

Model-free Learning of Regions of Attraction via Recurrent Sets

Yue Shen, Maxim Bichuch, and Enrique Mallada

Abstract—We consider the problem of learning an inner approximation of the region of attraction (ROA) of an asymptotically stable equilibrium point without an explicit model of the dynamics. Rather than leveraging approximate models with bounded uncertainty to find a (robust) invariant set contained in the ROA, we propose to learn sets that satisfy a more relaxed notion of containment known as recurrence. We define a set to be τ -recurrent (resp. k -recurrent) if every trajectory that starts within the set, returns to it after at most τ seconds (resp. k steps). We show that under mild assumptions a τ -recurrent set containing a stable equilibrium must be a subset of its ROA. We then leverage this property to develop algorithms that compute inner approximations of the ROA using counter-examples of recurrence that are obtained by sampling finite-length trajectories. Our algorithms process samples sequentially, which allows them to continue being executed even after an initial offline training stage. We further provide an upper bound on the number of counter-examples used by the algorithm, and almost sure convergence guarantees.

I. INTRODUCTION

The problem of estimating the region of attraction (ROA) of an asymptotically stable equilibrium point has a long-standing history in nonlinear control and dynamical systems theory [1]. 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 [2]. From a practical standpoint, having a representation of such region allows to test the limits of controller designs, which are usually based on (possibly linear) approximations of nonlinear systems [3], and provides a mechanism for safety verification of certain operating conditions [4] [5]. Unfortunately, it is known that finding an analytic form of the region of attraction is difficult and in general impossible [1, p. 122]. As a result, most effort in characterizing the ROA focus on finding inner approximations by means of invariant sets.

1) *Related Work*: 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, depending on whether accurate information about the dynamic model is present or not. Notably, at their core, almost all of the methods rely on finding an invariant set of the system. We briefly review such methods next.

Exact Models: When an exact description of the dynamics is available, it is possible to use this information via two complementary methodologies. Lyapunov methods utilize

the fact that Lyapunov functions are certificates of asymptotic stability and build inner approximations using its sublevel sets. Methods for finding such Lyapunov functions are surveyed in, e.g., [6]. In particular, [7] and [8] construct Lyapunov functions that are solutions of Zubov’s equation, and [9] searches for piece-wise linear Lyapunov functions that are found via linear programming. Similarly, piece-wise quadratic parameterizations of Lyapunov functions using LMI-based methods are considered in [10]. Finally, recent work [11] leverages the universal approximation property of neural networks to estimate the ROA of general nonlinear dynamical systems.

Alternatively, non-Lyapunov methods focus directly on the properties of the ROA. For example, trajectory reversing methods [12] [13] derive the boundary of ROA directly from the stable manifold of the equilibria on the boundary, and the reachable set method [14] generates a grid of sample points and classifies each of them by solving an optimal control problem.

Inexact Models: In the presence of uncertainty, robust ROA approximation methods [15]–[18] generalize Lyapunov approaches by finding a common Lyapunov function across the entire uncertainty set. Alternatively, learning-based methods utilize experimental data to estimate the region of attraction. When a Lyapunov function is provided, experimental data expand the Lyapunov function level set through, e.g., Gaussian processes [19], or a simple sampling approach [20]. To address the problem of simultaneously learning the Lyapunov function and the level set, [21] parameterizes the Lyapunov function as a neural network and iteratively trains it by sampling points that are outside of the current Lyapunov level set but come back in within T steps.

Notably, learning methods play a crucial role in model-free settings. In particular, similar to the Lyapunov methods, [22] uses trajectory data to fit values of a Lyapunov function by leveraging converse Lyapunov results. Perhaps most relevant to our paper is [23], which establishes a non-Lyapunov approach that determines the boundary of ROA directly from a support vector machine, trained from experimental data that is sampled via hybrid active learning techniques.

2) *Contributions*: In this paper, we provide a novel approach for learning inner approximations of the region of attraction of an asymptotically stable equilibrium point from sampled finite-length trajectories. We refer to such a method as “model-free” since it does not require an explicit description of the system but only requires a process that generates the sample trajectories.

Rather than focusing on learning invariant sets that require trajectories to always lie within the set, we propose to learn

Y. Shen and E. Mallada are with the Department of Electrical and Computer Engineering and M. Bichuch is with the Department of Applied Mathematics and Statistics, Johns Hopkins University, 3400 N Charles St., Baltimore, MD 21218, USA {yshen50, mbichuch, mallada}@jhu.edu.

sets that satisfy a more flexible notion of invariance. The contributions of this work are manifold:

- We propose the notion of recurrence as an alternative property that can be used to guarantee a set to be contained in the region of attraction.
- We show that under mild conditions, a compact set containing an asymptotically stable equilibrium point is a subset of the region of attraction if and only if it is recurrent.
- We leverage this property to develop several algorithms that can learn inner approximations of the region of attraction using counter-examples of recurrence that are based on finite-length trajectory samples.
- We further provide guarantees on the worst-case number of counter-examples required to compute a recurrent set.

3) *Organization*: The rest of the paper is organized as follows. In Section II, we formulate the problem we aim to solve, as well as revisit some classical results that will be leveraged in this work. The notion of recurrence to be used in this work is introduced in Section III, together with our first core set of results that show the relationship between recurrence and containment within the region of attraction. The proposed algorithms and the corresponding guarantees are given in Section IV. Numerical examples are provided in Section V and we conclude in Section VI.

II. PROBLEM FORMULATION

We consider a continuous time dynamical system

$$\dot{x}(t) = f(x(t)), \quad (1)$$

where $x(t) \in \mathbb{R}^d$ is the state at time t , and the map $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is continuously differentiable and (globally) Lipschitz. Given initial condition $x(0) = x_0$, we use $\phi(t, x_0)$ to denote the solution of (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 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 $\lim_{n \rightarrow \infty} t_n = \infty$ and $\lim_{n \rightarrow \infty} \phi(t_n, x_0) = y$. We will further use $\Omega(f)$ to denote the ω -limit set of (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 (1), then it follows that $x^* \in \Omega(f)$.

A. Region of Attraction

We would like then to learn the set of initial conditions that converge to x^* .

Definition 2 (Region of Attraction). *Given an invariant set $S \subseteq \Omega(f)$, the region of attraction (ROA) of S under (1) is defined as*

$$\mathcal{A}(S) := \left\{ x_0 \in \mathbb{R}^d \mid \liminf_{t \rightarrow \infty} d(\phi(t, x_0), S) = 0 \right\}, \quad (2)$$

where $d(\phi(t, x_0), S)$ is the distance from the solution $\phi(t, x_0)$ to the set S , i.e., $d(\phi(t, x_0), S) := \min_{x \in S} \|x - \phi(t, x_0)\|_2$.

When the set S is a singleton that contains exactly one point (say x), we abbreviate $\mathcal{A}(S) = \mathcal{A}(\{x\})$ as $\mathcal{A}(x)$.

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 [1].

Definition 3 (Positively Invariant Set). *A set $\mathcal{I} \subseteq \mathbb{R}^d$ is positively invariant w.r.t. (1) if and only if:*

$$x_0 \in \mathcal{I} \implies \phi(t, x_0) \in \mathcal{I}, \quad \forall t \in \mathbb{R}^+. \quad (3)$$

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. Unfortunately, without further assumptions, the set (2) may be a singleton, have zero measure, or be disconnected, making the problem of characterizing (2) from samples almost impossible. We thus make the following assumption.

Assumption 1. *The system (1) has an asymptotically stable equilibrium at $x^* \in \mathbb{R}^d$.*

Remark 1. *It follows from Assumption 1 that the (positively invariant) ROA $\mathcal{A}(x^*)$ is an open contractible set [24], i.e., the identity map of $\mathcal{A}(x^*)$ to itself is null-homotopic [25].*

As a result, Assumption 1 provides the necessary regularity conditions for $\mathcal{A}(x^*)$ to be of practical use. A natural approach under this setting is therefore to search for Lyapunov functions [1] 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 [26] that guarantees the existence of such a function:

Theorem 1 (Zubov's Existence Criterion). *A set \mathcal{A} containing x^* in its interior is the region of attraction of x^* under (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 $x \in \mathcal{A} \setminus \{x^*\}$, $h(x) > 0$ for $x \in \mathbb{R}^d \setminus \{x^*\}$.
- For every $\gamma_2 > 0$, there exists $\gamma_1 > 0$, $\alpha_1 > 0$ such that $V(x) > \gamma_1$, $h(x) > \alpha_1$, whenever $\|x\| \geq \gamma_2$.
- $V(x_k) \rightarrow 1$ for all sequences $\{x_k\}$ such that $x_k \rightarrow \partial\mathcal{A}$ or $\|x_k\| \rightarrow \infty$.
- V and h satisfy

$$(\mathcal{L}_f V)(x) = -h(x)(1 - V(x))\sqrt{1 + \|f(x)\|^2}, \quad (4)$$

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 1. *Under Assumption 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. Let V be the Zubov's function whose existence is guaranteed by Theorem 1. Thus by the definition of V , for $c \in (0, 1)$, $V_{\leq c} \subseteq \mathcal{A}(x^*)$. Further from (4), it follows that $(\mathcal{L}_f V)(x) < 0$, 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 $H : [0, 1] \times V_{\leq c} \rightarrow V_{\leq c}$ such that $H(0, x) = x$ and $H(1, x) = x^*$ for all $x \in V_{\leq c}$. Similar to [24], 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 [1]. 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 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. \square

The Zubov's function V of Theorem 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 (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 (3). This motivates the alternative proposed in the next section.

III. RECURRENT SETS

We now introduce the relaxed notion of invariance to be used in this paper, which we refer to here as recurrence. We will then illustrate how recurrent sets constitute a more flexible and general class of objects to study.

Definition 4 (Recurrent Set). *A set $\mathcal{R} \subseteq \mathbb{R}^d$ is recurrent w.r.t. (1), if for any point $x_0 \in \mathcal{R}$ and any time $t \geq 0$, there exists a time $t' \geq t$, such that $\phi(t', x_0) \in \mathcal{R}$.*

Note that a recurrent set, while not invariant, guarantees that solutions starting in this set will visit it back infinitely often. In particular, by Definition 3, a positively invariant set \mathcal{I} is recurrent. Thus, Definition 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. 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 results show that under mild assumptions, this should not be a source of concern.

Lemma 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. We will prove this statement by contradiction. Assume the result does not hold, i.e., there exist $x_0 \in \mathcal{R}$ and $t' \geq t$ such that $\phi(t', x_0) \notin \mathcal{R}$ for any $t > 0$. This, together with the definition of the recurrent set (Definition 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 1) that $\bar{x} = \lim_{i \rightarrow \infty} x_{n_i} \in \Omega(f) \cap \partial\mathcal{R}$, which contradicts with the assumption that $\partial\mathcal{R} \cap \Omega(f) = \emptyset$. \square

Theorem 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. (\implies): If \mathcal{R} is a compact recurrent set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$, Lemma 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 [1, p. 127] that the limit set $\Omega(x_0) \neq \emptyset$ and $\lim_{t \rightarrow \infty} d(\phi(t, x_0), \Omega(x_0)) = 0$. Therefore, we conclude $\Omega(f) \cap \mathcal{R} \supseteq \Omega(x_0) \cap \mathcal{R} \neq \emptyset$ 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 $\Omega_{\zeta}^{\mathcal{R}} \subset \mathcal{R}$.

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., $\liminf_{t \rightarrow \infty} d(\phi(t, x_0), \Omega(f) \cap \mathcal{R}) = 0$. 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. \square

Theorem 2 illustrates that recurrence necessarily implies 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.

Assumption 2. *The ω -limit set $\Omega(f)$ of (1) is composed of hyperbolic equilibrium points, with only one of them, say x^* , being asymptotically stable.*

Corollary 2. *Let assumptions 1 and 2 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. (\implies): By assumption \mathcal{R} is compact, and $\partial\mathcal{R} \cap \Omega(f) = \emptyset$. Then, Theorem 2 implies that if \mathcal{R} is recurrent then $\Omega(f) \cap \mathcal{R} \neq \emptyset$ and $\mathcal{R} \subset \mathcal{A}(\Omega(f) \cap \mathcal{R})$. This, together with the other assumption that $\Omega(f) \cap \mathcal{R} = \{x^*\}$, implies $\mathcal{R} \subset \mathcal{A}(x^*)$.

(\impliedby): This direction is trivial given Theorem 2. \square

Corollary 2 implies that from a practical standpoint, one may use recurrence as a mechanism for finding inner approximations for $\mathcal{A}(x^*)$. However, one limitation of the above results is that although \mathcal{R} is recurrent, we do not know a priori how long it may take for a trajectory to come back to \mathcal{R} after it leaves it. This motivates the following stricter notion of recurrence.

Definition 5 (τ -Recurrent Set). *A set $\mathcal{R} \subseteq \mathbb{R}^d$ is τ -recurrent w.r.t. (1), if for any point $x_0 \in \mathcal{R}$ and any time $t \geq 0$, there exists a $t' \in [t, t + \tau]$, such that $\phi(t', x_0) \in \mathcal{R}$.*

Theorem 3. *Let Assumption 1 hold, and consider a compact set $\mathcal{R} \subseteq \mathcal{A}(x^*)$ satisfying $x^* \in \text{int } \mathcal{R}$ 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. The proof of the theorem relies on Zubov's existence criterion stated in Theorem 1. Given \mathcal{R} , let us now define

$$\underline{c} := \min_{x \in \partial\mathcal{R}} V(x), \quad \bar{c} := \max_{x \in \partial\mathcal{R}} V(x),$$

$$\text{and } a := \max_{x \in C} \nabla V(x)^T f(x),$$

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

We first argue that $V_{\leq \underline{c}} := \{x : V(x) \leq \underline{c}\} \subseteq \mathcal{R}$. Let \underline{x} be the point in $\partial\mathcal{R}$ that achieves the minimum, i.e., $V(\underline{x}) = \underline{c}$, and let \mathcal{R}' be 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 $x' \in \partial\mathcal{R}$ with $V(x') < \underline{c}$; which contradicts the definition of \underline{c} . Thus, $x^* \in \mathcal{R}' \subseteq \mathcal{R}$.

Suppose then that $V_{\leq \underline{c}} \not\subseteq \mathcal{R}' \subseteq \mathcal{R}$, for any point $\tilde{x} \in V_{\leq \underline{c}} \setminus \mathcal{R}'$, $V(\phi(t, \tilde{x})) < \underline{c}$, for $t > 0$, and $\lim_{t \rightarrow \infty} \phi(t, \tilde{x}) = x^*$. Thus there exists $\tilde{t} > 0$ s.t. $V(\phi(\tilde{t}, \tilde{x})) < \underline{c}$ and $\phi(\tilde{t}, \tilde{x}) \in \partial\mathcal{R}$; contradiction. It follow then that $V_{\leq \underline{c}} \subseteq \mathcal{R}' \subseteq \mathcal{R}$.

Similarly, since the contradictable set $V_{\leq \bar{c}}$ contains every point in the boundary of \mathcal{R} , there cannot be any point in




Fig. 1: An visualization of the proof of Theorem 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 .

$x \in \mathcal{R}$ with $V(x) > \bar{c}$. We therefore get that the following inclusions must hold:

$$V_{\leq \underline{c}} \subseteq \mathcal{R} \subseteq V_{\leq \bar{c}}. \quad (5)$$

Finally, by (5), 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 $t \geq \bar{\tau} := \frac{\bar{c} - \underline{c}}{a}$ the Lyapunov value $V(\phi(t, x)) \leq \underline{c}$, which implies that $\phi(t, x) \in \mathcal{R}$ and result follows. \square

Note that the lower bound on τ in Theorem 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 1, we consider the set

$$\tilde{\mathcal{A}} := V_{\leq c} \setminus \{\text{int } \mathcal{B}_\delta + x^*\}, \quad (6)$$

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 (6) represents the Minkowski sum, and \mathcal{B}_δ is a closed δ ball centered at the origin, i.e., $\mathcal{B}_\delta = \{x \mid \|x\|_2 \leq \delta\}$. 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 $c \rightarrow 1$ in the case that $\mathcal{A}(x^*)$ is bounded.

Then, by denoting $\underline{c}(\delta)$ as the min 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.,

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

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

Theorem 4. *Let Assumption 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 $\tilde{x} \in V_{\leq \underline{c}(\delta)} \setminus \{\mathcal{B}_\delta + x^*\}$, $\lim_{t \rightarrow \infty} \phi(t, \tilde{x}) = x^*$ 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 $\phi(t, x) \in \mathcal{R}$. \square

IV. LEARNING RECURRENT SETS

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 (1):

$$x_n = \phi(n\tau_s, x_0), \quad x_0 \in \mathbb{R}^d, \quad n \in \mathbb{N}, \quad (7)$$

where $\tau_s > 0$ is the sampling period.

In this setting, we define the notion of discrete recurrence w.r.t. a length k trajectory:

Definition 6 (*k*-Recurrent Set). *A set $\mathcal{R} \subseteq \mathbb{R}^d$ is k -steps recurrent (k -recurrent for short) w.r.t. (7), if for any point $x_0 \in \mathcal{R}$ and any step index $n \geq 0$, there exists an $n' \in \{n, \dots, n+k\}$, such that $x_{n'} \in \mathcal{R}$.*

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. However, the converse is not necessarily true.*

To ensure one can find such a k -recurrent set, we consider again the specific set $\tilde{\mathcal{A}}$ defined in (6) that gives the following sufficient conditions for a set \mathcal{R} to be k -recurrent.

Theorem 5. *Let Assumption 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 $k > \bar{k} := \bar{\tau}(\delta)/\tau_s$, where $\bar{\tau}(\delta)$ is defined as in Theorem 4.*

Proof. Given Theorem 4, this result follows directly from $\phi(t, x) \in \mathcal{R}$ for all $x \in \mathcal{R}$ when $t \geq \bar{\tau}(\delta)$. \square

In the rest of the paper, we assume for simplicity that the ω -limit set $\Omega(f)$ of (1) is the set of all the hyperbolic equilibrium points with a unique asymptotically stable one (Assumption 2). We further assume w.l.o.g. that the asymptotically stable equilibrium is at the origin, i.e., $x^* = 0$. We briefly explain next the underlying mechanism that will be used to learn recurrent sets.

Algorithm Summary: We will restrict our search to a compact initial approximation $\hat{\mathcal{S}}^{(0)} \subset \mathbb{R}^d$ of the ROA satisfying $\hat{\mathcal{S}}^{(0)} \supseteq \mathcal{B}_\delta$. Precisely, we will seek to find a subset of the ROA within $\mathcal{A}(x^*) \cap \hat{\mathcal{S}}^{(0)}$ by computing k -recurrent sets \mathcal{R} that seek to satisfy the properties of Theorem 5. In this approach, starting from $\hat{\mathcal{S}}^{(0)}$, we sequentially generate a sequence of approximations $\hat{\mathcal{S}}^{(i)}$. For each $\hat{\mathcal{S}}^{(i)}$, we sample points $p_{ij} \in \hat{\mathcal{S}}^{(i)}$ and check whether a trajectory of length k that starts at p_{ij} returns to $\hat{\mathcal{S}}^{(i)}$ for each $j = 0, 1, 2, \dots$. A trajectory that does not return to $\hat{\mathcal{S}}^{(i)}$ within k steps is

a counter-example of k -recurrence. Once a counter-example is found, we update the approximation $\hat{\mathcal{S}}^{(i)}$ to $\hat{\mathcal{S}}^{(i+1)}$ and restart the sampling process. 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 k -recurrent set

```

Initialize  $\hat{\mathcal{S}}_0$  according to (9)
for Iteration  $i = 0, 1, \dots$  do
    for Iteration  $j = 0, 1, \dots$  do
        Generate random sample  $p_{ij} \in \hat{\mathcal{S}}^{(i)}$  uniformly
        if  $p_{ij}$  is a counter-example w.r.t  $\hat{\mathcal{S}}^{(i)}$  then
            Update  $\hat{\mathcal{S}}^{(i)}$  according to (10)
            break
        end
    end
end

```

A. Classification of sample points

We say that a sample point p_{ij} is a valid k -recurrent point w.r.t current approximation $\hat{\mathcal{S}}^{(i)}$ if starting from $x_0 = p_{ij}$,

$$\exists n \in \{1, \dots, k\}, \text{ s.t. } x_n \in \hat{\mathcal{S}}^{(i)}. \quad (8)$$

If (8) does not hold, we say p_{ij} is a *counter-example*. We will use such counter-examples to update our current set approximation $\hat{\mathcal{S}}^{(i)}$.

B. Construction of set approximations

In order to gradually update the sets $\hat{\mathcal{S}}^{(i)}$, we consider two parametric families of set approximations.

1) *Sphere approximation:* To construct a sphere approximation, we start by choosing a radius $\bar{b} > 0$ large enough such that the set

$$\hat{\mathcal{S}}^{(0)} := \{x \mid \|x\|_2 \leq b^{(0)} := \bar{b}\} \supseteq \mathcal{B}_\delta. \quad (9)$$

The sphere approximation for iteration i is then defined as $\hat{\mathcal{S}}^{(i)} := \{x \mid \|x\|_2 \leq b^{(i)}\}$. Finally, given a sample point $p_{ij} \in \hat{\mathcal{S}}^{(i)}$ we update $\hat{\mathcal{S}}^{(i)}$ based on the following criterion:

$$p_{ij} \text{ counter-example} \implies b^{(i+1)} = \|p_{ij}\|_2 - \varepsilon, \quad (10)$$

where $\varepsilon > 0$ is an algorithm parameter expressing the level of conservativeness in our update.

If the process reaches a value of $b^{(i)} < 0$, 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) *Polyhedron approximation:* One could also choose to construct $\hat{\mathcal{S}}^{(i)}$ using polyhedron approximation. To that end, we first construct a matrix $A := [a_1, \dots, a_n]^T \in \mathbb{R}^{n \times d}$, where each row vector a_n is a *normalized* ($\|a_l\| = 1 \forall l$) exploration direction indexed by $l \in \{1, \dots, n\}$. Precisely, we generate A randomly with row dimension n large enough such that for




Fig. 2: An illustration of the proof of Theorem 6. In particular, given an arbitrary point $p_{ij} \notin V_{\leq c}$, in the sphere case, it follows that $\|p_{ij}\|_2 \geq r$. And in the polyhedron case, the closest projection $\max_{l \in \{1, \dots, n\}} a_l^T p_{ij} \geq r/2$ under Assumption (11) that every pair of exploration directions are close enough.

any arbitrary direction $v \in \mathbb{R}^d$, there exists an exploration direction a_l with an angle $\theta_{v,l}$ between v and a_l satisfying

$$\theta_{v,l} := \arccos \left(\frac{v^T a_l}{\|v\| \|a_l\|} \right) \leq \frac{2}{3} \pi. \quad (11)$$

Note that the aforementioned process is analog to constructing an ε -net over a unit Euclidean hyper-sphere, for which several algorithms exist [27], [28]. Upper bounds on the size of n can also be found in the literature, see, e.g., [29].

The polyhedron approximation at iteration i is defined as $\hat{S}^{(i)} := \{x | Ax \leq b^{(i)}\}$, consisting on n inequalities aimed at approximating a k -recurrent set via counter examples. In this paper, keep A fixed for every iteration, and update the constraint coefficients $b^{(i)} = [b_1^{(i)}, \dots, b_n^{(i)}]^T \in \mathbb{R}^n$.

Similarly to the sphere approximation, we initialize $b^{(0)} := \bar{b} \mathbb{1}_n$, such that

$$\hat{S}^{(0)} = \{x | Ax \leq b^{(0)} := \bar{b} \mathbb{1}_n\} \supseteq \mathcal{B}_\delta, \quad (12)$$

which can be done for $\bar{b} > 0$ large enough; here $\mathbb{1}_n \in \mathbb{R}^n$ is the vector of all ones.

Afterwards $b^{(i)}$ is updated by sampling points $p_{ij} \in \hat{S}^{(i)}$ and checking the following criterion:

p_{ij} **counter-example**

$$\Rightarrow \begin{cases} b_{l^*}^{(i+1)} = a_{l^*}^T p_{ij} - \varepsilon \\ b_l^{(i+1)} = b_l^{(i)}, \quad \forall l = \{1, \dots, n\} \setminus l^*, \end{cases}$$

where $\varepsilon > 0$ is fixed and $l^* = \arg \max_{l \in \{1, \dots, n\}} \frac{a_l^T p_{ij}}{\|a_l\| \|p_{ij}\|}$ is the index of exploration direction that minimizes the angle between p_{ij} and a_l . If l^* consists of more than one index, we simply choose one at random. As before, we declare the search a failure whenever $b_l^{(i)} < 0$ for some l , since this implies that the equilibrium $x^* = 0$ is outside the set $\hat{S}^{(i)}$.

C. Bound on the number of updates

As mentioned before, the aforementioned search for approximations will fail if $b^{(i)}$ (sphere) or one of $b_l^{(i)}$ (polyhedron) becomes negative at some iteration $i > 0$. 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 $\hat{S}^{(i)}$ satisfying $\mathcal{B}_\delta \subseteq \hat{S}^{(i)} \subseteq V_{\leq c}$, then Theorem 5 guarantees that any sample $p_{ij} \in \hat{S}^{(i)}$ will lead to a k -recurrent trajectory, i.e., condition (8). As a result, the algorithm will stop updating at this point since we cannot find further counter-examples within $\hat{S}^{(i)}$. This means that, if it is possible for $\hat{S}^{(i)}$ to become a subset of $V_{\leq c}$, without violating the condition $\mathcal{B}_\delta \subseteq \hat{S}^{(i)}$, 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 6. *Let the initial approximation $\hat{S}^{(0)}$ satisfy $\mathcal{B}_\delta \subseteq \hat{S}^{(0)}$ and trajectory length $k > \bar{k}$, for \bar{k} as defined in Theorem 5. Then, given a counter-example p_{ij} , the resulting updated set satisfies $\mathcal{B}_\delta \subseteq \hat{S}^{(i+1)}$ whenever*

$$\varepsilon \leq \begin{cases} r - \delta & \text{with sphere approximation} \\ \frac{r}{2} - \delta & \text{with polyhedron approximation,} \end{cases} \quad (13)$$

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

Proof. Given an arbitrary counter-example p_{ij} w.r.t $\hat{S}^{(i)} \supseteq \mathcal{B}_\delta$, it follows that $p_{ij} \notin V_{\leq c}$ by Theorem 5; since otherwise, p_{ij} would generate a k -recurrent trajectory. Then, it follows from the definition of r that $\|p_{ij}\|_2 \geq r$, as illustrated in Figure 2. Further, let $\mathcal{B}_r := \{x | \|x\|_2 \leq r\} \subseteq V_{\leq c}$.

We now reason differently depending on the type of approximation.

(Sphere case): It then follows from $\|p_{ij}\|_2 \geq r$ that whenever $\varepsilon \leq r - \delta$, the update leads to $b^{(i+1)} = \|p_{ij}\|_2 - \varepsilon \geq r - \varepsilon \geq \delta$.

(Polyhedron case): It follows from (11), that for any point $p' \notin \mathcal{B}_r$, we have $\max_{l \in \{1, \dots, n\}} a_l^T p' \geq \|p'\| \cos(\frac{2}{3}\pi) \geq \frac{r}{2}$. Therefore, since by definition of \mathcal{B}_r , $p_{ij} \notin \mathcal{B}_r$ we conclude then that $b_{l^*}^{(i+1)} = a_{l^*}^T p_{ij} - \varepsilon \geq \frac{r}{2} - \varepsilon \geq \delta$.

Together with the fact that $\hat{S}^{(0)} \supseteq \mathcal{B}_\delta$, result follows. \square

Theorem 6 establishes that one can choose parameters k and ε so that the sequence of sets $\hat{S}^{(i)}$ never leads to $b^{(i)}$ or $b_l^{(i)}$ negative, 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 sets back to \hat{S}^0 .

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

Theorem 7. *Given the initial approximation $\hat{S}^{(0)} \in \mathcal{F}_{\bar{b}}$ and initial constant \bar{b} defined in (9) or (12), the total number*

of counter-examples encountered in Algorithm 1, with k -doubling after each failure, is bounded by $\frac{\bar{b}}{\varepsilon} \log_2 \bar{k}$ in the sphere case and $n \frac{\bar{b}}{\varepsilon} \log_2 \bar{k}$ in the polyhedron case.

Proof. Note that once a counter-example is encountered, we decrease the radius constraint (sphere case) or one of the exploration directions (polyhedron case) by at least ε . Therefore, $\hat{S}^{(i)} \in \mathcal{F}_{\bar{b}}$ for all $i \in \{1, 2, \dots\}$. And for any fixed k , our method can find at most $\frac{\bar{b}}{\varepsilon}$ counter-examples with the sphere approximation and $n \frac{\bar{b}}{\varepsilon}$ counter-examples with the polyhedron approximation 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, result follows. \square

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

D. Convergence guarantee

By Definition 6, a set \hat{S} is k -recurrent if every point $p \in \hat{S}$ satisfies (8). As shown before, certifying this property will enable us to guarantee that $\hat{S} \subset \mathcal{A}(0)$. However, it is infeasible to enforce condition (8) for every point in \hat{S} . Instead, we will show that under mild conditions, our algorithm converges to a \hat{S}^* satisfying $\text{int } \hat{S}^* \subseteq \mathcal{A}(0)$ with probability one.

In our algorithm, we generate samples p_j uniformly within some set \hat{S} , i.e., $p_j \stackrel{iid}{\sim} U(\hat{S})$ for all $j \in \{0, 1, 2, \dots\}$. We use \hat{S}_{counter} to denote the set that contains all the counter-examples in \hat{S} that certify \hat{S} being not k -recurrent, i.e.,

$$\hat{S}_{\text{counter}} := \{p \in \hat{S} \mid p \text{ is a counter-example}\}.$$

Given a random sample p_j we define the Bernoulli random variable X_j with $X_j = 1$ if $p_j \in \hat{S}_{\text{counter}}$ and $X_j = 0$ otherwise.

Lemma 2. *Consider a set $\hat{S} \in \mathcal{F}_{\bar{b}}$ satisfying $\partial \hat{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \hat{S} = \{0\}$, if $\text{int } \hat{S} \not\subseteq \mathcal{A}(0)$, then there exists a point $p \in \text{int } \hat{S} \setminus \mathcal{A}(0)$ and a time $t' > 0$ such that $\phi(t, p) \notin \hat{S}$, $\forall t > t'$. Moreover, the point p could be selected such that t' is arbitrarily close to zero.*

Proof. Let us consider a point $p \in \text{int } \hat{S} \setminus \mathcal{A}(0)$, we claim that either $p \notin \mathcal{A}(\Omega(f))$ or $p \in \mathcal{A}(\Omega(f) \setminus \{0\})$ is true since $\mathcal{A}(\Omega(f) \setminus \{0\}) \cap \mathcal{A}(0) = \emptyset$.

Consider first a point $p \notin \mathcal{A}(\Omega(f))$, we have $\phi(p, t) \rightarrow \infty$ as $t \rightarrow \infty$. Since \hat{S} is compact, the first result follows in this case.

In the other case that $p \in \mathcal{A}(\Omega(f) \setminus \{0\})$, it follows from the definition of the regions of attraction (Definition 2) that $\liminf_{t \rightarrow \infty} d(\phi(t, p), \Omega(f) \setminus \{0\}) = 0$. Note that by assumption \hat{S} is a compact set and $\{\Omega(f) \setminus \{0\}\} \cap \hat{S} = \emptyset$. Therefore, if the result does not follow, i.e., for all $t' > 0$ there exists a $t > t'$ such that $\phi(t, p) \in \hat{S}$, we can construct an infinite sequence $\{x_n\}_{n=0}^{\infty} \subset \partial \hat{S}$ as in the

proof of Lemma 1. Then, since $\partial \hat{S}$ 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 $\bar{x} \in \partial \hat{S}$. It follows then from the definition of ω -limit sets (Definition 1) that $\bar{x} = \lim_{i \rightarrow \infty} x_{n_i} \in \Omega(f) \cap \partial \hat{S}$, which contradicts with the assumption that $\partial \hat{S} \cap \Omega(f) = \emptyset$.

Now, since the first result follows, we can additionally let $q = \phi(t' - \delta, p)$ and conclude $\phi(t, q) \notin \hat{S}$, $\forall t > \delta$ with $\delta > 0$ arbitrarily close to zero, which implies the other result. \square

Lemma 3. *For any set $\hat{S} \in \mathcal{F}_{\bar{b}}$ satisfying $\partial \hat{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \hat{S} = \{0\}$, the volume of its counter-example part is positive, i.e., $\text{vol}(\hat{S}_{\text{counter}}) > 0$, whenever $\text{int } \hat{S} \not\subseteq \mathcal{A}(0)$.*

Proof. If $\text{int } \hat{S} \not\subseteq \mathcal{A}(0)$, then Lemma 2 implies that there exists a point $p \in \text{int } \hat{S} \setminus \mathcal{A}(0)$ and a time $t' \in (0, \tau_s)$ such that $\phi(t, p) \notin \hat{S}$, $\forall t > t'$. Then, for any $k = 1, 2, 3, \dots$, we can respectively construct a neighborhood of p that consists of counter-examples of k -recurrent.

Precisely, let us consider an arbitrary point $q \in \hat{S}$ and recall the assumption that the dynamical system (1) is globally L -Lipschitz, it follows from [1, p. 96] that the distance between solutions $\|\phi(t, p) - \phi(t, q)\| \leq \|p - q\| \exp(Lt)$ for all $t \geq 0$. Therefore, we choose q such that

$$\|p - q\| < \underline{d} := \min_{n=1, \dots, k} d(\phi(n\tau_s, p), \hat{S}) / \exp(Lkn\tau_s),$$

and claim $\|\phi(n\tau_s, p) - \phi(n\tau_s, q)\| \leq \|p - q\| \exp(Ln\tau_s) < d(\phi(n\tau_s, p), \hat{S})$, $\forall n = 1, \dots, k$, i.e., q is a counter-example of k -recurrence.

Finally, since $p \in \text{int } \hat{S}$, the aforementioned counter-example set $\{q \in \hat{S} \mid \|p - q\| < \underline{d}\}$ has positive volume, and thus the result follows. \square

Lemma 4. *For any set $\hat{S} \in \mathcal{F}_{\bar{b}}$ satisfying $\text{int } \hat{S} \not\subseteq \mathcal{A}(0)$, $\partial \hat{S} \cap \Omega(f) = \emptyset$ and $\Omega(f) \cap \hat{S} = \{0\}$, we have $\lim_{m \rightarrow \infty} \mathbb{P}(X_1 = \dots = X_m = 1) = 0$. That is, a counter-example is eventually sampled almost surely.*

Proof. Note that we have $\hat{S}_{\text{counter}} \subseteq \hat{S}$ and $\text{vol}(\hat{S}_{\text{counter}}) > 0$ by Lemma 3. Then, denoting the counter-example ratio as $\rho := \text{vol}(\hat{S}_{\text{counter}}) / \text{vol}(\hat{S})$, one can conclude $0 < \rho \leq 1$ and

$$\lim_{m \rightarrow \infty} \mathbb{P}(X_0 = \dots = X_m = 1) = \lim_{m \rightarrow \infty} (1 - \rho)^m = 0.$$

\square

We now leverage the results in Lemma 3 and Lemma 4 to obtain the following termination guarantee.

Theorem 8. *Consider $\hat{S}^{(0)} \in \mathcal{F}_{\bar{b}}$ with $\hat{S}^{(0)} \supseteq \mathcal{B}_{\delta}$ and $\Omega(f) \cap \hat{S}^{(0)} = \{0\}$. Then, after a finite number of iterations, the updates on $\hat{S}^{(i)}$ terminate at some \hat{S}^* whose interior is a non-empty subset of $\mathcal{A}(0)$ whenever $k > \bar{k}$ and (13) holds.*

Proof. Suppose that at any given iteration i the set $\text{int } \hat{S}^{(i)} \not\subseteq \mathcal{A}(0)$. Then it follows from Lemma 4 that a counter-example is eventually found almost surely, and a new set $\hat{S}^{(i+1)}$ is obtained. Also Theorem 7 implies the total number of such transitions is finite, since $\hat{S}^{(0)} \in \mathcal{F}_{\bar{b}}$ and $\mathcal{B}_{\delta} \subseteq \hat{S}^{(0)}$.

Now let $\hat{\mathcal{S}}^*$ denote the last updated approximation. Note that since there are not further updates to $\hat{\mathcal{S}}^*$ with probability one, this implies that $\text{vol}(\hat{\mathcal{S}}_{\text{counter}}^*) = 0$. We argue then that $\text{int } \hat{\mathcal{S}}^* \subseteq \mathcal{A}(0)$, since otherwise $\text{vol}(\hat{\mathcal{S}}_{\text{counter}}^*) > 0$, which contradicts the fact that $\hat{\mathcal{S}}^*$ is the last iteration. Finally, $\text{int } \hat{\mathcal{S}}^*$ is non-empty since Theorem 6 implies $\hat{\mathcal{S}}^* \supseteq \mathcal{B}_\delta$. \square

E. Multiple center point approximation

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

In particular, we consider $h \in \mathbb{N}^+$ center points x_q indexed by $q \in \{1, 2, \dots, h\}$, where the first center point as $x_1 = x^* = 0$. Then other centers, i.e., x_2, \dots, x_h , can be chosen uniformly within some region of interest or selected to be in some preferred place. At each center point x_q the sphere approximation is defined by $\hat{\mathcal{S}}_q^{(i)} := \{x \mid \|x - x_q\|_2 \leq b_q^{(i)}\}$, where $b_q^{(i)}$ represents the radius to be updated in the presence of counter-examples. As before we initialize $b_q^{(0)} = \bar{b}$. In the case of polyhedral approximations, we similarly define $b_q^{(i)} = [b_{q,1}^{(i)}, \dots, b_{q,n}^{(i)}]^T \in \mathbb{R}^n$, with $b_q^{(0)} = \bar{b}1^N$ and let $\hat{\mathcal{S}}_q^{(i)} := \{x \mid A(x - x_q) \leq b_q^{(i)}\}$.

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

Similar to Algorithm 1, sample points p_{ij} are generated uniformly within $\hat{\mathcal{S}}_{\text{multi}}^{(i)}$ in each sub-iteration $j = 1, 2, \dots$. In this multi-center case, p_{ij} is classified as a counter-example if starting from $x_0 = p_{ij}$, $x_n \notin \hat{\mathcal{S}}_{\text{multi}}^{(i)}$ for all $n \in \{1, \dots, k\}$. Once encountered a counter-example, we update $\hat{\mathcal{S}}_{\text{multi}}^{(i)}$ and restart sampling iteration j . In particular, given a counter-example $p_{ij} \in \hat{\mathcal{S}}_{\text{multi}}^{(i)}$, every approximations $\hat{\mathcal{S}}_q^{(i)}$ (sphere or polyhedron) satisfying $p_{ij} \in \hat{\mathcal{S}}_q^{(i)}$ are subjected to update respectively via the following criterion:

$$\begin{aligned} \text{(sphere)} \quad & b_q^{(i+1)} = \|p_{ij} - x_q\|_2 - \varepsilon \\ \text{(polyhedron)} \quad & \begin{cases} b_{q,l^*}^{(i+1)} = a_{l^*}(p_{ij} - x_q) - \varepsilon \\ b_{q,l}^{(i+1)} = b_{q,l}^{(i)}, \forall l \in \{1, \dots, n\} \setminus l^*, \end{cases} \end{aligned}$$

where $l^* = \arg \max_{l \in \{1, \dots, n\}} \frac{a_l^T (p_{ij} - x_q)}{\|a_l\| \|p_{ij} - x_q\|}$. Again, we choose one at random if l^* consists of more than one index.

Then, those approximations not containing p_{ij} are updated as $\hat{\mathcal{S}}_q^{(i+1)} = \hat{\mathcal{S}}_q^{(i)}$. Note that the parameter ε is strictly positive. Thus, for all center points $x_q \notin \mathcal{A}(0)$, the corresponding constraint parameters $b_q^{(i)}$ could decrease to negative values and result in $\hat{\mathcal{S}}_q^{(i)} = \emptyset$ without affecting our results.

In this multi-center setting, we use \mathcal{F}_b^h to denote the parametric family of h closed balls (resp. polytopes) defined

by $\cup_{q=1}^h \mathcal{S}_q$, where $\mathcal{S}_q = \{x : \|x - x_q\|_2 \leq b_q\}$ (resp. $\mathcal{S}_q = \{x : A(x - x_q) \leq b_q\}$), for $b_q \in [0, \bar{b}]$ (resp. $b_q \in [0, \bar{b}]^n$) and $x_q \in \mathbb{R}^d$ indexed by $q = \{1, \dots, h\}$.

Theorem 9. *For any iteration $i \in \mathbb{N}^+$, the multi-center approximation $\hat{\mathcal{S}}_{\text{multi}}^{(i)}$ is non-vanishing, i.e., $\hat{\mathcal{S}}_{\text{multi}}^{(i)} \supseteq \mathcal{B}_\delta$, if $k > \bar{k}$ and condition (13) is satisfied. The total number of counter-examples encountered, with k -doubling after each failure, is bounded by $h \frac{\bar{b}}{\varepsilon} \log_2 \bar{k}$ and $nh \frac{\bar{b}}{\varepsilon} \log_2 \bar{k}$ in the sphere and polyhedron case respectively. Moreover, the last updated multi-center approximation $\hat{\mathcal{S}}_{\text{multi}}^*$ satisfies $\text{int } \hat{\mathcal{S}}_{\text{multi}}^* \neq \emptyset$ and $\text{int } \hat{\mathcal{S}}_{\text{multi}}^* \subseteq \mathcal{A}(0)$ whenever $\Omega(f) \cap \hat{\mathcal{S}}_{\text{multi}}^{(0)} = \{0\}$.*

Proof. By definition $\hat{\mathcal{S}}_{\text{multi}}^{(i)} \supseteq \hat{\mathcal{S}}_1^{(i)}$ for all $i \in \mathbb{N}^+$, Theorem 6 therefore implies $\hat{\mathcal{S}}_{\text{multi}}^{(i)} \supseteq \hat{\mathcal{S}}_1^{(i)} \supseteq \mathcal{B}_\delta$ under (13). The bound on the total number of counter-examples follows as in Theorem 7, since every additional approximation $\hat{\mathcal{S}}_q^{(i)} \in \mathcal{F}_{\bar{b}}$ for all $q \in \{1, \dots, h\}$ and iteration $i \in \{1, \dots\}$. Finally, by generalizing Lemma 2-4 and Theorem 8 to the scope of $\mathcal{F}_{\bar{b}}^h$, the last statement follows. \square

V. EXPERIMENTS

We illustrate the accuracy of the proposed methodology by approximating the region of attraction of the following autonomous dynamical system:

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

The black dotted area in Figure 3 represents the complement of ROA of the origin, which is computed by testing a mesh grid of points. A point is marked black if it does not converge to the equilibrium after $t = 30$. In our algorithm we set $\varepsilon = 0.1$, $k = 50$ and $\tau_s = 0.5$. To estimate the number of iterations until convergence, we stop our algorithm when all black dots are excluded from our current approximation.

The outcomes of our approximation are marked in green. In particular, Figure 3 (left two panels) shows the outcome of applying Algorithm 1 using a sphere and a $n = 200$ directions polyhedron approximation. To address the problem of under-estimation, as shown in the right two panels of Figure 3, we can generate random center points. Fifty spheres or ten polyhedrons approximates with random center points give a good approximation. Detailed statistics of our algorithms for the aforementioned scenarios are provided in Table I. Notably, the number of counter-examples and the steps simulated per sample is small, which illustrates the efficiency of our algorithm.

VI. CONCLUSIONS AND FUTURE WORK

We consider the problem of learning the region of attraction of a stable equilibrium point. We propose the use of a more flexible notion of invariance known as recurrence. We provide necessary and sufficient conditions for a recurrent set to be an inner approximation of the ROA. Our algorithms are sequential and only incur a limited number of counter-examples. Future work includes extending our framework to other families of approximations and control design.

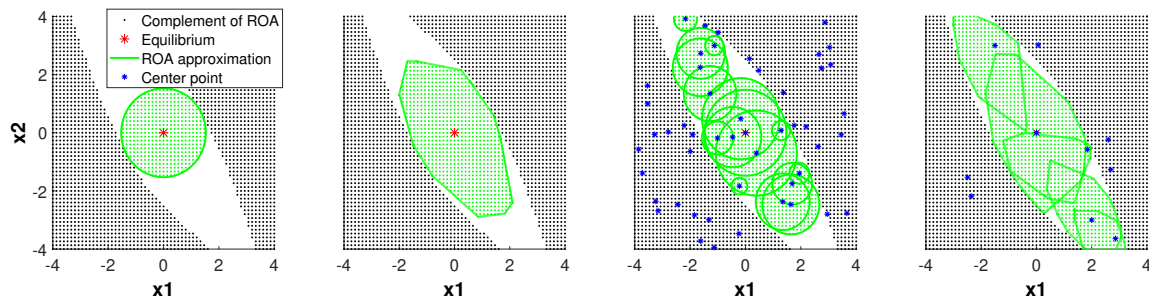


Fig. 3: The region of attraction approximations for one center point (left two) and multiple random center points (right two) in the sphere and polyhedron case, respectively.

Approximate method	# of counter examples	# of samples	# of steps simulated	Average # of steps per sample
1-center sphere approximation	14	7024	7935	1.39
1-center polyhedron approximation	94	23130	28127	1.22
50-center sphere approximations	191	17481	53756	3.07
10-center polyhedron approximations	370	46819	66399	1.41

TABLE I: Performance statistics for different configurations of our algorithm.

REFERENCES

- [1] H. K. Khalil, "Nonlinear systems; 3rd ed." 2002.
- [2] S. Willard, *General Topology*, ser. Addison-Wesley series in mathematics. Dover Publications, 2004.
- [3] Y. Li, S. Das, and N. Li, "Online optimal control with affine constraints," 2020.
- [4] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," 2020.
- [5] A. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," pp. 70–76, 2017.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado, "Learning lyapunov functions for hybrid systems," 2020.
- [12] 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.
- [13] 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.
- [14] R. Baier and M. Gerdt, "A computational method for non-convex reachable sets using optimal control," in *2009 European Control Conference, ECC 2009*, 08 2009.
- [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] 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.
- [24] E. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, ser. Texts in Applied Mathematics. Springer New York, 2013.
- [25] J. R. Munkres, *Topology / James R. Munkres.*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, Inc., 2000.
- [26] R. D. Driver, "Methods of an lyapunov and their application (vi zubov)," *SIAM Review*, vol. 7, no. 4, p. 570, 1965.
- [27] D. Haussler and E. Welzl, " ϵ -nets and simplex range queries," *Discrete & Computational Geometry*, vol. 2, pp. 127–151, 1987.
- [28] N. H. Mustafa, "Computing Optimal Epsilon-Nets Is as Easy as Finding an Unhit Set," in *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, 2019, pp. 87:1–87:12.
- [29] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," 2011.