Data-driven Practical Stabilization of Nonlinear Systems via Chain Policies: Sample Complexity and Incremental Learning

Roy Siegelmann and Enrique Mallada

***Abstract*— We propose a method for data-driven practical stabilization of nonlinear systems with provable guarantees, based on the novel concept of *Nonparametric Chain Policies (NCPs)*. The approach employs a normalized nearest-neighbor rule to assign, at each state, a finite-duration control signal derived from stored data, after which the process repeats. Unlike recent works that model the system as linear, polynomial, or polynomial fraction, we only require the system to be locally Lipschitz. Our analysis builds on the framework of Recurrent Lyapunov Functions (RLFs), which enable data-driven certification of stability using standard norm functions instead of requiring the explicit construction of a classical Lyapunov function. To extend this framework, we introduce the concept of Recurrent Control Lyapunov Functions (R-CLFs), which can certify the existence of an NCP that practically stabilizes an arbitrarily small c-neighborhood of an equilibrium point. We also provide an explicit sample complexity guarantee of $\mathcal{O}\big((3/\rho)^d \log(R/c)\big)$ number of trajectories—where R is the domain radius, d the state dimension, and ρ a system-dependent constant. The proposed Chain Policies are nonparametric, thus allowing new verified data to be readily incorporated into the policy to either improve convergence rate or enlarge the certified region. Numerical experiments illustrate and validate these properties.**

I. Introduction

Data-driven control methods offer an alternative paradigm for synthesizing controllers directly from trajectory observations, potentially bypassing the need for accurate system models while reducing computational burden and conservativeness of classical control synthesis [1], [2]. Recent years have witnessed significant progress in data-driven control. These approaches, as well as their level of maturity, depend considerably on the underlying system properties. For *linear systems*, the field has substantially matured: LMI-based formulas [3] and convex programs [4], [5] can transform trajectories into stabilizing feedback controllers with robustness [3], [5]–[7], performance [4], [5], [8] and sample complexity guarantees [9], [10].

For *nonlinear systems*, several approaches have been proposed, with methods highly dependent on the implicit assumptions made on the nonlinear system class and the control synthesis methodology. One prolific line of works considers dynamics formed from dictionary-based hypothesis classes—e.g., using polynomials [11]–[13], fractions of polynomials [14], or general nonlinear functions [15]—and formulate semidefinite programs that render policies with a wide variety of guarantees, including contraction-based stability [16]–[18] or robustness [13]. Other methods employ general learning techniques to learn models or policies and leverage intrinsic system properties to provide different guarantees, e.g., Koopman operator methods that exploit spectral properties [19], sample complexity analysis for stochastic dynamics [20], and conformal prediction approaches for statistical robustness [21].

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

To address these challenges, we introduce the concept of *Nonparametric Chain Policies* (NCPs), a data-driven approach that requires only Lipschitz assumptions on the system dynamics while providing explicit sample complexity guarantees for practical stabilization. NCPs employ a normalized nearest-neighbor rule to assign finite-duration control signals from a stored library of verified trajectories, enabling direct use of data without parametric modeling or optimization re-solving when new data arrives. The library itself is updateable. Our theoretical guarantees build on the framework of Recurrent Lyapunov Functions [22], which we extend here for the control setting by introducing Recurrent Control Lyapunov Functions (R-CLFs, Section III)—a natural relaxation of standard Control Lyapunov Functions [23].

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

1) **Explicit sample complexity:** NCPs achieve practical exponential stabilization using $\mathcal{O}\big((3/\rho)^d \log(R/c)\big)$ sample trajectories, with transparent scaling in dimension d, radius R, precision c, and system-dependent parameter ρ.
2) **Incremental learning:** The nonparametric nature of NCPs allows for new verified data to be seamlessly incorporated to expand a certified region, or improve performance, without discarding previous guarantees or

Roy Siegelmann is with the Department of Biological Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA. Email: rsiege@mit.edu. Enrique Mallada is with the Department of Electrical and Computer Engineering, Johns Hopkins University, 3400 N Charles St., Baltimore, MD 21218, USA. Email: mallada@jhu.edu.

This work was supported by the NSF Global Centers program under Grant No. 2330450 and by the DOE Office of Science (ASCR) under Award No. 826565.

re-solving optimization problems.

3) **Performance-complexity trade-offs:** The framework explicitly controls the trade-off between sample requirements and performance through a parameter ρ that relates theoretical optimal performance and practical performance guaranteed by the NCP.

Organization. Section II introduces preliminaries. Section III presents Recurrent Control Lyapunov Functions (R-CLFs) and stability guarantees. Section IV defines Nonparametric Chain Policies and establishes sample complexity results. Section V demonstrates the approach on nonlinear benchmarks, and Section VI concludes.

Due to space limitations, several detailed proofs are omitted; complete proofs and additional technical results are available in the extended version of this paper on arXiv [24].

Notation. $\|\cdot\|$ denotes an arbitrary norm on $\mathbb{R}^n$. We denote by $\mathrm{cl}(S)$ the closure of a set S and $\mathrm{int}(S)$ the interior of S. Given $x \in \mathbb{R}^n$ and $r > 0$, we define the closed ball of radius r centered at x as $B_r(x) := \{y \in \mathbb{R}^n \mid \|y - x\| \le r\}$. For a scalar $a \in \mathbb{R}$, we write $[a]_+ := \max\{a, 0\}$. For a set $S \subseteq \mathbb{R}^n$ and a point $x \in \mathbb{R}^n$, the distance from x to S is defined as $\mathrm{d}(x, S) := \inf_{y \in S} \|y - x\|$, with signed distance defined as

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

II. Preliminaries

We consider a nonlinear control system:

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

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

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

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

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

More generally, for a sequence of control signals $u_n \in \mathcal{U}^{(0,\tau_n]}$, with $\tau_n > 0$, $\forall n \in \mathbb{N}$, we further use $u_{[n]} := u_0u_1 \dots u_n$, and $u_{[\infty]} = \lim_{n\to\infty} u_{[n]}$. In some occasions we slightly abuse notation by using u interchangeably to represent instantaneous inputs in U and signals in $\mathcal{U}^{(a,b]}$; the intended meaning will always be clear from context.

For an initial state $x \in \mathbb{R}^n$ and control signal $u \in \mathcal{U}^{(0,a]}$, we denote by $\phi(t, x, u)$ the solution of (1) for $t \in (0, a]$. We further assume the following regularity conditions for (1).

Assumption 1 (Forward Completeness)**.** *The solutions of the control system* (1) *are* ***forward complete****. Specifically, for each initial condition $x \in \mathbb{R}^n$ and every control signal $u \in \mathcal{U}$, the trajectory $\phi(t, x, u)$ exists and remains bounded for all $t \ge 0$.*

Assumption 2 (Uniform Lipschitz Continuity)**.** *The vector field $f(x, u)$ of system* (1) *is locally Lipschitz continuous in x, uniformly with respect to u. More precisely, for every compact set $S \subseteq \mathbb{R}^n$, there exists a constant $L_S \ge 0$ such that*

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

A. Practical Exponential Stabilizability

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

Definition 1 (Equilibrium Point)**.** *A point $x^* \in \mathbb{R}^n$ is an* ***equilibrium point*** *of system* (1) *if there exists a control input $u^* \in U$ such that $f(x^*, u^*) = 0$.*

Definition 2 ((Practical) Exponential Stabilizability)**.** *Let $S \subseteq \mathbb{R}^n$. The equilibrium x^* of system* (1) *is said to be:*

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

$$\|\phi(t, x, u) - x^*\| \le K e^{-\lambda t} \|x - x^*\|, \quad \forall t \ge 0; \tag{2}$$

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

$$\|\phi(t, x, u) - x^*\| \le K e^{-\lambda t} \|x - x^*\| + c, \quad \forall t \ge 0, \tag{3}$$

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

It is well known from the topological entropy literature that it is impossible to exponentially stabilize a system, i.e., achieve (2), using a finite number of control signals [25]. We will therefore aim to enforce the weaker notion of practical exponential stability, i.e., (3), which follows the terminology of [26], [27].

B. Recurrent Lyapunov Functions

To provide guarantees for our data–driven stabilization framework, we build on the theory of *Recurrent Lyapunov Functions (RLFs)* [22], [28]. Unlike classical Lyapunov functions, which require a decrease all along trajectories, RLFs only require a decrease at a sequence of *bounded times*. This relaxation broadens the class of certificates available, while still ensuring exponential stability.

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

Definition 3 (Containment Times)**.** *Given a set $S \subset \mathbb{R}^n$, an initial state $x \in \mathbb{R}^n$, and an input $u \in \mathcal{U}$, the set of* containment times *is*

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

For constants $a, b > 0$ we define

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

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

We now recall the definition of an RLF in the *autonomous* case, where the trajectory is uniquely determined by the initial condition. In this case we denote by $\phi(t, x)$ the system's flow.

Definition 4 (Recurrent Lyapunov Function)**.** *Let $S \subset \mathbb{R}^n$ be a compact set with $x^* \in \mathrm{int}(S)$. A continuous function*

$V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ *is called a* Recurrent Lyapunov Function (RLF) *over* S *with rate* $\alpha > 0$ *and horizon* $\tau > 0$ *if*

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

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

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

Definition 5 (Reachable Tube). *For the control system* (1)*, a constant* $\tau > 0$*, and a set* $S \subset \mathbb{R}^n$*, we denote the* τ*-reachable tube from* S *within* τ *units of time by*

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

III. Recurrent Control Lyapunov Functions

As mentioned above, our guarantees rely on the theory of Recurrent Lyapunov Functions (RLFs) from [22]. In this section, we extend this notion to the control setting, introducing Recurrent Control Lyapunov Functions (R-CLFs), and illustrate how they can be used to certify practical stabilizability. Though RLFs have been shown to to certify stability, asymptotic stability and exponential stability [22], our focus here is on practical exponential stability and thus we will use the following definition.

Definition 6 (Recurrent Control Lyapunov Function (R-CLF)). *Consider the control system* (1) *with equilibrium* $x^* \in \mathbb{R}^n$*. Let* $S \subseteq \mathbb{R}^n$ *be a set satisfying* $x^* \in \mathrm{int}(S)$*. A continuous function* $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ *is a* ***Recurrent Control Lyapunov Function (R-CLF)*** *over* S *if the following conditions hold:*

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

$$a_1\|x - x^*\| \leq V(x) \leq a_2\|x - x^*\|, \quad \forall x \in S. \tag{4}$$

(ii) ***Control*** α***-Exponential*** (τ,δ)***-Recurrence****: There exist constants* $\tau, \alpha > 0$ *and* $\delta \geq 0$ *such that for every* $x \in S$*, there exists* $u \in \mathcal{U}^{(0,\tau]}$ *satisfying*

$$\min_{t \in T_S(x,u;\tau)} e^{\alpha t}(V(\phi(t,x,u)) - \delta) \leq [V(x) - \delta]_+ . \tag{5}$$

Lemma 1 characterizes the long term behavior of the control system (1) under the controls $u \in \mathcal{U}$ built upon concatenation of controls satisfying property *(ii)* of Definition 6.

Lemma 1 (Characterization of R-CLF). *Let assumptions 1 and 2 hold. Consider an equilibrium* x^* *of* (1) *and a* compact *set* S *satisfying* $x^* \in \mathrm{int}(S)$*. A function* $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ *satisfying* (4) *is a Recurrent Control Lyapunov Function (R-CLF) over* S if and only if *there exists parameters* $\alpha, \tau > 0$ *and* $\delta \geq 0$ *such that for any* $x \in S$ *there is a sequence* $\{t_n\}_{n \in \mathbb{N}}$ *and* $u \in \mathcal{U}$ *satisfying the following conditions:*

$$\lim_{n \to +\infty} t_n = +\infty, \ \textit{with} \quad t_{n+1} - t_n \in (0,\tau], \tag{6a}$$

$$\phi(t_n, x, u) \in S\,, \quad \textit{and} \tag{6b}$$

$$V(\phi(t_n,x,u)) - \delta \leq \begin{cases} e^{-\alpha t_n}(V(x) - \delta)\,, & n \leq \bar{n}\,, \\ 0\,, & \textit{o.w.}\,, \end{cases} \tag{6c}$$

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

Proof Sketch of Lemma 1. **Necessity** ($\Rightarrow$)**:** Assume V is an R-CLF over S. Then, given $x \in S$, one can use (5) to build a sequence $\{x_n\}_{n \geq 0}$ and $\{u_n \in \mathcal{U}^{(0,\tau_n]}\}_{n \geq 0}$ satisfying, $t_0 = 0$, $x_0 = x \in S$, and, for all $n \geq 0$,

$$t_{n+1} - t_n := \tau_n \in (0,\tau], \quad x_{n+1} := \phi(\tau_n, x_n, u_n) \in S,$$

where τ_n is the maximum t that minimizes (5) when $x = x_n$. Once $V(x_n) \leq \delta$, since $x_n \in S$, (5) ensures that $V(x_{n+1}) \leq \delta$. We define $\bar{n}$ to be that last n for which $V(x_n) > \delta$. Then for all $n \leq \bar{n}$,

$$V(x_n) - \delta \leq e^{-\alpha(t_n - t_{n-1})}(V(x_{n-1}) - \delta) \leq e^{-\alpha t_n}(V(x_0) - \delta).$$

Defining $u = \lim_{n \to \infty} u_{[n]}$, leads to $x_n = \phi(t_n, x, u)$ and thus showing (6b) and (6c). Finally, since $\tau_n \in (0,\tau]$ for all n, the sequence $\{t_n\}$ is strictly increasing. If t_n converged to a finite limit t^*, then by continuity of ϕ and compactness of S we would have $x_n \to x^* := \phi(t^*, x, u) \in S$. Applying the above inequality for $x^* \in S$ would produce,

$$e^{\alpha(t^* - t_n)}(V(x^*) - \delta) \leq V(x_n) - \delta, \quad \forall n \geq 0.$$

Thus, for n large enough, $\tau_n < \tau_n' := t^* - t_n \leq \tau$, making τ_n' a valid minizier of (5), contradicting the maximality of τ_n. Therefore $t_n \to \infty$, establishing (6a).

Sufficiency ($\Leftarrow$)**:** To show that V is a R-CLF, it is sufficient for any $x \in S$ to restrict the corresponding $u \in \mathcal{U}$ that satisfies (6) to the interval $(0,\tau]$ and choosing t_1 from the sequence (6) to show (5). ☐

We will leverage Lemma 1 to prove (practical) exponential convergence of trajectories. To that end, we need to bound how much a trajectory can travel between instances of exponential convergence (6c). The following lemma provides a mechanism to obtain such bounds. The proof is based on Grönwall's Lemma [29, Lemma A.1], and can be found in [30].

Lemma 2 (Containment Lemma). *Let Assumption 2 hold. Consider a compact set* $S \subset \mathbb{R}^n$ *and a constant* $\tau > 0$*. Then, for any* $x \in S, u \in \mathcal{U}$ *the following holds:*

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

where $L := L_{\mathcal{R}^\tau(S)}$*,* $F_S = \max_{u \in U} \max_{x \in S} \|f(x,u)\|$.

We are now ready to show that R-CLFs from Definition 6 guarantee exponential stabilizability.

Theorem 1 (R-CLF Implies (Practical) Exponential Stabilizability). *Consider the control system* (1) *with equilibrium* $x^* \in \mathbb{R}^n$*, and let* S *be a set satisfying* $S \subseteq \mathbb{R}^n$ *and* $x^* \in \mathrm{int}(S)$*. Let Assumption 1 and Assumption 2 hold, and* $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ *be a Recurrent Control Lyapunov Function over* S*, with constants* $\alpha, \tau > 0$*,* $\delta \geq 0$*, and linear bound constants* $a_1, a_2 > 0$ *from* (4)*.*

Then, the equilibrium x^* *is (practically) exponentially stabilizable on* S *(when* $\delta > 0$*). In particular, for every initial condition* $x \in S$*, a control signal* $u \in \mathcal{U}$ *satisfying* (6) *for* *some sequence of times* $\{t_n\}_{n\in\mathbb{N}}$ *ensures:*

$$\|\phi(t,x,u)-x^*\| \le Ke^{-\lambda t}\|x-x^*\|+c, \quad \forall t\ge 0, \tag{7}$$

where $\lambda{:=}\alpha$, $K{:=}\frac{a_2}{a_1}e^{\alpha\tau}(1{+}L\tau e^{L\tau})$, *and* $c{:=}\frac{\delta}{a_1}(1{+}L\tau e^{L\tau})$, *with* $L{:=}L_{R^\tau(S)}$.

Proof Sketch of Theorem 1. Fix $x\in S$. By Lemma 1, there exist $u\in\mathcal{U}$ and times $\{t_n\}$ with $x_n := \phi(t_n,x,u)\in S$ satisfying

$$V(x_n)\le e^{-\alpha t_n}(V(x)-\delta)+\delta$$

for all $n\le\bar{n}$, where $\bar{n}$ denotes the last index with $V(x_n)\ge\delta$. Using the bounds on V gives

$$r_n := \|x_n-x^*\| \le \tfrac{a_2}{a_1}e^{-\alpha t_n}\|x-x^*\|+\tfrac{\delta}{a_1}.$$

For any $t\in(t_n,t_{n+1}]$, Lemma 2 yields

$$\|\phi(t,x,u)-x^*\|\le(1+L\tau e^{L\tau})r_n.$$

Combining these estimates and using $t\le t_{n+1}\le t_n+\tau$ gives

$$\|\phi(t,x,u)-x^*\|\le Ke^{-\alpha t}\|x-x^*\|+c, \quad t\in(t_n,t_{n+1}],$$

with $K=\frac{a_2}{a_1}e^{\alpha\tau}(1+L\tau e^{L\tau})$, $c=\frac{\delta}{a_1}(1+L\tau e^{L\tau})$.

If $\bar{n}=\infty$ we are done. Otherwise, $V(x_n)\le\delta$ for $n>\bar{n}$, so $\|x_n-x^*\|\le\delta/a_1$, and the same bound implies $\|\phi(t,x,u)-x^*\|\le c$ for $t>t_{\bar{n}}$. The desired result, i.e., (7), follows. □

Theorem 1 states that the existence of an R-CLF implies that x^* can be made practically exponentially stable. At the core of its proof is the fact that one can find a function V that satisfies the recurrent condition (5). A key observation of [28], is that condition (5) can be met by a norm, provided τ and α are properly chosen (c.f. [28, Theorem 6]). The caveat is that in order to make R-CLFs practically useful, one would need to store, for each $x\in S$, a suitable $u:(0,\tau]\to U$ that ensures (5). In the next section we demonstrate a powerful property of R-CLFs, that when $\delta>0$, only a finite number of such signals are needed.

IV. Non-Parametric Chain Policies

In the previous section, we introduced Recurrent Control Lyapunov Functions (R-CLFs) to characterize exponential stabilizability via carefully selected control signals. Here, we propose *nonparametric chain policies*, a systematic data-driven approach for generating these stabilizing signals. The method aligns closely with recent developments in topological entropy regarding the minimal complexity required to accomplish control tasks (see, e.g., [25], [27], [31]). A distinctive feature is that control signals need not be generated online; instead, they are explicitly stored in a *control alphabet* [30].

Definition 7 (Control Alphabet)**.** *A **control alphabet** is a finite collection of control signals of size* $M>0$

$$\mathcal{A}:=\{v_i:(0,\tau_i]\to U\}_{i=0}^{M},$$

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

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

Definition 8 (Assignment Set)**.** *An **assignment set** is a finite collection of verification triplets*

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

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

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

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

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

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

where $r_{\mathcal{K}}(x) := \min_{(x_i,r_i,v_i)\in\mathcal{K}}\frac{\|x-x_i\|}{r_i}$, and $\iota_{\mathcal{K}}(x)=0$ corresponds to selecting the default control v_0.

Remark 1. *We designate* $v_0\in\mathcal{A}$ *as the* default control*. Unless otherwise stated, we take* $v_0(t)=u^*\in U$, $\forall t\in[0,\tau_0)$*, where* u^* *is any equilibrium control satisfying Definition 1.*

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

Definition 9 (Nonparametric Chain Policy)**.** *Given an assignment set* $\mathcal{K}$ *and default control* v_0*, the **nonparametric chain policy** is given by the map* $\pi_{\mathcal{K}}:\mathbb{R}^n\to\mathcal{A}$*:*

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

Remark 2. *The policy* $\pi_{\mathcal{K}}$ *induces an infinite-horizon control signal* $u_{\mathcal{K},x}\in\mathcal{U}$ *through concatenation. Starting with* $x_0=x$ *and the empty signal* $u_{[0]}=\emptyset$*, for each* $n\ge 0$ *define*

$$u_{[n+1]}=u_{[n]}\,v_{\iota_{\mathcal{K}}(x_n)},\quad x_{n+1}=\phi\big(\tau_{\iota_{\mathcal{K}}(x_n)},\,x_n,\,v_{\iota_{\mathcal{K}}(x_n)}\big). \tag{8}$$

The resulting control is then: $u_{\mathcal{K},x}:=\lim_{n\to\infty}u_{[n]}$.

A. Convergence Guarantees of NCPs

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

Theorem 2 (Practical Exponential Stabilization via Chain Policies)**.** *Consider an equilibrium point* $x^*\in\mathbb{R}^n$ *of (1), and let* $S\subseteq\mathbb{R}^n$ *be a set with* $x^*\in\operatorname{int}(S)$*. Let* $\pi_{\mathcal{K}}$ *denote a nonparametric chain policy associated with the assignment set* $\mathcal{K}=\{(x_i,r_i,v_i)\}_{i=1}^{N}$ *and a default control* $v_0\in\mathcal{A}$*, define*

$\tau := \max\{\tau_0, \tau_1, \dots, \tau_N\}$*, and let* $L := L_{\mathcal{R}^\tau(S)}$*. Suppose the following hold:*

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

$$B_\varepsilon(x^*) \subset B_{\varepsilon(1+L\tau e^{L\tau})}(x^*) \subset \mathrm{int}(S), \tag{9a}$$

$$\mathrm{cl}(S \setminus B_\varepsilon(x^*)) \subseteq \mathrm{Supp}(\mathcal{K}). \tag{9b}$$

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

$$e^{\alpha\tau_i}\big(\|\phi(\tau_i, x_i, v_i) - x^*\| + r_i e^{L\tau_i}\big) \le \|x_i - x^*\| - r_i, \tag{10a}$$

$$\mathrm{sd}(\phi(\tau_i, x_i, v_i), S) + r_i e^{L\tau_i} \le 0, \tag{10b}$$

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

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

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

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

Let $x_0 = x \in S$, $t_0 = 0$, and $u_{[0]} = \emptyset$. Define the sequences $\{x_n\}$, $\{t_n\}$, and $u_{[n]}$ according to (8), i.e.,

$$x_{n+1} = \phi\big(\tau_{\iota_{\mathcal{K}}(x_n)}, x_n, v_{\iota_{\mathcal{K}}(x_n)}\big), \qquad u_{[n+1]} = u_{[n]} v_{\iota_{\mathcal{K}}(x_n)},$$
$$u_{\mathcal{K},x} = \lim_{n\to\infty} u_{[n]}, \qquad t_{n+1} := t_n + \tau_{\iota_{\mathcal{K}}(x_n)}, \quad \forall n \ge 0.$$

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

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

so condition (6a) holds. Moreover, by induction one shows that for all $n \ge 1$, $\phi(t_n, x, u_{\mathcal{K},x}) = x_n = \phi\big(\tau_{\iota_{\mathcal{K}}(x_{n-1})}, x_{n-1}, v_{\iota_{\mathcal{K}}(x_{n-1})}\big)$.

We claim that if $x_n \in S$ then $x_{n+1} \in S$. Suppose first that $x_n \in S \setminus B_\varepsilon(x^*)$. By the covering condition (9b), there exists $(x_i, r_i, v_i) \in \mathcal{K}$ such that $x_n \in B_{r_i}(x_i)$ and the verification condition (10) holds. In particular, by (10b),

$$\begin{aligned}\mathrm{sd}(x_{n+1}, S) &= \mathrm{sd}(\phi(\tau_i, x_n, v_i), S)\\ &\le \mathrm{sd}(\phi(\tau_i, x_i, v_i), S) + r_i e^{L\tau_i} \le 0,\end{aligned}$$

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

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

$$\begin{aligned}\|x_{n+1} - x^*\| &= \|\phi(\tau_0, x_n, v_0) - x^*\|\\ &\le \varepsilon + d(\phi(\tau_0, x_n, v_0), B_\varepsilon(x^*))\\ &\le \varepsilon + F_{B_\varepsilon(x^*)}\tau e^{L\tau_0} \le \varepsilon(1 + L\tau_0 e^{L\tau_0}),\end{aligned} \tag{11}$$

so by (9a) and $\tau_0 \le \tau$ we conclude $x_{n+1} \in S$. Thus, $x_n \in S$ implies $x_{n+1} \in S$, i.e., condition (6b) holds.

Verification of (6c). Let $\delta := \varepsilon(1 + L\tau e^{L\tau})$ and $\bar{n} := \inf\{n : \|x_n - x^*\| \le \delta\}$. If $\iota_{\mathcal{K}}(x_n) = 0$ and $\|x_n - x^*\| \le \varepsilon \le \delta$, then by (11) we have $\|x_{n+1} - x^*\| \le \delta$.

If $\iota_{\mathcal{K}}(x_n) = i \ne 0$, then from (10a) and $x_n \in B_{r_i}(x_i)$,

$$\begin{aligned}e^{\alpha(t_{n+1}-t_n)}\|x_{n+1} - x^*\| &\le e^{\alpha\tau_i}\big(\|\phi(\tau_i, x_i, v_i) - x^*\| + r_i e^{L\tau_i}\big)\\ &\le \|x_i - x^*\| - r_i \;\le\; \|x_n - x^*\|,\end{aligned}$$

which implies

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

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

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

Hence for $n \ge \bar{n}$, $\|x_n - x^*\| \le \delta$, while for $n < \bar{n}$ the excess above δ decays exponentially. This verifies condition (6c).

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

B. Existence and Sample Complexity of NCPs

Theorem 2 establishes that nonparametric chain policies can guarantee practical exponential stability of a region around an equilibrium point that is appropriately covered by data points from $\mathcal{K}$. However, it is a priori not clear how many data points are needed to construct such policy, or even whether such a policy exists. The next result provides conditions for existence of Chain Policies as well as a bound on the sample complexity of such policies, i.e., the sizes of the assignment set $\mathcal{K}$ and alphabet $\mathcal{A}$ required to construct such policy.

Theorem 3 (Existence and Sample Complexity of Chain Policies)**.** *Consider the control system* (1) *with equilibrium* $x^* \in \mathbb{R}^n$*. Let* $S = B_R(x^*)$ *with* $R > 0$ *such that* x^* *is* λ*-exponentially stabilizable on* S *with gain* $K > 0$*, and choose* ε *s.t.* $R > \varepsilon > 0$*. Fix* $\alpha \in (0, \lambda)$ *and choose* $\tau > \frac{\ln K}{\lambda - \alpha}$*. Assign*

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

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

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

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

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

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

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

Proof Sketch of Theorem 3

Since x^* is λ-exponentially stabilizable on S, for each grid center $x_i \in S$ there exists a control signal v_i such that

$$\|\phi(t, x_i, v_i) - x^*\| \le Ke^{-\lambda t}\|x_i - x^*\|, \qquad t \ge 0.$$

Fix $\alpha \in (0, \lambda)$ and $\tau > \frac{\ln K}{\lambda - \alpha}$, and let $L := L_{\mathcal{R}^\tau(S)}$. Define

$$\rho := \frac{1 - Ke^{-(\lambda-\alpha)\tau}}{1 + e^{(L+\alpha)\tau}}, \qquad \bar{r}_i := \rho \|x_i - x^*\|.$$

Using the decay exponential estimate at $t = \tau$ gives

$$e^{\alpha\tau}\big(\|\phi(\tau, x_i, v_i) - x^*\| + \bar{r}_i e^{L\tau}\big) \leq \|x_i - x^*\| - \bar{r}_i,$$

so $(x_i, \bar{r}_i)$ satisfies the verification condition (10a). Hence any radius $0 < r_i \leq \bar{r}_i$ also satisfies the verification condition.

It therefore suffices to cover $B_R(x^*) \setminus B_\varepsilon(x^*)$ with balls $B_{r_i}(x_i)$ satisfying $r_i \leq \bar{r}_i$. To this end, partition the region into annuli $A_l := B_{R_l}(x^*) \setminus B_{R_{l-1}}(x^*)$, $R_l = (2\varepsilon)3^{l-1}$. The number of annuli satisfies $n = \lceil \log_3\big(\frac{R}{\varepsilon}\big) \rceil$.

Each annulus A_l can be covered by $3^d - 1$ cubes of side R_l. For any center $x_i \in A_l$ we have $\|x_i - x^*\| \geq 3^{l-1}\varepsilon$, and admissibility requires $r_i \leq \bar{r}_i = \rho\|x_i - x^*\|$. Starting from cubes of radius $R_l/2$, we refine each cube by splitting it into 3^d subcubes until

$$r_i = \frac{R_l}{2}\, 3^{-m} \leq \rho\, 3^{l-1}\varepsilon \leq \rho ||x_i - x^*||,$$

which requires $m = \left\lceil \log_3\Big(\frac{1}{\rho}\Big) \right\rceil$. Thus each initial cube produces at most $(3^d - 1)3^{dm}$ centers, yielding the bound $(3/\rho)^d$ per annulus. Multiplying $n \times m$ gives claimed bound. □

Remark 3 (Performance–Complexity Trade-off)**.** *The parameter ρ captures a fundamental trade-off: choosing α close to the intrinsic rate λ improves performance but drives $\rho \to 0$, inflating sample complexity as $O((3/\rho)^d)$. Smaller α reduces data requirements at the cost of slower convergence.*

Remark 4 (Scalability)**.** *The sample complexity in Theorem 3 scales exponentially with the state dimension d. This dependence arises from the Lipschitz-based analysis used to obtain uniform guarantees over the state space. While such scaling is common without additional structural assumptions on $f(x, u)$, we do not establish a matching lower bound. Whether improved dependence on d can be achieved under additional assumptions remains an interesting direction for future work, along with Robustness to model uncertainty or disturbances.*

C. Incremental Learning of NCPs

The performance–complexity trade-off in Remark 3, together with the existence and sample complexity result of Theorem 3, suggests a practical methodology for progressively improving performance. By sampling trajectories more finely and refining the covering set $\mathcal{K}$, one can construct NCPs that certify larger rates α (by reducing the effective radius r), thereby narrowing the gap to the optimal rate λ. Performance thus improves as the assignment set is enriched with additional verified points, enabling incremental learning.

Beyond improving rates, NCPs also support incremental *expansion of the certified region*. The next result formalizes this: previously verified assignments can be combined with new ones to enlarge the domain over which stability is guaranteed.

Theorem 4 (Incremental Learning of $\mathcal{K}$)**.** *Consider an equilibrium point $x^* \in \mathbb{R}^n$ of* (1) *and a set $S \subseteq \mathbb{R}^n$ satisfying $x^* \in \mathrm{int}(S)$. Let $\pi_\mathcal{K}$ be a nonparametric policy with assignment set $\mathcal{K} = \{(x_i, r_i, v_i)\}_{i=1}^N$ and default control $v_0 \in \mathcal{A}$ satisfying properties* (i)–(iii) *of Theorem 2 with parameters α, δ, τ, L, ε. Take $x_j \in \mathbb{R}^n \setminus S$, $r_j > 0$, and $v_j \in \mathcal{U}^{[0,\tau_j)}$ s.t. $B_{r_j}(x_j) \cup S = \emptyset$. Define the enlarged set $S' := S \cup B_{r_j}(x_j)$, and let $L_j = L_{\mathcal{R}^{\tau_j}(B_{r_j}(x_j))}$, and $L' = \max\{L_j, L\}$.*

Whenever the following two conditions are satisfied:

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

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

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

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

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

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

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

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

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

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

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

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

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

Proof. The proof is omitted due to page limits. and can be found in [24]. □

V. Numerical Experiments

Using the sufficiency of NCPs of Theorem 2 along with the grid construction of Theorem 3, we next introduce Algorithm 1 to design NCPs to stabilize a region. With this algorithm, we present case studies in different classic control stabilization problems, which demonstrate useful features of NCPs.

A. Unicycle

Consider the unicycle model with state (x, y, θ) and controls velocity v and angular velocity ω:

$$\dot{x} = v\cos\theta, \quad \dot{y} = v\sin\theta, \quad \dot{\theta} = \omega.$$

With $v \in [0, 1]$ and $\omega \in [-1, 1]$, we derive NCPs for two different norms, $V_1 = \max\{|x|, |y|, |\theta|\}$ and $V_2 = \sqrt{x^2 + y^2 + 0.01\theta^2}$. For both choices, the method stabilizes the entire region $(x, y) \in [-20, 20]^2$, $\theta \in (-\pi, \pi]$ (Fig. 1).

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

Algorithm 1 Constructing Nonparametric Chain Policy (NCP)

Require: Region $S \subset \mathbb{R}^n$, target rate $\alpha > 0$, maximum horizon $\tau_{\max} > 0$, tolerance $\varepsilon > 0$

```
Estimate L := L_{R_{τmax}(S)}
Construct an initial covering G = {(x_i, r_i)} of S \ B_ε(x*)
for each (x_i, r_i) ∈ G do
    Compute candidate control v_i over [0, τmax] (for example, by MPPI [32], [33])
    if there exists τ_i ≤ τmax satisfying (10a) then
        Store (x_i, r_i, v_i, τ_i)
    else
        Subdivide B_{r_i}(x_i) into 3^d sub-balls of radius r_i/3
        Add sub-balls to G, removing the original (x_i, r_i)
    end if
end for
Retain only assignments satisfying (10b)
Attempt bootstrapping (Theorem 4) to unverified (x_i, r_i) ∈ G, storing new successful (x_i, r_i, v_i, τ_i)
return assignment set K
```

B. Inverted Pendulum

The inverted pendulum consists of a mass m at the end of a rigid rod of length l pivoting about a fixed point. Let θ denote the angle from the vertical, with $\theta = 0$ corresponding to the inverted equilibrium. The dynamics are

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

where g is the gravitational acceleration and $u(t)$ is the control torque.

Figure 3 demonstrates the refinement capabilities of NCPs, such that by adding data (simulated by splitting all balls once), the rate of convergence achieved is significantly increased.

VI. Conclusions

In this work we proposed a method for data-driven (practical) stabilization of nonlinear systems using nonparametric Chain Policies. The approach leverages a normalized nearest-neighbor rule to assign, at each state, a finite-duration control signal, after which the process repeats. The method is grounded in the notion of Recurrent Lyapunov Functions (RLFs) as well as their control extension Control-RLFs, which enable certification of stability using standard norm function.

Our analysis establishes that:

1) NCP Policies achieve practical exponential convergence to a c-neighborhood with sample complexity scaling as $O\big((3/\rho)^d \log(R/c)\big)$, with radius R and precision c.
2) The framework supports incremental growth: new assignments can be added to expand the verified region while preserving previously established guarantees.
3) Controller refinement is monotone: more data only improves convergence rates and enlarges certified region.

These results position Chain Policies as a flexible, data-driven approach to certified stabilization, offering rigorous guarantees together with the ability to expand incrementally as new data becomes available.

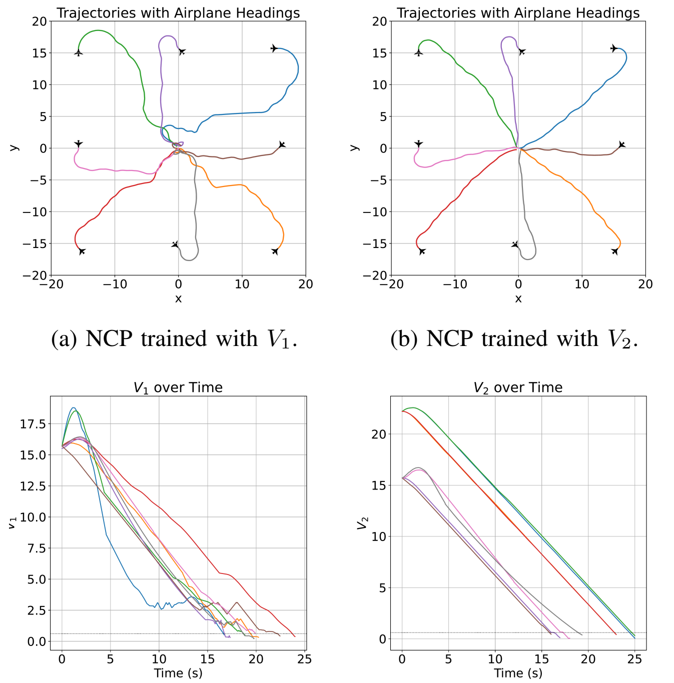

(a) NCP trained with V_1.

(b) NCP trained with V_2.

(c) $V_1(\phi(t,x,u))$.

(d) $V_2(\phi(t,x,u))$.

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

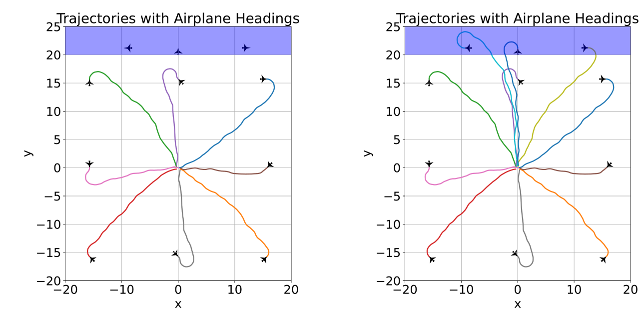

(a) Before learning

(b) After learning.

Fig. 2: **Incremental Learning of Unicycle Policy.** Extending the state space from the previously learned region in the y-direction. Subfigure (a) contains the phase plot before learning, while subfigure (b) contains the phase plot after. The new region is learned without forgetting, such that parts of the trajectory in the old region use previously designed controls.

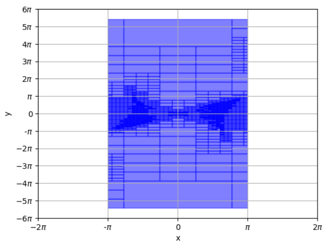

(a) Base NCP verification.

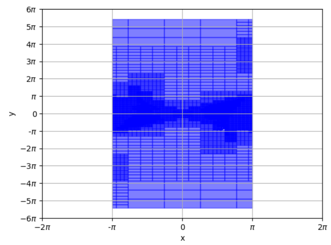

(b) NCP refinement.

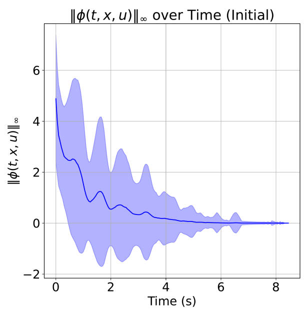

(c) $\|\phi(t,x,u))\|_\infty$, initial.

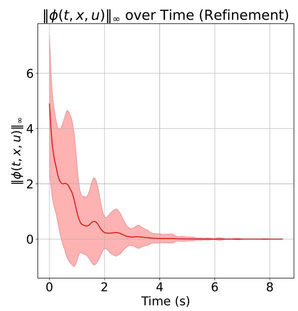

(d) $\|\phi(t,x,u))\|_\infty$, refinement.

Fig. 3: **Additional Data Refinement Facilitates Improved NCP Performance.** Plot (a) contains the balls used to verify the region $(\theta, \dot{\theta}) \in (-\pi, \pi] \times [-5\pi, 5\pi]$ for the inverted pendulum. Plot (b) is a refinement of plot (a), wherein all balls were split once more and re-verified. The minimum verified rate of convergence for trajectories α goes from 0.003 to 0.0145, and the average verified α goes from 1.815 to 3.149. Plot (c) demonstrates the average norm over time of 400 sample trajectories under each schema. We have $\tau_{\max} = 1.5$, $\varepsilon = 0.01$, $L = 5$, and $c = \varepsilon(1 + L\tau_{\max}e^{L\tau_{\max}}) \simeq 0.072$.

References

[1] K. Zhou, J. Doyle, and K. Glover, "Robust and optimal control," *Control Engineering Practice*, vol. 4, no. 8, pp. 1189–1190, 1996.

[2] E. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, ser. Texts in Applied Mathematics. Springer New York, 2013.

[3] C. De Persis and P. Tesi, "Formulas for data-driven control: Stabilization, optimality, and robustness," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 909–924, 2019.

[4] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the deepc," in *2019 18th European control conference (ECC)*. IEEE, 2019, pp. 307–312.

[5] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE transactions on automatic control*, vol. 66, no. 4, pp. 1702–1717, 2020.

[6] J. Berberich, A. Koch, C. W. Scherer, and F. Allgöwer, "Robust data-driven state-feedback design," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 1532–1538.

[7] H. J. Van Waarde, M. K. Camlibel, and M. Mesbahi, "From noisy data to feedback controllers: Nonconservative design via a matrix s-lemma," *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 162–175, 2020.

[8] C. De Persis and P. Tesi, "Low-complexity learning of linear quadratic regulators from noisy data," *Automatica*, vol. 128, p. 109548, 2021.

[9] X. Chen and E. Hazan, "Black-box control for linear dynamical systems," in *Conference on Learning Theory*. PMLR, 2021, pp. 1114–1143.

[10] S. W. Werner and B. Peherstorfer, "On the sample complexity of stabilizing linear dynamical systems from data," *Foundations of Computational Mathematics*, vol. 24, no. 3, pp. 955–987, 2024.

[11] T. Dai and M. Sznaier, "A semi-algebraic optimization approach to data-driven control of continuous-time nonlinear systems," *IEEE Control Systems Letters*, vol. 5, no. 2, pp. 487–492, 2020.

[12] M. Guo, C. De Persis, and P. Tesi, "Data-driven stabilization of nonlinear polynomial systems with noisy data," *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 4210–4217, 2021.

[13] H. El-Kebir and M. Ornik, "Sum-of-squares data-driven robustly stabilizing and contracting controller synthesis for polynomial nonlinear systems," *arXiv preprint arXiv:2503.07438*, 2025.

[14] R. Strässer, J. Berberich, and F. Allgöwer, "Data-driven control of nonlinear systems: Beyond polynomial dynamics," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 4344–4351.

[15] N. Monshizadeh, C. De Persis, and P. Tesi, "A versatile framework for data-driven control of nonlinear systems," *IEEE Transactions on Automatic Control*, 2025.

[16] A. Oliveira, J. Zheng, and M. Sznaier, "Convex data-driven contraction with riemannian metrics," *IEEE Control Systems Letters*, 2025.

[17] Z. Hu, C. De Persis, and P. Tesi, "Enforcing contraction via data," *IEEE Transactions on Automatic Control*, 2025.

[18] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview," *Annual Reviews in Control*, vol. 52, pp. 135–169, 2021.

[19] B. Huang, X. Ma, and U. Vaidya, "Data-driven nonlinear stabilization using koopman operator," in *The Koopman Operator in Systems and Control: Concepts, Methodologies, and Applications*. Springer, 2020, pp. 313–334.

[20] Y. Chen and U. Vaidya, "Sample complexity for nonlinear stochastic dynamics," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 3526–3531.

[21] T.-W. Hsu and H. Tsukamoto, "Statistical guarantees in data-driven nonlinear control: Conformal robustness for stability and safety," *IEEE Control Systems Letters*, 2025.

[22] 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 *62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 12 2023, pp. 6665–6672.

[23] E. D. Sontag, "A lyapunov-like characterization of asymptotic controllability," *SIAM journal on control and optimization*, vol. 21, no. 3, pp. 462–471, 1983.

[24] R. Siegelmann and E. Mallada, "Data-driven practical stabilization of nonlinear systems via chain policies: Sample complexity and incremental learning," *arXiv preprint arXiv:2510.03982*, 2025.

[25] F. Colonius, "Minimal bit rates and entropy for exponential stabilization," *SIAM Journal on Control and Optimization*, vol. 50, no. 5, pp. 2988–3010, 2012.

[26] B. Hamzi and A. J. Krener, "Practical stabilization of systems with a fold control bifurcation," in *New Trends in Nonlinear Dynamics and Control and Their Applications*. Springer, 2004, pp. 37–48.

[27] F. Colonius and B. Hamzi, "Entropy for practical stabilization," *SIAM Journal on Control and Optimization*, vol. 59, no. 3, pp. 2195–2222, 2021.

[28] R. Siegelmann, Y. Shen, F. Paganini, and E. Mallada, "Stability analysis and data-driven verification via recurrent lyapunov functions," 07 2025, submitted.

[29] H. K. Khalil, "Nonlinear systems; 3rd ed." 2002.

[30] H. Sibai and E. Mallada, "Recurrence of nonlinear control systems: Entropy, bit rates, and finite alphabet controllers," *Nonlinear Analysis: Hybrid Systems*, vol. 59, p. 101649, 2026.

[31] F. Colonius and C. Kawan, "Invariance entropy for control systems," *SIAM J. Control Optim.*, vol. 48, no. 3, pp. 1701–1721, 2009.

[32] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.

[33] ——, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.