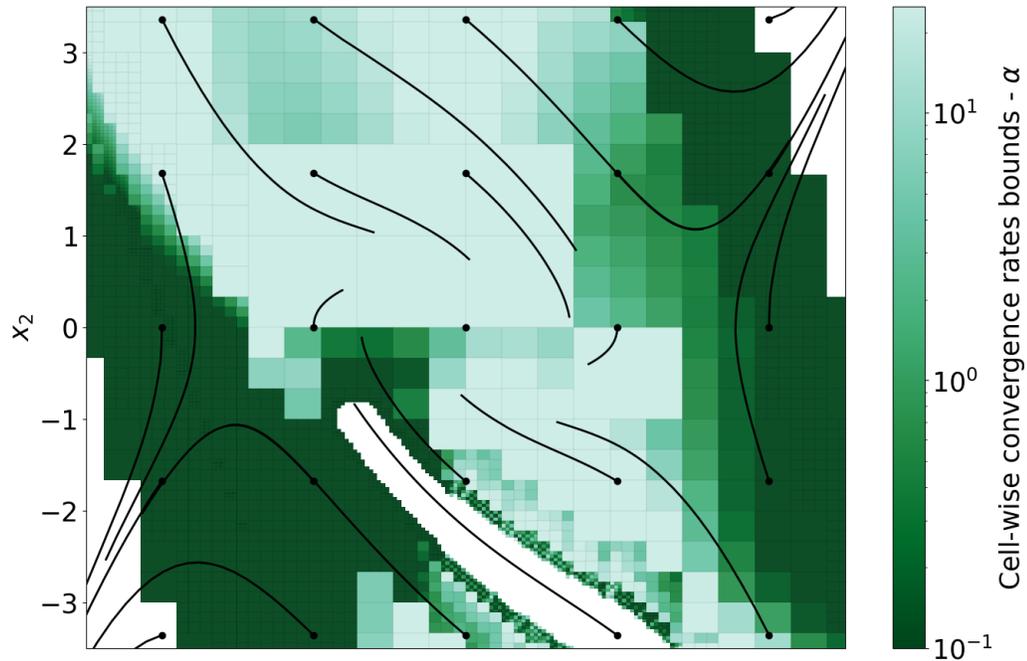
Figure 3-6. The verified boundary-approaching rates for trajectories originating from both the safe ($\mathcal{X}_\tau \cap h \geq 0$) and the unsafe ($\mathcal{X}_\tau \cap h < 0$) regions as they approach the local safety boundary defined by $h_{=0} \cap \mathrm{int}\,\mathcal{X}$. In the upper panel, darker green signifies a higher safety level. While in the lower panel, darker brown denotes a higher safety level.

minimizes the error $err(g)$ for all cells g within the grid G. This results in a detailed and finely-tuned safety analysis across the entire region. Conversely, Algorithm 6, depicted in Figure 3-8, optimizes a singular error value, err, leading to fewer cell divisions while still ensuring the worst-case boundary-approaching rates are maintained, indicated by the strategic placement of red boxes to denote critical areas. The comparative efficiency of these two methodologies is further explored in Figure 3-9, which provides a clear visual comparison of their operational efficiency.
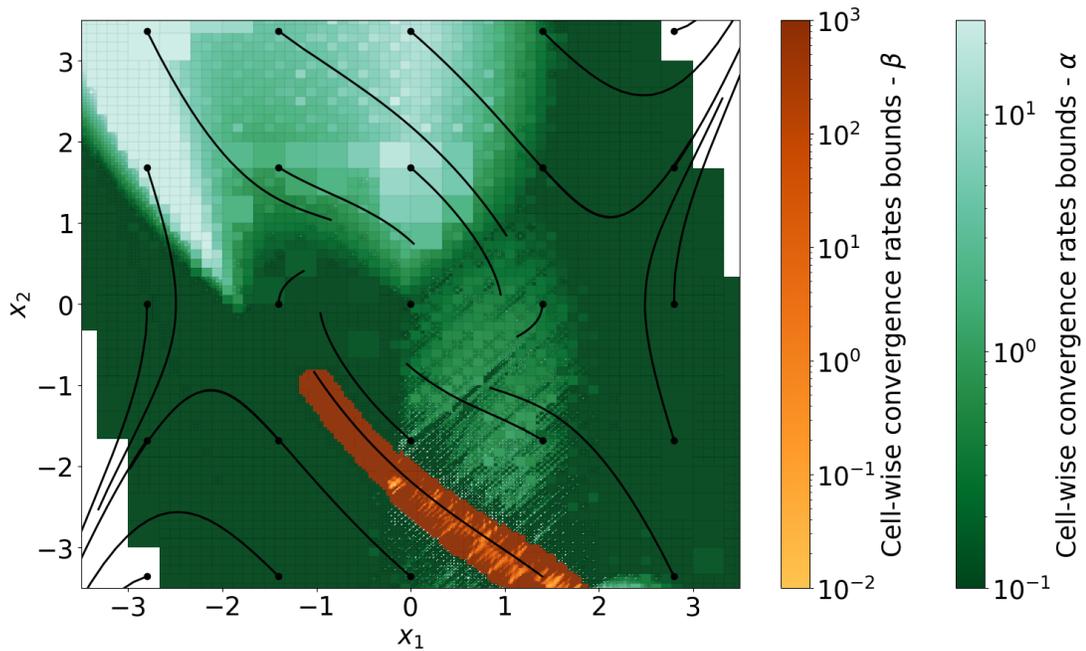
Figure 3-7. This figure displays the detailed outcomes of Algorithm 5, showcasing its high precision resolution. The algorithm focuses on reducing the error measurement $err(g)$ for each cell g in grid G, resulting in a finer resolution of safety assessments across the grid.
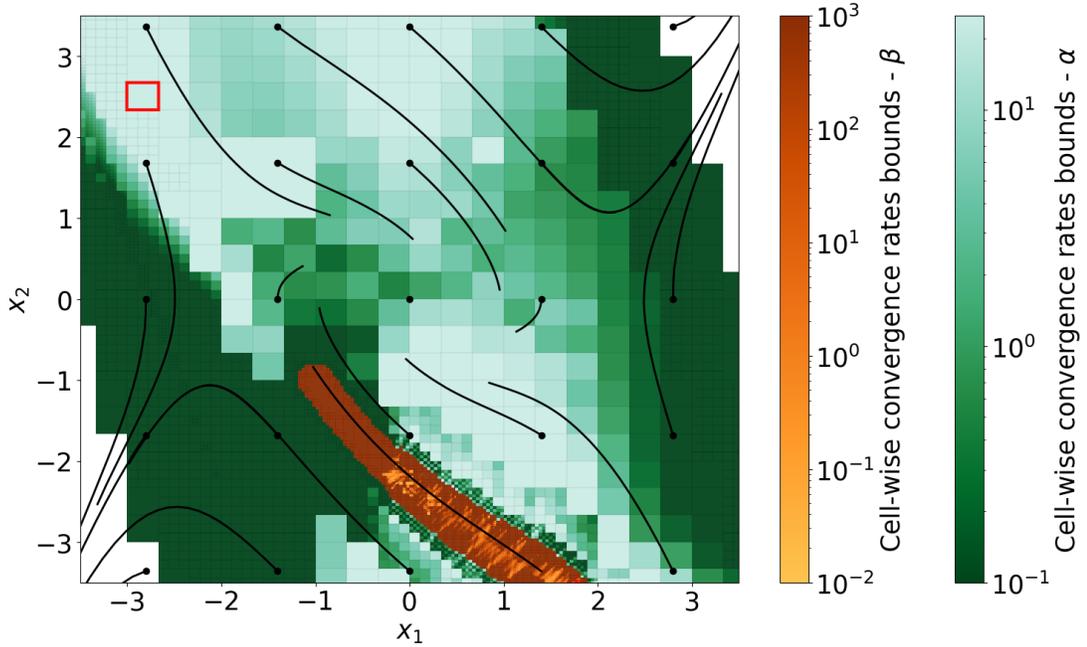
Figure 3-8. Depicted here are the results of Algorithm 6, which aims to reduce a single error value, err, across the entire grid. The red boxes highlight the cells that give the worst α and β.
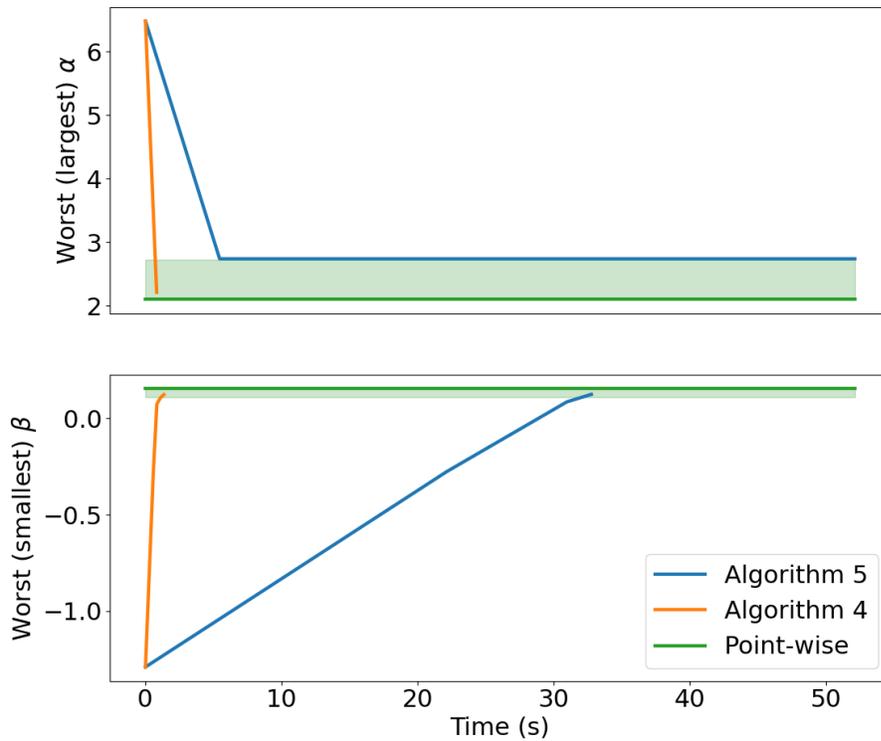
Figure 3-9. This figure provides a comparative analysis of the efficiency between Algorithm 5 and Algorithm 6, illustrating the trade-offs between grid resolution and computational simplicity.

3.4 Appendix

Proof of Lemma 3.1

Proof. Rearranging condition (3.4) gives:

$$g(x) := L_f h(x) - (-\alpha)h(x) \geq 0, \ \forall x \in D_0$$

We then have the following differential equation:

$$\frac{dh(\phi(t,x))}{dt} = L_f h(\phi(t,x))$$
$$= g(\phi(t,x)) + (-\alpha)h(\phi(t,x)) \tag{3.54}$$

Solving (3.54) gives:

$$h(\phi(t,x)) = e^{-\alpha t}(\int_0^t g(\phi(s,x))e^{\alpha s}ds + h(x)). \tag{3.55}$$

Note that condition (3.4) ensures $h_{\geq 0}$ is invariant and the function value $h(\phi(t,x))$ is strictly increasing along the trajectory starting from $x \in h_{<0} \cap D_0$. Therefore, D_0 is also an invariant set, and thus

$$\int_0^t g(\phi(s,x))ds \geq 0, \ \forall t \geq 0, x \in D_0. \tag{3.56}$$

By applying (3.56) to (3.55), one can conclude that:

$$h(\phi(t,x)) \geq e^{-\alpha t}h(x), \quad \forall t \geq 0, \ x \in D_0,$$

and the result follows. □

Proof of Theorem 3.3

Proof. (i): Starting from any initial state $x \in h_{=0}$, condition (3.5) requires $L_f h(x) \geq 0$, i.e., condition (3.2). According to Theorem 3.1, this guarantees the invariance of the super-level set $h_{\geq 0}$.

(ii): As demonstrated in the proof of Lemma 3.1, the following two conditions hold for all $x \in D_0 \cap h_{<0}$:

$$g(x) = L_f h(x) - (-\beta)h(x) \geq 0,$$
$$h(\phi(t, x)) = e^{-\beta t}\left(\int_0^t g(\phi(s, x))e^{\beta s}ds + h(x)\right). \tag{3.57}$$

Note that the super-level set $h_{\geq 0}$ is invariant, and the function value $h(\phi(t, x))$ increases strictly along the trajectory starting from any $x \in D_0 \cap h_{<0}$. Hence, D_0 is an invariant set.

Now, starting from $x \in h_{<0} \cap D_0$, if there is a time t' with $h(\phi(t', x)) \geq 0$, then condition (3.6) is automatically satisfied for any time $t \geq t'$. Thus, w.l.o.g. we concentrate on the scenario where $h(\phi(t, x)) < 0$ for all $t \geq 0$. In this case, $\phi(t, x)$ remains within $D_0 \cap h_{<0}$, and consequently:

$$\int_0^t g(\phi(s, x))ds \geq 0, \quad \forall t \geq 0. \tag{3.58}$$

By applying (3.58) to (3.57), we successfully verify (3.7).

(iii): This result is trivially obtained by applying Lemma 3.1 with $D_0 = h_{\geq 0}$. $\quad\square$

Proof of Theorem 3.4

Proof. ($\Longrightarrow$): This direction follows directly from the definition of the invariant set. Precisely, the super-level set $h_{\geq 0}$ being invariant implies:

$$\phi(t, x) \in h_{\geq 0} \implies h(\phi(t, x)) \geq 0,$$

for all $t \geq 0$ and $x \in h_{\geq 0}$.

($\Longleftarrow$): Suppose $h_{\geq 0}$ is not invariant. Then there must exist an initial state $x \in h_{\geq 0}$ and a time instant $t^* \geq 0$, such that $\phi(t^*, x) \notin h_{\geq 0}$. We then use t' to denote the last time the trajectory $\phi(t, x)$ stays within the closed set $h_{\geq 0}$ before t^*, i.e.,

$$x' := \phi(t', x) \in h_{=0} \quad \text{and} \quad h(\phi(t, x)) < 0, \quad \forall t \in (t', t^*].$$

86

This contradicts condition (3.8), which requires $h(\phi(t, x'))$ to always remain non-negative when starting from $h(x') = h(\phi(t', x)) = 0$. Therefore, the result follows.

$\square$

Proof of Theorem 3.5

Proof. *(i):* Suppose $h_{\geq 0}$ is not invariant. As stated in the proof of Theorem 3.4, there exists a $t \geq 0$ and $x \in h_{=0}$ such that:

$$h(\phi(s, x)) < 0, \quad \forall s \in (0, t].$$

However, condition (3.9) requires that $h(\phi(t, x)) \geq h(x) - \int_0^t \zeta(h(\phi(s, x)))ds$, and the right-hand side would be non-negative in this case. This is a contradiction.

(ii): We first evaluate the following Taylor expansion for t around 0:

$$\int_0^t \zeta(h(\phi(s, x)))ds = 0 + \zeta(h(\phi(0, x)))t + o(t).$$

This, together with condition (3.9), implies the following bound on the convergence rate expressed in the form of the lower-right Dini derivative [69]:

$$
\begin{aligned}
D_+ h(x) &:= \liminf_{t \to 0} \frac{h(\phi(t, x)) - h(x)}{t} \\
&\geq \liminf_{t \to 0} \frac{-\int_0^t \zeta(h(\phi(s, x)))ds}{t} \\
&= \liminf_{t \to 0} \frac{0 - \zeta(h(\phi(0, x)))t - o(t)}{t} \\
&= -\zeta(h(x)), \quad \forall x \in D_0.
\end{aligned}
$$

$\square$

Proof of Theorem 3.6

Proof. (i): Suppose $h_{\geq 0}$ is not invariant. As stated in the proof of Theorem 3.4, there exists a $t \geq 0$ and $x \in h_{=0}$ such that:

$$h(\phi(s, x)) < 0, \quad \forall s \in (0, t].$$

This contradicts condition (3.10), which requires the function value $h(\phi(t, x)) \geq h(x)e^{-\alpha t} = 0$ for all $t \geq 0$. Therefore, the result follows.

(ii): This part follows trivially given condition (3.10). □

Proof of Theorem 3.7

Proof. ($\Longrightarrow$): This direction follows directly from the definition of a τ-recurrent set. Precisely, the super-level set $h_{\geq 0}$ being τ-recurrent implies that for all $x \in h_{\geq 0}$,

$$\exists t' \in (0, \tau] \text{ s.t. } \phi(t', x) \in h_{\geq 0} \implies \max_{t \in (0, \tau]} h(\phi(t, x)) \geq 0.$$

($\Longleftarrow$): Suppose the closed set $h_{\geq 0}$ is not τ-recurrent. There must exist an initial state $x \in h_{=0}$ such that $\phi(t, x) \notin h_{\geq 0}$ for all $t \in (0, \tau]$. In this case, $h(\phi(t, x)) < h(x) = 0$ for all $t \in (0, \tau]$, which contradict with condition (3.11). Therefore, $h_{\geq 0}$ is τ-recurrent. □

Proof of Lemma 3.2

Proof. Given $x \in D_0$, we build the time sequence $\{t_n\}_{n \in \mathbb{N}}$ satisfying (3.13) and (3.14), following an inductive method similar to that detailed in our previous work [1, Lem 1].

[Base case]: For the base case, we have $t_0 = 0$, $x_0 = x \in D_0$, and define t_1 as follows:

$$t_1 = \max\{\arg\max_{t \in (0, \tau]}\{h(\phi(t, x_0)) + \int_0^t \zeta(h(\phi(s, x_0)))ds\}\}; \tag{3.59}$$

note that the second maximum exists by condition (3.12), and is no smaller than $h(x_0)$; if there are multiple maximizing times, t_1 is defined as the largest. By construction, $t_1 - t_0 \in (0, \tau]$, and the function h evaluated at $x_1 := \phi(t_1, x_0)$ satisfies:

$$h(x_1) \geq h(x_0) - \int_0^{t_1} \zeta(h(\phi(s, x_0)))ds = h(x_0) + \delta_0,$$

thus confirming (3.14) with the left argument on the $\max$ of the right-hand side.

To prove the right argument on the $\max$ of equation (3.14), one need to first show that:

$$h(x_{n+1}) \begin{cases} > h(x_n) \text{ if } x_n \in h_{<0} \cap D_0 & \text{(3.60a)} \\ \geq 0 \qquad \text{if } x_n \in h_{\geq 0}, & \text{(3.60b)} \end{cases}$$

Let us first consider the case $x_0 \in h_{<0} \cap D_0$. If $h(x_1) \geq 0$, then (3.60a) follows trivially. If $h(x_1) < 0$, we argue that $h(\phi(t, x_0)) < 0$ and $\zeta(h(\phi(t, x_0))) < 0$ for all $t \in [0, t_1]$; otherwise, t_1 would not maximize (3.59). Hence, $\delta_0 > 0$, and we verified $h(x_1) \geq h(x_0) + \delta_0 > h(x_0)$, thereby satisfying (3.60a).

In the other case that $x_0 \in h_{\geq 0}$, we demonstrate that $h(x_1) \geq 0$ by contradiction. Suppose $h(x_0) \geq 0$ and $h(x_1) < 0$, we use t' to denote the last time the trajectory $\phi(t, x_0)$ stays within the closed set $h_{\geq 0}$ before t_1, i.e.,

$$x' := \phi(t', x_0) \in h_{=0} \quad \text{and} \quad h(\phi(t, x_0)) < 0, \ \forall t \in (t', t_1].$$

This contradicts with the fact that t_1 is a maximizer of (3.59), since $h(x') = 0 > h(x_1)$ and $\int_0^{t'} \zeta(h(\phi(s, x_0))) ds > \int_0^{t_1} \zeta(h(\phi(s, x_0))) ds$. Therefore, we have $h(x_1) \geq 0$, and thus (3.60b) follows.

[Inductive step]: Note that condition (3.14) further implies $x_1 = \phi(t_1, x_0) \in D_0$, since $h(x_0) \geq -c$. Thus, the inductive construction proceeds in a similar manner: given $t_1 < t_2 < \cdots t_n$, with $x_n := \phi(t_n, x) \in D_0$, define $t_{n+1} - t_n$ as:

$$\max\{\arg\max_{t \in (0, \tau]} h(\phi(t, x_n)) + \int_0^t \zeta(h(\phi(s, x_n))) ds\}. \tag{3.61}$$

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

$$x_{n+1} := \phi(t_{n+1}, x_0) = \phi(t_{n+1} - t_n, x_n)$$

satisfies the conditions in (3.14).

[Divergence of t_n]: It remains to show that $t_n \to \infty$, which we argue by contradiction. If, instead, the strictly increasing sequence of times was bounded, we

would have $t_n \uparrow t^*$. Note that $\phi(t^*, x_0)$ is well defined since the dynamical system (2.1) is forward complete. Also, by the continuity of $\phi(\cdot, x)$:

$$v_n := h(\phi(t_n, x_0)) + \int_0^{t_n} \zeta(h(\phi(s, x_0)))ds$$

$$\to v^* := h(\phi(t^*, x_0)) + \int_0^{t^*} \zeta(h(\phi(s, x_0)))ds.$$

Note that it follows from the first inequality of (3.14) that:

$$h(x_{n+1}) \geq h(x_n) + \delta_n$$

$$\geq h(x_n) - \int_0^{t_{n+1}-t_n} \zeta(h(\phi(s, x_n)))ds$$

$$= h(x_n) - \int_{t_n}^{t_{n+1}} \zeta(h(\phi(s, x_0)))ds$$

$$\geq h(x_n) - \int_0^{t_{n+1}} \zeta(h(\phi(s, x_0)))ds + \int_0^{t_n} \zeta(h(\phi(s, x_0)))ds$$

$$\implies v_{n+1} \geq v_n, \quad \forall n \in \mathbb{N}.$$

Therefore, $\{v_n\}$ is non-decreasing, and we further conclude that $v^* \geq v_n$ for all $n \in \mathbb{N}$. Now pick n such that $t_n \geq t^* - \tau$. This means that $s^* := t^* - t_n \in (0, \tau]$ is in the feasible set for the maximization in (3.61), which by definition gives as maximum v_{n+1}, achieved at $t_{n+1} - t_n$.

Now, since $v^* = h(s^*, x_n) + \int_0^{s^*} \zeta(h(\phi(s, x_n)))ds \geq v_{n+1}$, this means s^* also qualifies as a maximizer, and in fact $s^* = t^* - t_n > t_{n+1} - t_n$. This contradicts the definition of $t_{n+1} - t_n$ given in (3.61), because it would not be the largest maximizing time. Thus, the sequence must be divergent, establishing the claim. $\square$

Proof of Theorem 3.8

Proof. (i): Suppose the closed set $h_{\geq 0}$ is not τ-recurrent. Then, there exists an initial condition $x \in h_{=0}$ such that $\phi(t, x) \notin h_{\geq 0}$ for all $t \in (0, \tau]$. In this case, $h(\phi(t, x)) < 0$ and $\zeta(h(\phi(t, x))) < 0$ for all $t \in (0, \tau]$. Note that this contradicts condition (3.12), which requires its left-hand side to be nonnegative starting from $h(x) = 0$. Therefore, $h_{\geq 0}$ is τ-recurrent.

(ii): Starting from $h(x_n) < 0$, if $h(x_{n+1}) \geq 0$, a positive step size $h(x_{n+1}) - h(x_n) \geq -h(x_n) > 0$ can be ensured automatically. Conversely, if $h(x_{n+1}) < 0$, it follows from Lemma 3.2 that $h(x_{n+1}) \geq h(x_n) + \delta_n$ with $\delta_n > 0$. Thus, a positive step size $h(x_{n+1}) - h(x_n) \geq \delta_n > 0$ is also guaranteed.

(iii): Given $h(x_n) > 0$, the inequality (3.14) requires:

$$h(x_{n+1}) \geq h(x_n) + \delta_n \quad \text{and} \quad h(x_{n+1}) \geq 0.$$

Then, by rearranging terms, one can conclude $h(x_{n+1}) - h(x_n) \geq \max\{\delta_n, -h(x_n)\}$.

$\square$

Proof of Lemma 3.3

Proof. Given $x_0 = x \in D_0$, we build the time sequence $\{t_n\}_{n \in \mathbb{N}}$ satisfying (3.16) and (3.17) again by induction.

[Base case]: For the base case, we have $t_0 = 0$ and define t_1 as follows:

$$t_1 = \max\{\arg\max_{t \in (0,\tau]}\{e^{\beta t}[h(\phi(t, x_0))]_- + e^{\alpha t}[h(\phi(t, x_0))]_+\}\}$$

note that the second maximum exists by condition (3.15), and is no smaller than $h(x_0)$; if there are multiple maximizing times, t_1 is defined as the largest. By construction, $t_1 - t_0 \in (0, \tau]$, and the function h evaluated at $x_1 := \phi(t_1, x_0)$ satisfies:

$$e^{\beta t_1}[h(x_1)]_- + e^{\alpha t_1}[h(x_1)]_+ \geq h(x_0). \tag{3.62}$$

Note that whenever $h(x_0) \geq 0$, (3.62) requires $h(x_1) \geq 0$ and thus

$$h(x_1) \geq e^{-\alpha t_1}h(x_0).$$

In the case that $h(x_0) < 0$, we have:

$$h(x_1) \geq e^{-\beta t_1}h(x_0).$$

Therefore, (3.17) follows. Finally, $h(x_0) \geq -c$ and by (3.17) $h(x_1) > h(x_0)$, we have $x_1 = \phi(t_1, x_0) \in D_0$, which finishes the proof of the base case of the induction.

[Inductive step]: The inductive step construction proceeds in a similar manner: given $t_1 < t_2 < \cdots t_n$, with $x_k := \phi(t_k, x_0) \in D_0$, $0 \leq k \leq n$. Now, define $t_{n+1} - t_n$ as:

$$\max\{\arg\max_{t \in (0,\tau]}\{e^{\beta t}[h(\phi(t, x_n))]_- + e^{\alpha t}[h(\phi(t, x_n))]_+\}\} \tag{3.63}$$

Note that $t_{n+1} - t_n \in (0, \tau]$ as required. A similar proof to the base case then shows that

$$x_{n+1} := \phi(t_{n+1}, x_0) = \phi(t_{n+1} - t_n, x_n) \in D_0$$

and (3.17) is satisfied.

[Divergence of t_n]: It remains to show that $t_n \to \infty$, which we argue by contradiction. If, instead, the strictly increasing sequence of times was bounded, we would have $t_n \uparrow t^*$. Note that $\phi(t^*, x_0)$ is well defined since the dynamical system (2.1) is forward complete. Also, by the continuity of $\phi(\cdot, x)$:

$$v_n := e^{\beta t_n}[h(\phi(t_n, x_0))]_- + e^{\alpha t_n}[h(\phi(t_n, x_0))]_+$$
$$\to v^* := e^{\beta t^*}[h(\phi(t^*, x_0))]_- + e^{\alpha t^*}[h(\phi(t^*, x_0))]_+.$$

Note that it follows from (3.17) that:

$$e^{\beta t_{n+1}}[h(\phi(t_{n+1}, x_0))]_- + e^{\alpha t_{n+1}}[h(\phi(t_{n+1}, x_0))]_+$$
$$\geq e^{\beta t_n}[h(\phi(t_n, x_0))]_- + e^{\alpha t_n}[h(\phi(t_n, x_0))]_+$$
$$\implies v_{n+1} \geq v_n, \quad \forall n \in \mathbb{N}.$$

Therefore, $\{v_n\}$ is non-decreasing, and we further conclude that $v^* \geq v_n$ for all $n \in \mathbb{N}$. Now pick n such that $t_n \geq t^* - \tau$. This means that $s^* := t^* - t_n \in (0, \tau]$ is in the feasible set for the maximization in (3.63), which by definition gives as maximum v_{n+1}, achieved at $t_{n+1} - t_n$.

Now, since

$$v^* = e^{\beta s^*}[h(\phi(s^*, x_n))]_- + e^{\alpha s^*}[h(\phi(s^*, x_n))]_+ \geq v_{n+1},$$

92

this means s^* also qualifies as a maximizer, and in fact $s^* = t^* - t_n > t_{n+1} - t_n$. This contradicts the definition of $t_{n+1} - t_n$ given in (3.63), since it would not be the largest maximizing time. Thus, the sequence must be divergent, establishing the claim. □

Proof of Theorem 3.9

Proof. (i): Suppose the closed set $h_{\geq 0}$ is not τ-recurrent. There exists an initial condition $x \in h_{=0}$ such that $\phi(t, x) \notin h_{\geq 0}$ for all $t \in (0, \tau]$. In this case, $h(\phi(t, x)) < 0$ for all $t \in (0, \tau]$. Note that this contradicts condition (3.15), which requires $\max_{t \in (0, \tau]} e^{\beta t} h(\phi(t, x)) \geq 0$ starting from $h(x) = 0$. Therefore, $h_{\geq 0}$ is τ-recurrent.

(ii): Given $h(x_n) < 0$, then inequality (3.17) requires $h(x_{n+1}) \geq e^{-\beta \Delta t_n} h(x_n)$. Thus,

$$h(x_{n+1}) - h(x_n) = (e^{-\beta \Delta t_n} - 1) h(x_n) > 0.$$

Now, starting from $x \in h_{<0} \cap D_0$, we have

$$0 > h(x_n) \geq e^{-\beta(t_n - t_{n-1})} h(x_{n-1})$$
$$\geq e^{-\beta(t_n - t_{n+1} + t_{n+1} - t_{n-2})} h(x_{n-2}) \geq h(x) e^{-\beta t_n}.$$

whenever $x_n, ..., x_1 \in h_{<0} \cap D_0$. If there exists a $n' \in \{1, \ldots, n\}$ such that $h(x_{n'}) \geq 0$, (3.18) still follows since

$$h(x_n) \geq h(x_{n'}) e^{-\hat{\alpha}(t_n - t_{n'})} \geq 0 > h(x) e^{-\beta t_n}.$$

(iii): Whenever $h(x_n) > 0$, then inequality (3.17) requires $h(x_{n+1}) \geq e^{-\alpha \Delta t_n} h(x_n)$. Thus,

$$h(x_{n+1}) - h(x_n) = (e^{-\alpha \Delta t_n} - 1) h(x_n) < 0.$$

In the case that $x \in h_{\geq 0}$, we have:

$$h(x_n) \geq e^{-\alpha(t_n - t_{n-1})} h(x_{n-1})$$
$$\geq e^{-\alpha(t_n - t_{n+1} + t_{n+1} - t_{n-2})} h(x_{n-2})$$
$$\geq h(x) e^{-\alpha t_n} \geq 0,$$

i.e., (3.19) follows. □

Proof of Theorem 3.10

Proof. Let us first consider the case that $x \in h_{<0} \cap D_0$. In this case, (3.23) is automatically satisfied if there is a time t' with $h(\phi(t', x)) \geq 0$. Therefore, we focus on the case where $h(\phi(t, x)) < 0$ for all $t \in (0, \hat{\tau}]$. In this case, we have:

$$0 > \max_{t \in (0, \hat{\tau}]} e^{\hat{\beta}t} \mathrm{sd}(\phi(t, x), h_{\leq 0})$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{\hat{\beta}t} \frac{1}{\alpha_1} h(\phi(t, x)) \tag{3.64}$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{(\hat{\beta} - \beta)t} \frac{1}{\alpha_1} h(x) \tag{3.65}$$

$$\geq e^{(\hat{\beta} - \beta)\hat{\tau}} \frac{\alpha_2}{\alpha_1} \mathrm{sd}(x, h_{\leq 0}) \tag{3.66}$$

$$\geq \mathrm{sd}(x, h_{\leq 0}), \tag{3.67}$$

where (3.65) comes from the definition of the IEBF (3.10) and (3.64) and (3.66) are based on the sector containment assumption (3.22). Note that $\mathrm{sd}(x, h_{\leq 0}) < 0$, and thus (3.67) is true whenever $e^{(\hat{\beta} - \beta)\hat{\tau}} \frac{\alpha_2}{\alpha_1} \leq 1$, which is achieved by choosing $\hat{\tau} \geq \log(\alpha_2/\alpha_1)/(\beta - \hat{\beta})$.

Next, starting from $x \in h_{\geq 0}$, observe that $h(\phi(t, x)) \geq 0, \forall t \geq 0$, since Theorem 3.6 ensures $h_{\geq 0}$ is an invariant set. Therefore, we proceed similarly:

$$\max_{t \in (0, \hat{\tau}]} e^{\hat{\alpha}t} \mathrm{sd}(\phi(t, x), h_{\leq 0})$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{\hat{\alpha}t} \frac{1}{\alpha_2} h(\phi(t, x)) \tag{3.68}$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{(\hat{\alpha} - \alpha)t} \frac{1}{\alpha_2} h(x) \tag{3.69}$$

$$\geq e^{(\hat{\alpha} - \alpha)\hat{\tau}} \frac{\alpha_1}{\alpha_2} \mathrm{sd}(x, h_{\leq 0}) \tag{3.70}$$

$$\geq \mathrm{sd}(x, h_{\leq 0}) \geq 0, \tag{3.71}$$

where (3.69) comes from the definition of the IEBF (3.10), (3.68) and (3.70) come from the sector containment assumption (3.21). Now, since $\mathrm{sd}(x, h_{\leq 0}) \geq 0$, (3.71) is true whenever $e^{(\hat{\alpha} - \alpha)\hat{\tau}} \frac{\alpha_1}{\alpha_2} \geq 1$, which can be achieved by choosing $\hat{\tau} \geq \log(\alpha_1/\alpha_2)/(\alpha - \hat{\alpha})$. Consequently, by choosing $\hat{\tau}$ as specified in (3.24), we guarantee (3.23) in all possible scenarios. □

94

Proof of Theorem 3.11

Proof. We first note that since $h_{\geq 0} \subset S$ and $\partial S \cap h_{=0} = \emptyset$, the inequality $\bar{\delta} \geq \underline{\delta} > 0$ naturally holds. On top of this, definitions (3.26a) and (3.26b) together imples:

$$0 < \underline{\delta} \leq \mathrm{sd}(x, h_{\geq 0}) - \mathrm{sd}(x, S) \leq \bar{\delta}, \ \forall\, x \in D_0.$$

Consequently, we have

$$0 > -\underline{\delta} \geq \mathrm{sd}(x, h_{\leq 0}) - \hat{h}(x) \geq -\bar{\delta}, \ \forall\, x \in D_0, \tag{3.72}$$

since $\mathrm{sd}(x, h_{\leq 0}) = -\mathrm{sd}(x, h_{\geq 0})$ and $\hat{h}(x) = -\mathrm{sd}(x, S)$.

Since h is a sector contained IEBF, Theorem 3.10 ensures that the function $\mathrm{sd}(\cdot, h_{\leq 0})$ is a REBF satisfying (3.23) with any $\hat{\tau} \geq \tau^* := \max\{\frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha}-\alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta-\hat{\beta}}\}$. Theorem 3.9 part (ii-iii) further establishes a sequence $\{t_n\}_{n\in\mathbb{N}}$ with $t_0 = 0$ such that for each state $x_n := \phi(t_n, x)$, we have

$$\mathrm{sd}(x_n, h_{\leq 0}) \geq \begin{cases} \mathrm{sd}(x, h_{\leq 0})e^{-\hat{\beta}t_n} \text{ if } x \in h_{<0} \cap D_0 & \text{(3.73a)} \\ \mathrm{sd}(x, h_{\leq 0})e^{-\hat{\alpha}t_n} \text{ if } x \in h_{\geq 0} & \text{(3.73b)} \end{cases}$$

for all $n \in \mathbb{N}$.

Now, starting from any initial state $x \in \hat{h}_{<0} \cap D_0$, we have $h(x) < 0$ since $\hat{h}_{\geq 0} = S \supset h_{\geq 0}$. Therefore,

$$e^{\hat{\beta}t_n}\hat{h}(\phi(t_n, x)) \geq e^{\hat{\beta}t_n}(\mathrm{sd}(\phi(t_n, x), h_{\leq 0}) + \underline{\delta}) \tag{3.74}$$

$$\geq \mathrm{sd}(x, h_{\leq 0}) + e^{\hat{\beta}t_n}\underline{\delta} \tag{3.75}$$

$$\geq \hat{h}(x) - \bar{\delta} + e^{\hat{\beta}t_n}\underline{\delta} \tag{3.76}$$

$$\geq \hat{h}(x) \tag{3.77}$$

where (3.74) and (3.76) are based on (3.72), (3.75) comes from (3.73a), and (3.77) follows whenever $t_n \geq \log(\bar{\delta}/\underline{\delta})/\hat{\beta}$.

Next, starting from any $x \in \hat{h}_{\geq 0}$, we have:

$$e^{\hat{\alpha}t_n}\hat{h}(\phi(t_n, x)) \geq e^{\hat{\alpha}t_n}(\mathrm{sd}(\phi(t_n, x), h_{\leq 0}) + \underline{\delta}) \tag{3.78}$$

$$\geq \mathrm{sd}(x, h_{\leq 0}) + e^{\hat{\alpha}t_n}\underline{\delta} \tag{3.79}$$

$$\geq \hat{h}(x) - \bar{\delta} + e^{\hat{\alpha}t_n}\underline{\delta} \tag{3.80}$$

$$\geq \hat{h}(x) \geq 0, \quad \text{if } h(x) \geq 0, \tag{3.81}$$

or

$$e^{\hat{\alpha} t_n} \hat{h}(\phi(t_n, x)) \geq e^{\hat{\beta} t_n} \hat{h}(\phi(t_n, x)) \tag{3.82}$$

$$\geq \hat{h}(x) \geq 0, \quad \text{if } h(x) < 0. \tag{3.83}$$

where (3.78) and (3.80) are based on (3.72), (3.79) comes from (3.73b), and (3.81)

follows whenever $t_n \geq \log(\bar{\delta}/\underline{\delta})/\hat{\alpha}$. In the case that $h(x) < 0$, (3.82) and (3.83) are

based on (3.77) and the assumption that $\hat{\alpha} \geq \hat{\beta}$.

Finally, by combining these conditions, we can verify (3.25) whenever $t_n \geq \hat{\delta} :=$

$\log(\bar{\delta}/\underline{\delta})/\min\{\hat{\alpha}, \hat{\beta}\}$. Note that,

$$\lim_{n \to \infty} t_n = \infty \text{ and } t_{n+1} - t_n \in (0, \tau^*], \quad \forall n \in \mathbb{N}.$$

Therefore, by choosing $\hat{\tau} \geq \tau^* + \hat{\delta}$, there must be a $t_n \in [\hat{\delta}, \hat{\tau}]$ ensures (3.25), and the

result follows. □

Proof of Theorem 3.12

Proof. Suppose that $\mathcal{R}_\tau(S)$ is not invariant, there must exist a $y \in \mathcal{R}_\tau(S)$ and a $t_1 > 0$

such that $\phi(t_1, y) \notin \mathcal{R}_\tau(S)$. By the definition of the reachable set, there also exists

a $x \in S$ and a $t_2 \in [0, \tau)$, such that $\phi(t_2, x) = y$. Since $\mathcal{R}_\tau(S) \supseteq S$, we conclude

$\phi(t_1, y) = \phi(t_1 + t_2, x) \notin S$.

We then use t' to denote the last time the trajectory $\phi(t, x)$ stays within the closed

set S, i.e.,

$$x' := \phi(t', x) \in S \quad \text{and} \quad \phi(t, x) \notin S \ \forall t \in (t', t_1 + t_2].$$

Note that we must have $t_1 + t_2 - t' \leq \tau$ since S is τ-recurrent. This contradicts the

assumption that $\phi(t_1, y) \notin \mathcal{R}_\tau(S)$ since $\phi(t_1, y) = \phi(t_1 + t_2 - t', x') \in \mathcal{R}_\tau(S)$. □

Proof of Lemma 3.4

Proof. Starting from a point $x \in S$, we use d^* to denote the maximum signed

distance from the trajectory $\phi(t, x)$ to the set S within τ seconds, and t^* to denote

the time this maximum distance is achieved, i.e.,

$$d^* := \max_{t \in [0,\tau]} \mathrm{sd}(\phi(x,t), S), \; t^* := \arg\max_{t \in [0,\tau]} \mathrm{sd}(\phi(x,t), S).$$

Note that the assumption that the set S is τ-recurrent implies that any trajectory starting from S can leave S for at most τ seconds. Therefore, we only need to consider the maximum signed distance within τ seconds, which would be equivalent to the maximum distance of any trajectory starting from S over unbounded time.

If $d^* \leq 0$, it follows that $d^* \leq 0 \leq c$, and the result follows trivially. We then consider the case that $d^* > 0$. In this case, we further use t' to denote the last time before t^* such that the trajectory $\phi(x,t) \in S$, i.e.,

$$\phi(t', x) \in S \quad \text{and} \quad \phi(t, x) \notin S \; \forall t \in (t', t^*].$$

To show $\mathrm{sd}(\phi(t, x), S) \leq c_1$, we have:

$$\mathrm{sd}(\phi(t, x), S) \leq d^* = \mathrm{sd}(\phi(t^*, x), S)$$
$$\leq \|\phi(t^*, x) - \phi(t', x)\| = \int_{t'}^{t^*} \|f(\phi(s, x))\| ds$$
$$\leq F(\mathcal{R}_\tau(S))(t^* - t')$$
$$\leq F(\mathcal{R}_\tau(S))\tau.$$

To show $\mathrm{sd}(\phi(t, x), S) \leq c_2$, we have:

$$\mathrm{sd}(\phi(t, x), S) \leq d^* = \mathrm{sd}(\phi(t^*, x), S)$$
$$\leq \int_{t'}^{t^*} \|f(\phi(s, x))\| ds$$
$$\leq \int_{t'}^{t^*} \|f(\phi(s, x)) - F(\mathrm{P}_S(\phi(s, x)))\| + \|F(\mathrm{P}_S(\phi(s, x)))\| ds$$
$$\leq \int_{t'}^{t^*} \mathrm{sd}(\phi(s, x), S)L(\mathcal{R}_\tau(S)) + F(\partial S) ds$$
$$= \int_{t'}^{t^*} \mathrm{sd}(\phi(s, x), S)L(\mathcal{R}_\tau(S)) ds + F(\partial S)t$$

Then, by applying the Grönwall's inequality [3, Lemma 2.1] with $\lambda = F(\partial S)t, \mu = L(\partial S), y(t) = \mathrm{sd}(\phi(t, x), S)$, we have:

$$\mathrm{sd}(\phi(t, x), S) \leq F(\partial S)\tau e^{L(\mathcal{R}_\tau(S))\tau}, \forall x \in S.$$

A combination of these two conditions implies (3.29). $\qquad\square$

97

Proof of Lemma 3.5

Proof. Given an initial point $x_0 = \mathcal{X}_\tau$, we build the time sequence $\{t_n\}_{n \in \{1,...,N\}}$ satisfying (3.32) and (3.33) again by induction.

[Base case]: For the base case, we have $t_0 = 0$ and define t_1 as follows:

$$t_1 = \max\{\arg\max_{t \in (0,\tau]}\{e^{\beta t}[h(\phi(t, x_0))]_- + e^{\alpha t}[h(\phi(t, x_0))]_+\}\}$$

note that the second maximum exists by condition (3.31), and is no smaller than $h(x_0)$; if there are multiple maximizing times, t_1 is defined as the largest. By construction, $t_1 - t_0 \in (0, \tau]$, and the function h evaluated at $x_1 := \phi(t_1, x_0)$ satisfies:

$$e^{\beta t_1}[h(x_1)]_- + e^{\alpha t_1}[h(x_1)]_+ \geq h(x_0).$$

Note that whenever $h(x_0) \geq 0$, (3.62) requires $h(x_1) \geq 0$ and thus

$$h(x_1) \geq e^{-\alpha t_1}h(x_0).$$

In the case that $h(x_0) < 0$, we have:

$$h(x_1) \geq e^{-\beta t_1}h(x_0).$$

Therefore, (3.33) follows, which finishes the proof of the base case of the induction.

[Inductive step]: The inductive step construction proceeds in a similar manner: given $t_1 < t_2 < \cdots t_N$, with $x_k := \phi(t_k, x_0) \in \mathcal{X}_\tau$, $0 \leq k \leq N - 1$. Now, define $t_{n+1} - t_n$ as:

$$\max\{\arg\max_{t \in (0,\tau]}\{e^{\beta t}[h(\phi(t, x_n))]_- + e^{\alpha t}[h(\phi(t, x_n))]_+\}\} \tag{3.84}$$

Note that $t_{n+1} - t_n \in (0, \tau]$ as required. A similar proof to the base case then shows that

$$x_{n+1} := \phi(t_{n+1}, x_0) = \phi(t_{n+1} - t_n, x_n) \in \mathcal{X}_\tau$$

and (3.33) is satisfied.

[Divergence of t_n]: It remains to show that $t_N > T$, which we argue by contradiction. If, instead, the strictly increasing sequence of times was bounded before T, we would have $t_n \uparrow t^* \leq T$. Note that $\phi(t^*, x_0)$ is well defined since the dynamical system (2.1) is forward complete. Also, by the continuity of $\phi(\cdot, x)$:

$$
\begin{aligned}
v_n &:= e^{\beta t_n}[h(\phi(t_n, x_0))]_- + e^{\alpha t_n}[h(\phi(t_n, x_0))]_+ \\
&\to v^* := e^{\beta t^*}[h(\phi(t^*, x_0))]_- + e^{\alpha t^*}[h(\phi(t^*, x_0))]_+.
\end{aligned}
$$

Note that it follows from (3.33) that:

$$
\begin{aligned}
&e^{\beta t_{n+1}}[h(\phi(t_{n+1}, x_0))]_- + e^{\alpha t_{n+1}}[h(\phi(t_{n+1}, x_0))]_+ \\
&\geq e^{\beta t_n}[h(\phi(t_n, x_0))]_- + e^{\alpha t_n}[h(\phi(t_n, x_0))]_+ \\
&\implies v_{n+1} \geq v_n, \quad \forall n \in \{0, \ldots, N-1\}.
\end{aligned}
$$

Therefore, $\{v_n\}$ is non-decreasing, and we further conclude that $v^* \geq v_n$ for all $n \in \mathbb{N}$. Now pick n such that $t_n \geq t^* - \tau$. This means that $s^* := t^* - t_n \in (0, \tau]$ is in the feasible set for the maximization in (3.84), which by definition gives as maximum v_{n+1}, achieved at $t_{n+1} - t_n$.

Now, since

$$
v^* = e^{\beta s^*}[h(\phi(s^*, x_n))]_- + e^{\alpha s^*}[h(\phi(s^*, x_n))]_+ \geq v_{n+1},
$$

this means s^* also qualifies as a maximizer, and in fact $s^* = t^* - t_n > t_{n+1} - t_n$. This contradicts the definition of $t_{n+1} - t_n$ given in (3.84), since it would not be the largest maximizing time. Thus, the sequence must be divergent before achieving T, establishing the claim. $\qquad \square$

Proof of Theorem 3.16

Proof. *(i):* Suppose for the sake of contradiction that the closed set $h_{\geq 0} \cap \mathcal{X}_\tau$ is not locally τ-recurrent within $\mathcal{X}$. Then, there exists an initial condition $x \in h_{=0} \cap \mathcal{X}_\tau$ such that $\phi(t, x) \notin h_{\geq 0} \cap \mathcal{X}$ for all $t \in (0, \tau]$. In this case, $h(\phi(t, x)) < 0$, for all $t \in (0, \tau]$. Note that this contradicts condition (3.31), which requires that

$\max_{t \in (0,\tau]} e^{\beta t} h(\phi(t,x)) \geq 0$ starting from all $x \in h_{=0} \cap \mathcal{X}_\tau$. Therefore, $h_{\geq 0}$ is locally τ-recurrent w.r.t $\mathcal{X}$.

(ii): If $x_n \in h_{<0} \cap \mathcal{X}_\tau$ and $\phi(t,x) \in \mathcal{X}, \forall t \in [0, t_n + \tau]$, then inequality (3.33) requires that $h(x_{n+1}) \geq e^{-\beta \Delta t_n} h(x_n)$. Thus,

$$h(x_{n+1}) - h(x_n) \geq (e^{-\beta \Delta t_n} - 1) h(x_n) > 0.$$

Now, starting from some $x \in h_{<0} \cap \mathcal{X}_\tau$, we have

$$0 > h(x_n) \geq e^{-\beta(t_n - t_{n-1})} h(x_{n-1})$$
$$\geq e^{-\beta(t_n - t_{n+1} + t_{n+1} - t_{n-2})} h(x_{n-2}) \geq h(x) e^{-\beta t_n}.$$

for all $n \in \{0, \ldots, N\}$ whenever $x_n, \ldots, x_1 \in h_{<0} \cap \mathcal{X}$. If there exists an $n' \in \{1, \ldots, n\}$ such that $h(x_{n'}) \geq 0$, (3.34) still follows since

$$h(x_n) \geq h(x_{n'}) e^{-\alpha(t_n - t_{n'})} \geq 0 > h(x) e^{-\beta t_n}.$$

(iii): Whenever $x_n \in h_{>0} \cap \mathcal{X}_\tau$ and $\phi(t,x) \in \mathcal{X}, \forall t \in [0, t_n + \tau]$, then inequality (3.17) requires that $h(x_{n+1}) \geq e^{-\alpha \Delta t_n} h(x_n)$. Thus,

$$h(x_{n+1}) - h(x_n) \geq (e^{-\alpha \Delta t_n} - 1) h(x_n) < 0.$$

For any $x \in h_{\geq 0}$, we have:

$$h(x_n) \geq e^{-\alpha(t_n - t_{n-1})} h(x_{n-1})$$
$$\geq e^{-\alpha(t_n - t_{n+1} + t_{n+1} - t_{n-2})} h(x_{n-2})$$
$$\geq h(x) e^{-\alpha t_n} \geq 0,$$

for all $n \in \{0, \ldots, N\}$ whenever $x_n, \ldots, x_1 \in h_{<0} \cap \mathcal{X}$, i.e., (3.35) follows. $\square$

Proof of Theorem 3.17

Proof. Let us first consider the case that $x \in h_{<0} \cap \mathcal{X}_{\hat{\tau}}$. In this case, (3.39) is automatically satisfied if there is a $t' \leq \hat{\tau}$ with $h(\phi(t', x)) \geq 0$. Therefore, we focus on the case

where $h(\phi(t, x)) < 0$ for all $t \in (0, \hat{\tau}]$. In this case, we have:

$$0 > \max_{t \in (0, \hat{\tau}]} e^{\hat{\beta} t} \text{sd}|_{\mathcal{X}}(\phi(t, x), h_{\leq 0})$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{\hat{\beta} t} \frac{1}{\alpha_1} h(\phi(t, x)) \tag{3.85}$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{(\hat{\beta} - \beta) t} \frac{1}{\alpha_1} h(x) \tag{3.86}$$

$$\geq e^{(\hat{\beta} - \beta) \hat{\tau}} \frac{\alpha_2}{\alpha_1} \text{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \tag{3.87}$$

$$\geq \text{sd}|_{\mathcal{X}}(x, h_{\leq 0}), \tag{3.88}$$

where (3.86) comes from the definition of the LIEBF (3.30) and (3.85) and (3.87) are based on the sector containment assumption (3.38). Note that $\text{sd}|_{\mathcal{X}}(x, h_{\leq 0}) < 0$, and thus (3.88) is true if $e^{(\hat{\beta} - \beta) \hat{\tau}} \frac{\alpha_2}{\alpha_1} \leq 1$, which is achieved by choosing $\hat{\tau} \geq \log(\alpha_2 / \alpha_1) / (\beta - \hat{\beta})$.

Next, whenever $x \in h_{\geq 0} \cap \mathcal{X}_{\hat{\tau}}$, Theorem 3.15 implies $h(\phi(t, x)) \geq 0, \forall t \in [0, \hat{\tau}]$. Therefore, we proceed similarly:

$$\max_{t \in (0, \hat{\tau}]} e^{\hat{\alpha} t} \text{sd}|_{\mathcal{X}}(\phi(t, x), h_{\leq 0})$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{\hat{\alpha} t} \frac{1}{\alpha_2} h(\phi(t, x)) \tag{3.89}$$

$$\geq \max_{t \in (0, \hat{\tau}]} e^{(\hat{\alpha} - \alpha) t} \frac{1}{\alpha_2} h(x) \tag{3.90}$$

$$\geq e^{(\hat{\alpha} - \alpha) \hat{\tau}} \frac{\alpha_1}{\alpha_2} \text{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \tag{3.91}$$

$$\geq \text{sd}|_{\mathcal{X}}(x, h_{\leq 0}) \geq 0, \tag{3.92}$$

where (3.90) comes from the definition of the LIEBF (3.30), (3.89) and (3.91) come from the sector containment assumption (3.37). Now, since $\text{sd}(x, h_{\leq 0}) \geq 0$, (3.92) is true whenever $e^{(\hat{\alpha} - \alpha) \hat{\tau}} \frac{\alpha_1}{\alpha_2} \geq 1$, which can be achieved by choosing $\hat{\tau} \geq \log(\alpha_1 / \alpha_2) / (\alpha - \hat{\alpha})$. Consequently, by choosing $\hat{\tau}$ as specified in (3.40), we guarantee (3.39) in all possible scenarios. $\square$

Proof of Theorem 3.18

Proof. We first note that since $h_{\geq 0} \cap \mathcal{X} \subset S$ and $\partial S \cap h_{=0} = \emptyset$, the inequality $\bar{\delta} \geq \underline{\delta} > 0$ naturally holds. On top of this, definitions (3.42a) and (3.42b) together imples:

$$0 < \underline{\delta} \leq \mathrm{sd}|_{\mathcal{X}}(x, h_{\geq 0}) - \mathrm{sd}|_{\mathcal{X}}(x, S) \leq \bar{\delta}, \ \forall x \in \mathcal{X}.$$

Consequently, we have

$$0 > -\underline{\delta} \geq \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) - \hat{h}(x) \geq -\bar{\delta}, \ \forall x \in \mathcal{X}, \tag{3.93}$$

since $\mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) = -\mathrm{sd}|_{\mathcal{X}}(x, h_{\geq 0})$, $\hat{h}(x) = -\mathrm{sd}|_{\mathcal{X}}(x, S)$.

Since h is a locally sector contained LIEBF, Theorem 3.17 ensures that the function $\mathrm{sd}|_{\mathcal{X}}(\cdot, h_{\leq 0})$ is a REBF satisfying (3.39) with any $\hat{\tau} \geq \tau^* := \max\{\frac{\log(\alpha_2/\alpha_1)}{\hat{\alpha}-\alpha}, \frac{\log(\alpha_2/\alpha_1)}{\beta-\hat{\beta}}\}$. Theorem 3.16 part (ii-iii) further establishes a sequence $\{t_n\}_{n\in\{1,\ldots,N\}}$ with $t_0 = 0$ such that for each state $x_n := \phi(t_n, x)$, we have

$$\mathrm{sd}(x_n, h_{\leq 0}) \geq \begin{cases} \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0})e^{-\hat{\beta}t_n} \text{ if } x \in h_{<0} \cap \mathcal{X} & \text{(3.94a)} \\ \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0})e^{-\hat{\alpha}t_n} \text{ if } x \in h_{\geq 0} \cap \mathcal{X} & \text{(3.94b)} \end{cases}$$

for all $n \in \{0, \ldots, N\}$.

Now, starting from any initial state $x \in \hat{h}_{<0} \cap \mathcal{X}_{\hat{\tau}} \subseteq \mathcal{X}$, we have $h(x) < 0$ since $\hat{h}_{\geq 0} = S \supset h_{\geq 0} \cap \mathcal{X}$. Therefore,

$$e^{\hat{\beta}t_n}\hat{h}(\phi(t_n, x)) \geq e^{\hat{\beta}t_n}(\mathrm{sd}|_{\mathcal{X}}(\phi(t_n, x), h_{\leq 0}) + \underline{\delta}) \tag{3.95}$$

$$\geq \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) + e^{\hat{\beta}t_n}\underline{\delta} \tag{3.96}$$

$$\geq \hat{h}(x) - \bar{\delta} + e^{\hat{\beta}t_n}\underline{\delta} \tag{3.97}$$

$$\geq \hat{h}(x) \tag{3.98}$$

where (3.95) and (3.97) are based on (3.93), (3.96) comes from (3.94a), and (3.98) follows whenever $t_n \geq \log(\bar{\delta}/\underline{\delta})/\hat{\beta}$.

Next, starting from any $x \in \hat{h}_{\geq 0} \cap \mathcal{X}_{\hat{\tau}}$, we have:

$$e^{\hat{\alpha}t_n}\hat{h}(\phi(t_n, x)) \geq e^{\hat{\alpha}t_n}(\mathrm{sd}|_{\mathcal{X}}(\phi(t_n, x), h_{\leq 0}) + \underline{\delta}) \tag{3.99}$$

$$\geq \mathrm{sd}|_{\mathcal{X}}(x, h_{\leq 0}) + e^{\hat{\alpha}t_n}\underline{\delta}$$

$$\geq \hat{h}(x) - \bar{\delta} + e^{\hat{\alpha}t_n}\underline{\delta} \tag{3.100}$$

$$\geq \hat{h}(x) \geq 0, \quad \text{if } h(x) \geq 0, \tag{3.101}$$

or

$$e^{\hat{\alpha} t_n} \hat{h}(\phi(t_n, x)) \geq e^{\hat{\beta} t_n} \hat{h}(\phi(t_n, x)) \tag{3.102}$$

$$\geq \hat{h}(x) \geq 0, \quad \text{if } h(x) < 0. \tag{3.103}$$

where (3.99) and (3.100) are based on (3.93), (3.79) comes from (3.94b), and (3.101) follows whenever $t_n \geq \log(\bar{\delta}/\underline{\delta})/\hat{\alpha}$. In the case that $h(x) < 0$, (3.102) and (3.103) are based on (3.98) and the assumption that $\hat{\alpha} \geq \hat{\beta}$.

Finally, by combining these conditions, we can verify (3.41) whenever $t_n \geq \hat{\delta} := \log(\bar{\delta}/\underline{\delta})/\min\{\hat{\alpha}, \hat{\beta}\}$. Note that, starting from any $x \in \mathcal{X}_{\hat{\tau}}$, we have $t_N \geq \hat{\tau} - \tau^*$ and:

$$t_{n+1} - t_n \in (0, \tau^*], \quad \forall n \in \{0, \dots, N-1\}.$$

Therefore, by choosing $\hat{\tau} \geq \tau^* + \hat{\delta}$, there must be a $t_n \in [\hat{\delta}, \hat{\tau}]$ ensures (3.41), and the result follows. □

Proof of Lemma 3.6

Proof. We will leverage the fact that $\|x - y\| \leq r$ to show equation (3.43) in the following four cases:

(1): Whenever $x, y \notin S$:

$$
\begin{aligned}
\text{sd}(x, S) = \|x - P_{\partial S}(x)\| &\leq \|x - P_{\partial S}(y)\| \\
&\leq \|x - y + y - P_{\partial S}(y)\| \\
&\leq \|x - y\| + \|y - P_{\partial S}(y)\| \\
&= \|x - y\| + \text{sd}(y, S) \\
&\leq r + \text{sd}(y, S) \\
\implies \text{sd}(y, S) &\geq \text{sd}(x, S) - r
\end{aligned}
$$

(2): Whenever $x \in S$ and $y \notin S$:

$$
\begin{aligned}
\text{sd}(x, S) = -\|x - P_{\partial S}(x)\| &\leq \|x - P_{\partial S}(x)\| \\
&\leq \|x - P_{\partial S}(y)\| \\
&\leq \|x - y + y - P_{\partial S}(y)\| \\
&\leq \|x - y\| + \|y - P_{\partial S}(y)\| \\
&= \|x - y\| + \text{sd}(y, S) \\
&\leq r + \text{sd}(y, S) \\
\implies \text{sd}(y, S) &\geq \text{sd}(x, S) - r
\end{aligned}
$$

103

(3): Whenever $x, y \in S$:

$$
\begin{aligned}
-\mathrm{sd}(y, S) = \|y - \mathrm{P}_{\partial S}(y)\| &\leq \|y - \mathrm{P}_{\partial S}(x)\| \\
&\leq \|y - x + x - \mathrm{P}_{\partial S}(x)\| \\
&\leq \|y - x\| + \|x - \mathrm{P}_{\partial S}(x)\| \\
&= \|y - x\| - \mathrm{sd}(x, S) \\
&\leq r - \mathrm{sd}(x, S) \\
\implies \mathrm{sd}(y, S) &\geq \mathrm{sd}(x, S) - r
\end{aligned}
$$

(4): Whenever $x \notin S$ and $y \in S$:

In this case, we use p to denote the point where ∂S intersects with the line between x and y. If there is more than one intersection, let p be a random one. Note that such a point p must exist since $x \notin S$ but $y \in S$. Then:

$$
\begin{aligned}
r \geq \|x - y\| = \|x - p\| + \|p - y\| &\geq \|y - \mathrm{P}_{\partial S}(y)\| + \|x - \mathrm{P}_{\partial S}(x)\| \\
&= -\mathrm{sd}(y, S) + \mathrm{sd}(x, S) \\
\implies \mathrm{sd}(y, S) &\geq \mathrm{sd}(x, S) - r
\end{aligned}
$$

A combination of these four conditions implies (3.43). Result follows. $\qquad \square$

Proof of Theorem 3.20

Proof. (i): If the initial states satisfies the condition in (3.45), then it follows from Corollary 3.1 that:

$$
\mathrm{sd}(\phi(t, y), \mathcal{X}_u) \geq \mathrm{sd}(\phi(t, x), \mathcal{X}_u) - re^{Lt} > 0,
$$

for all $t \in [0, \tau]$ and for all $y \in \mathcal{B}_r(x)$. Therefore, $\mathcal{B}_r(x) \cap \mathcal{R}_\tau^{-1}(\mathcal{X}_u) = \emptyset$.

(ii): If the initial states satisfies the condition in (3.46) instead, then let us use t^* to denote the time at which $\mathrm{sd}(\phi(t^*, x), \mathcal{X}_u) < -re^{Lt^*}$. Again, it follows from Corollary 3.1 that:

$$
\mathrm{sd}(\phi(t^*, y), \mathcal{X}_u) \leq \mathrm{sd}(\phi(t^*, x), \mathcal{X}_u) + re^{Lt^*} < 0,
$$

for all $y \in \mathcal{B}_r(x)$. Consequently, $\mathcal{B}_r(x) \subseteq \mathcal{R}_\tau^{-1}(\mathcal{X}_u)$. $\qquad \square$

Proof of Lemma 3.7

Proof. Again, we leverage the fact that $\|x - y\| \leq r$ to show equation (3.49) in the following four cases:

(1): Whenever $x, y \notin S$:

$$
\begin{aligned}
\mathrm{sd}|_{\mathcal{X}}(x, S) = \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| &\leq \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&\leq \|x - y + y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&\leq \|x - y\| + \|y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&= \|x - y\| + \mathrm{sd}|_{\mathcal{X}}(y, S) \\
&\leq r + \mathrm{sd}|_{\mathcal{X}}(y, S) \\
\Longrightarrow \ \mathrm{sd}|_{\mathcal{X}}(y, S) &\geq \mathrm{sd}|_{\mathcal{X}}(x, S) - r
\end{aligned}
$$

(2): Whenever $x \in S$ and $y \notin S$:

$$
\begin{aligned}
\mathrm{sd}(x, S) = -\|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| &\leq \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| \\
&\leq \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&\leq \|x - y + y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&\leq \|x - y\| + \|y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| \\
&= \|x - y\| + \mathrm{sd}(y, S) \\
&\leq r + \mathrm{sd}(y, S) \\
\Longrightarrow \ \mathrm{sd}(y, S) &\geq \mathrm{sd}(x, S) - r
\end{aligned}
$$

(3): Whenever $x, y \in S$:

$$
\begin{aligned}
-\mathrm{sd}|_{\mathcal{X}}(y, S) = \|y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| &\leq \|y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| \\
&\leq \|y - x + x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| \\
&\leq \|y - x\| + \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| \\
&= \|y - x\| - \mathrm{sd}|_{\mathcal{X}}(x, S) \\
&\leq r - \mathrm{sd}|_{\mathcal{X}}(x, S) \\
\Longrightarrow \ \mathrm{sd}|_{\mathcal{X}}(y, S) &\geq \mathrm{sd}|_{\mathcal{X}}(x, S) - r
\end{aligned}
$$

(4): Whenever $x \notin S$ and $y \in S$:

In this case, we use p to denote the point where $\partial S \cap \mathrm{int}\mathcal{X}$ intersects with the line between x and y. If there is more than one intersection, let p be a random one. Since $x, y \in \mathcal{X}$ and $\mathcal{X}$ is convex, the entire line between x and y is a subset of $\mathcal{X}$.

Therefore, such a point p must exist given $x \notin S$ but $y \in S$. Then:

$$
\begin{aligned}
r \geq \|x - y\| &= \|x - p\| + \|p - y\| \\
&\geq \|y - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(y)\| + \|x - \mathrm{P}_{\partial S \cap \mathrm{int}\mathcal{X}}(x)\| \\
&= -\mathrm{sd}|_{\mathcal{X}}(y, S) + \mathrm{sd}|_{\mathcal{X}}(x, S) \\
\implies \mathrm{sd}|_{\mathcal{X}}(y, S) &\geq \mathrm{sd}|_{\mathcal{X}}(x, S) - r
\end{aligned}
$$

A combination of these four conditions implies (3.49).

For (3.50), we first note that $\phi(t, x), \phi(t, y) \in \mathcal{X}$, $\forall t \in [0, \tau]$ since $x, y \in \mathcal{X}_\tau$. Then, under the Lipschiz assumption, i.e., (3.44), (3.50) follows from (3.49). $\qquad\square$

Proof of Theorem 3.22

Proof. (i): For the right-hand side of (3.51), we first conclude from Lemma 3.7 that:

$$
h(x) + r = -\mathrm{sd}|_{\mathcal{X}}(x, S) + r \geq -\mathrm{sd}|_{\mathcal{X}}(y, S) = h(y). \tag{3.104}
$$

We then let t^* to be the time when the left-hand side of (3.51) is maximized, i.e.,

$$
t^* = \arg\max_{t \in (0, \tau]} e^{\alpha t}[\hat{h}_r^-(t, x)]_+ + e^{\beta t}[\hat{h}_r^-(t, x)]_-.
$$

At this maximization time $t^* \in (0, \tau]$, we have:

$$
\begin{aligned}
e^{\alpha t^*}&[h(\phi(t^*, y))]_+ + e^{\beta t^*}[h(\phi(t^*, y))]_- \\
&\geq e^{\alpha t^*}[\hat{h}_r^-(t^*, x)]_+ + e^{\beta t^*}[\hat{h}_r^-(t^*, x)]_- && (3.105) \\
&\geq h(x) + r && (3.106) \\
&\geq h(y) && (3.107)
\end{aligned}
$$

where (3.105) comes from Lemma 3.7, (3.106) comes from (3.51), and (3.107) comes from (3.104).

(ii): For the right-hand side of (3.52), we again conclude from Lemma 3.7 that:

$$
h(x) - r = -\mathrm{sd}|_{\mathcal{X}}(x, S) - r \leq -\mathrm{sd}|_{\mathcal{X}}(y, S) = h(y). \tag{3.108}
$$

We then let t^* to be the time when the left-hand side of (3.52) is maximized, i.e.,

$$
t^* = \arg\max_{t \in (0, \tau]} e^{\alpha t}[\hat{h}_r^+(t, x)]_+ + e^{\beta t}[\hat{h}_r^+(t, x)]_-.
$$

Again, at the maximization time $t^* \in (0, \tau]$, we have:

$$e^{\alpha t^*}[h(\phi(t^*, y))]_+ + e^{\beta t^*}[h(\phi(t^*, y))]_-$$

$$\leq e^{\alpha t^*}[\hat{h}_r^+(t^*, x)]_+ + e^{\beta t^*}[\hat{h}_r^+(t^*, x)]_- \tag{3.109}$$

$$< h(x) - r \tag{3.110}$$

$$\leq h(y) \tag{3.111}$$

where (3.109) follows from Lemma 3.7, (3.110) follows from (3.52), and (3.111) follows from (3.108). $\qquad\square$

3.5 Conclusion

In this chapter, we introduce a novel methodology designed to characterize safety within a model-free context. This approach uniquely relies on sampled system trajectories, offering a practical way to assess safety without the need for detailed models of the system dynamics.

Our first step is to relax the notion of set invariance in the context of characterizing the safe state space region. To this end, we systematically relax the classic differential barrier conditions into integral conditions and further into recurrent conditions. We also thoroughly explore the interconnections between these conditions, and we establish sufficient conditions under which a τ-recurrent set, induced by the recurrent conditions, can be utilized to confirm safety.

We then detail a data-driven algorithm built on the theoretical insights previously discussed, aimed at effectively characterizing safety within a designated region, $\mathcal{X}$. This algorithm requires only a limited number of finite-length trajectory samples, and it capitalizes on the capabilities of highly parallelizable processing units, facilitating the simultaneous processing of multiple trajectory samples.

The flexibility of this method is further highlighted by its ability to incrementally search for safety levels through the analysis of safety-boundary-approaching rates of trajectories originating from $\mathcal{X}$.

We conclude this chapter with a numerical example that demonstrates the practical effectiveness of our approach.

Chapter 4

Conclusions

This thesis explores the development of rigorous methodologies for certifying the safety of high-dimensional dynamical systems in a model-free context, particularly focusing on learning regions of attraction and safe state space regions without full system models. The key innovation is applying the concept of recurrence to relax the stringent constraints typically imposed by invariance in safety certification. By leveraging recurrence, this work demonstrates that recurrent sets can be effectively used to characterize safety with enhanced efficiency and accuracy.

Theoretically, the thesis establishes necessary and sufficient conditions for utilizing recurrent approaches to characterize safety. This theoretical framework lays the groundwork for a deeper understanding of how recurrence can serve as a reliable proxy for invariance, thus simplifying the complex task of safety verification in dynamic systems.

Building on the theoretical insights, the thesis introduces practical, data-driven algorithms that leverage a finite number of finite-length sampled trajectories to determine safe regions within a dynamical system. This algorithm is optimized for computational efficiency and can be effectively implemented using parallel processing units, making it applicable in real-world scenarios where safety is paramount.

The methodology is validated through various numerical examples that illus-

trate the practicality and effectiveness of the proposed approaches. By utilizing sampled system trajectories, the research provides a robust framework for the safety verification of dynamic systems.

Bibliography

[1] R. Siegelmann, Y. Shen, F. Paganini, and E. Mallada, "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*, 2023, pp. 6665–6672.

[2] H. Sibai and E. Mallada, "Recurrence of nonlinear control systems: Entropy and bit rates," in *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control*, 2024, pp. 1–9.

[3] H. K. Khalil, "Nonlinear systems; 3rd ed." 2002.

[4] S. Willard, *General Topology*, ser. Addison-Wesley series in mathematics. Dover Publications, 2004.

[5] Y. Li, S. Das, and N. Li, "Online optimal control with affine constraints," 2020.

[6] A. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," pp. 70–76, 2017.

[7] I. Mitchell and C. Tomlin, "Overapproximating reachable sets by hamilton-jacobi projections," *Journal of Scientific Computing*, vol. 19, 11 2002.

[8] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference*, 2019, pp. 3420–3431.

[9] P. Giesl and S. Hafstein, "Review on computational methods for lyapunov functions," *Discrete and Continuous Dynamical Systems - B*, vol. 20, no. 8, pp. 2291–2331, 2015.

[10] A. Vannelli and M. Vidyasagar, "Maximal lyapunov functions and domains of attraction for autonomous nonlinear systems," *Automatica*, vol. 21, no. 1, pp. 69–80, 1985.

[11] M. Hassan and C.Storey, "Numerical determination of domains of attraction for electrical power systems using the method of zubov," *International Journal of Control*, vol. 34, no. 2, pp. 371–381, 1981.

[12] P. Julian, J. Guivant, and A. Desages, "A parametrization of piecewise linear lyapunov functions via linear programming," *International Journal of Control*, vol. 72, no. 7-8, pp. 702–715, 1999.

[13] R. Goebel, A. Teel, T. Hu, and Z. Lin, "Conjugate convex lyapunov functions for dual linear differential inclusions," *IEEE Transactions on Automatic Control*, vol. 51, no. 4, pp. 661–666, 2006.

[14] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado, "Learning lyapunov functions for hybrid systems," 2020.

[15] B. Xue, N. Zhan, and Y. Li, "Robust regions of attraction generation for state-constrained perturbed discrete-time polynomial systems," 2020.

[16] R. Ambrosino and E. Garone, "Robust stability of linear uncertain systems through piecewise quadratic lyapunov functions defined over conical partitions," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012.

[17] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado, "Learning region of attraction for nonlinear systems," 2021.

[18] U. Topcu, A. K. Packard, P. Seiler, and G. J. Balas, "Robust region-of-attraction estimation," *IEEE Transactions on Automatic Control*, vol. 55, no. 1, pp. 137–142, 2009.

[19] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, "Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 4661–4666.

[20] E. Najafi, R. Babuska, and G. Lopes, "A fast sampling method for estimating the domain of attraction," *Nonlinear Dynamics*, vol. 86, pp. 823–834, 2016.

[21] S. M. Richards, F. Berkenkamp, and A. Krause, "The lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems," 2018.

[22] B. K. Colbert and M. M. Peet, "Estimating the region of attraction using stable trajectory measurements," 2018.

[23] C. A. Montenegro G., S. Leudo, and R. Sanfelice, "A data-driven approach for certifying asymptotic stability and cost evaluation for hybrid systems," in *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '24. New York, NY, USA: Association for Computing Machinery, 2024.

[24] A. Abate, D. Ahmed, M. Giacobbe, and A. Peruffo, "Formal synthesis of lyapunov neural networks," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 773–778, 2021.

[25] R. Genesio, M. Tartaglia, and A. Vicino, "On the estimation of asymptotic stability regions: State of the art and new proposals," *IEEE Transactions on Automatic Control*, vol. 30, no. 8, pp. 747–755, 1985.

[26] H.-D. Chiang, M. Hirsch, and F. Wu, "Stability regions of nonlinear autonomous dynamical systems," *IEEE Transactions on Automatic Control*, vol. 33, no. 1, pp. 16–27, 1988.

[27] X.-S. Wang, J. D. Turner, and B. P. Mann, "A model-free sampling method for estimating basins of attraction using hybrid active learning (hal)," 2020.

[28] E. R. Vieira, A. Sivaramakrishnan, S. Tangirala, E. Granados, K. Mischaikow, and K. E. Bekris, "Morals: Analysis of high-dimensional robot controllers via topological tools in a latent space," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 27–33.

[29] M. NAGUMO, "Uber die lage der integralkurven gewouml;hnlicher differentialgleichungen," *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, vol. 24, pp. 551–559, 1942.

[30] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.

[31] H. Kong, F. He, X. Song, W. N. N. Hung, and M. Gu, "Exponential-condition-based barrier certificate generation for safety verification of hybrid systems," 2013.

[32] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," *IFAC Proceedings Volumes*, vol. 40, no. 12, pp. 462–467, 2007, 7th IFAC Symposium on Nonlinear Control Systems.

[33] M. Z. Romdlony and B. Jayawardhana, "Uniting control lyapunov and control barrier functions," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 2293–2298.

[34] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[35] E. D. Sontag, "A 'universal' construction of artstein's theorem on nonlinear stabilization," *Systems and Control Letters*, vol. 13, no. 2, pp. 117–123, 1989.

[36] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, p. 54–61, 2015.

[37] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," in *2020 59th IEEE Conference on Decision and Control*, 2020.

[38] M. Srinivasan, A. Dabholkar, S. Coogan, and P. Vela, "Synthesis of control barrier functions using a supervised machine learning approach," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 03 2020.

[39] O. Botchkarev and S. Tripakis, "Verification of hybrid systems with linear differential inclusions using ellipsoidal approximations," in *Hybrid Systems: Computation and Control,Lecture Notes in Computer Science.* Springer Berlin, Heidelberg, Dec 1999, p. 73–88.

[40] A. Kurzhanski and P. Varaiya, "On ellipsoidal techniques for reachability analysis. part i: External approximations," *Optimization Methods & Software - OPTIM METHOD SOFTW*, vol. 17, pp. 177–206, 04 2002.

[41] A. Bemporad, F. D. Torrisi, and M. Morari, "Optimization-based verification and stability characterization of piecewise affine and hybrid systems," in

Hybrid Systems: Computation and Control,Lecture Notes in Computer Science, Dec 1999, p. 45–58.

[42] A. Chutinan and B. Krogh, "Verification of polyhedral-invariant hybrid automata using polygonal flow pipe approximations," *Int. Workshop Hybrid Syst.: Comput. Control*, 11 2000.

[43] Y. Yazarel and P. Pappas, "Geometric programming relaxations for linear system reachability," *American Control Conference,American Control Conference*, Dec 2003.

[44] A. Chutinan and B. Krogh, "Computational techniques for hybrid system verification," *IEEE Transactions on Automatic Control*, p. 64–75, Dec 2002.

[45] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations," *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, 1988.

[46] C. Tomlin, I. Mitchell, A. Bayen, and M. Oishi, "Computational techniques for the verification of hybrid systems," *Proceedings of the IEEE*, p. 986–1001, Jun 2003.

[47] I. Mitchell and C. J. Tomlin, "Level set methods for computation in hybrid systems," in *Hybrid Systems: Computation and Control*, N. Lynch and B. H. Krogh, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 310–323.

[48] I. Mitchell, A. M. Bayen, and C. J. Tomlin, "Validating a hamilton-jacobi approximation to hybrid system reachable sets," in *Hybrid Systems: Computation and Control*, M. D. Di Benedetto and A. Sangiovanni-Vincentelli, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 418–432.

[49] B. Riečan, "Variation on a poincaré theorem," *Sets and Systems*, vol. 232, pp. 39–45, 2013, set Theory and Applications Selected papers from FSTA 2012.

[50] M. KAC, "On the notion of recurrence in discrete stochastic processes," *Bull. Amer. Math. Soc.*, vol. 53, pp. 1002–1010, 10 1946.

[51] I. Antoniou, "Caratheodory and the foundations of thermodynamics and statistical physics," *Foundations of Physics*, vol. 32, pp. 627–641, 04 2002.

[52] R. Easton, *Chain transitivity and the domain of influence of an invariant set.* Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 95–102.

[53] C. C. Conley, *Isolated invariant sets and the Morse index / Charles Conley.*, ser. Regional conference series in mathematics ; no. 38. Providence: Published for the Conference Board of the Mathematical Sciences by the American Mathematical Society, 1978.

[54] P. Pageault, "Lyapunov functions : a weak KAM approach," Theses, Ecole normale supérieure de lyon - ENS LYON, Nov. 2011.

[55] M. Soyer, S. Olaru, and Z. Fang, "From constraint satisfactions to periodic positive invariance for discrete-time systems," in *2020 59th IEEE Conference on Decision and Control*, 2020, pp. 4547–4552.

[56] R. Comelli, A. H. González, A. Ferramosca, S. Olaru, M. M. Seron, and E. Kofman, "Simplified design of practically stable mpc schemes," *Systems & Control Letters*, vol. 180, p. 105626, 2023.

[57] Y. Shen, M. Bichuch, and E. Mallada, "On the value of energy storage in generation cost reduction," in *19th IEEE European Control Conference (ECC)*, 5 2020, pp. 1526–1532.

[58] M. Almassalkhi, S. Brahma, N. Nazir, H. Ossareh, P. Racherla, S. Kundu, S. P. Nandanoori, T. Ramachandran, A. Singhal, D. Gayme, C. Ji, E. Mallada, Y. Shen, P. You, and D. Anand, "Hierarchical, grid-aware, and economically optimal coordination of distributed energy resources in realistic distribution systems," *Energies*, vol. 13, no. 23, pp. 1–35, 12 2020.

[59] E. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, ser. Texts in Applied Mathematics.   Springer New York, 2013.

[60] J. R. Munkres, *Topology / James R. Munkres.*, 2nd ed.   Upper Saddle River, NJ: Prentice Hall, Inc., 2000.

[61] R. D. Driver, "Methods of am lyapunov and their application (vi zubov)," *SIAM Review*, vol. 7, no. 4, p. 570, 1965.

[62] Y. Shen, M. Bichuch, and E. Mallada, "Model-free learning of regions of attraction via recurrent sets," in *2022 IEEE 61st Conference on Decision and Control (CDC)*.   IEEE, 2022, pp. 4714–4719.

[63] M. Lienen and S. Günnemann, "torchode: A parallel ODE solver for pytorch," in *The Symbiosis of Deep Learning and Differential Equations II, NeurIPS*, 2022.

[64] F. Bullo, *Contraction Theory for Dynamical Systems*.   Kindle Direct Publishing, 2023.

[65] M. Anghel, F. Milano, and A. Papachristodoulou, "Algorithmic construction of lyapunov functions for power system stability analysis," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, no. 9, pp. 2533–2546, 2013.

[66] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.

[67] S. Prajna and A. Rantzer, "Convex programs for temporal verification of nonlinear dynamical systems," *SIAM Journal on Control and Optimization*, p. 999–1021, Dec 2006.

[68] T. M. Apostol, *One-variable calculus, with an introduction to linear algebra*. Blaisdell Waltham, Toronto, London, 1967, vol. 1.

[69] R. Kannan and C. Krueger, *Advanced Analysis: On the Real Line*, ser. Universitext - Springer-Verlag. Springer New York, 1996.