

Neural Lyapunov Redesign

Arash Mehrjou

Max Planck Institute for Intelligent Systems & ETH Zürich

AMEHRJOU@ETHZ.CH

Mohammad Ghavamzadeh

Google Research

GHAVAMZA@GOOGLE.COM

Bernhard Schölkopf

Max Planck Institute for Intelligent Systems

BS@TUE.MPG.DE

Abstract

Learning controllers merely based on a performance metric has been proven effective in many physical and non-physical tasks in both control theory and reinforcement learning. However, in practice, the controller must guarantee some notion of safety to ensure that it does not harm either the agent or the environment. Stability is a crucial notion of safety, whose violation can certainly cause unsafe behaviors. Lyapunov functions are effective tools to assess stability in nonlinear dynamical systems. In this paper, we combine an improving Lyapunov function with automatic controller synthesis in an iterative fashion to obtain control policies with large safe regions. We propose a two-player collaborative algorithm that alternates between estimating a Lyapunov function and deriving a controller that gradually enlarges the stability region of the closed-loop system. We provide theoretical results on the class of systems that can be treated with the proposed algorithm and empirically evaluate the effectiveness of our method using an exemplary dynamical system.

Keywords: Lyapunov function, Controller Synthesis, Actor-Critic, Neural Networks

1. Introduction

Studying the stability region of autonomous systems and designing controllers (policies) to drive a non-autonomous system towards a target behavior are of fundamental importance in many disciplines, such as aviations ([Liao and Wang, 2004](#)), autonomous driving ([Wen-Xing and Li-Dong, 2018](#)), and robotics ([Pierson and Gashler, 2017](#)). An indisputable goal for a controller is to stabilize the system. Unlike the global nature of linear systems, stability is a local property in nonlinear systems. Knowledge of the stability region is essential in many applications, such as stability of power systems ([Xin et al., 2007](#)), design of associative memory in artificial neural networks ([Hopfield, 1994](#)), robotics ([Westervelt et al., 2003](#)), and biology ([Baer et al., 2006](#)).

Controllers that enhance the stability region of a system, also known as Region of Attraction (RoA), are highly desired as they make more clever use of the inherent nonlinear structure of the system. For example, an autonomous driving system will remain safe under more diverse and potentially harsh conditions ([Imani Masouleh and Limebeer, 2018](#)).

A great number of methods for designing a controller for nonlinear systems and determining their RoA ([Isidori, 2014](#); [Khalil and Grizzle, 2002](#)) have been proposed in the literature mainly based on Lyapunov's theory of stability ([Liapunoff, 1907](#)). However, the design of a stabilizing controller has been ad-hoc for every class of nonlinear systems. Our work is inspired by a classic method called *Lyapunov Redesign* in control theory that uses a given Lyapunov function to design a controller

such that the closed-loop system becomes stable when assessed by that Lyapunov function. Here, we relax the necessity to know the Lyapunov function by learning it together with the controller.

With the recent interest in data-driven control, machine learning tools have been employed to learn Lyapunov function for nonlinear systems (Richards et al., 2018). A given Lyapunov function is used in (Berkenkamp et al., 2016, 2017) to derive a controller with a safety guarantee. Chang et al. (2019) solve an optimization problem at every stage to find the states that violate the Lyapunov condition. Our work proposes improvements in various directions. We improve the Lyapunov learning algorithm of Richards et al. (2018) by a theoretically motivated method and show that *hyperbolicity* is a needed feature for nonlinear systems whose controllers are learned iteratively. Moreover, Berkenkamp et al. (2016, 2017) assume the Lyapunov function is given while our work co-learns the Lyapunov function and the controller. Chang et al. (2019)’s algorithm needs to be solved until the end that can be too costly as a global optimization problem has to be solved multiple times. However, our work is a growing method that improves the controller while the system is in action.

We bring together tools from control theory and machine learning to build an iterative algorithm that redesigns both controller and the Lyapunov function to enlarge the stability region. Our contributions are as follows: **1)** Improving the learning of the Lyapunov function, **2)** Interlacing the Lyapunov learning with the policy update to enlarge the RoA iteratively, and **3)** Providing theoretical results for the tractable class of systems and analysing the learning signal. Section 2 goes over the preliminary materials and Section 3 describe the problem. The proposed algorithm and its theoretical discussions are presented in Section 4. Finally, empirical evaluation comes in Section 5, followed by related work and conclusions in Section 6. To remain within the page limit, some proofs, theoretical discussions, and extended experimental results are provided in the Appendix, which can be found in the extended version of the paper <https://arxiv.org/pdf/2006.03947.pdf>. The developed software is open-sourced and is available at https://github.com/amehrjou/neural_lyapunov_redesign.

2. Preliminaries

Here we present the definitions, notations, and theoretical results that are used in later sections.

System: A discrete-time¹ time-invariant disturbance-free nonlinear dynamical system is described by

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \tag{1}$$

where $k \in \mathbb{Z}$ is the discrete time index, $\mathbf{x}_k \in \mathcal{X} \subseteq \mathbb{R}^d$ and $\mathbf{u}_k \in \mathcal{U} \subseteq \mathbb{R}^p$ are the state and control signals. We consider a fully observable regime, where the states are available to a feedback controller, i.e., $\mathbf{u}_k = \pi(\mathbf{x}_k)$, and π is the feedback law or policy. Hence, (1) can be written as a time-invariant autonomous (TIA) system $\mathbf{x}_{k+1} = f_\pi(\mathbf{x}_k)$ where f_π is the time-independent *dynamics* function. By assuming Lipschitz continuity for f and π , a unique solution to this system for every initial state exists that is captured by the so-called *flow* function $\Phi(\mathbf{x}, \cdot) : \mathbb{Z} \rightarrow \mathcal{X}$, with $\Phi(\mathbf{x}, 0) = \mathbf{x}$.

Sets: For a TIA system with dynamics function f , a state vector $\bar{\mathbf{x}}$ is called an *equilibrium point* if it is a fixed-point for f , i.e., $\bar{\mathbf{x}} = f(\bar{\mathbf{x}})$. A state vector is called a *regular point* if it is not an *equilibrium point*. Let $J_f(\mathbf{x})$ be the Jacobian of f at \mathbf{x} . If $J_f(\mathbf{x})$ has no eigenvalue with modulus one, \mathbf{x} is called a *hyperbolic equilibrium point*. A hyperbolic equilibrium point is asymptotically stable when the eigenvalues of its corresponding Jacobian have modulus less than one; otherwise, it is an *unstable*

1. Including discretized continuous-time systems

equilibrium point. A system whose all equilibrium points are hyperbolic is called a hyperbolic system. A set $M \subseteq \mathcal{X}$ is called an *invariant* set, if $f(M) = M$, i.e., every trajectory starting in M remains in M , for $k \in \mathbb{Z}$. A point $\mathbf{p} \in \mathbb{R}^d$ is said to be in the ω -limit set (or α -limit set) of \mathbf{x} , if for every $\epsilon > 0$ and $N > 0$ ($N < 0$), there exists a $k > N$ ($k < N$) such that $\|\mathbf{x}_k - \mathbf{p}\| < \epsilon$. The stable and unstable manifolds of $\bar{\mathbf{x}}$ are defined as the set of points whose ω -limit (α -limit) set is $\bar{\mathbf{x}}$ and denoted by $W^s(W^u)$. Both W^s and W^u are proved to be invariant sets (Palis and De Melo, 2012).

Stability: A fixed-point $\bar{\mathbf{x}}$ is said to be an asymptotically stable equilibrium, if $\lim_{k \rightarrow \infty} \mathbf{x}_k = \bar{\mathbf{x}}$. Nonlinear systems often have a local stability region (RoA) that is defined for the stable equilibrium $\bar{\mathbf{x}}$ as $\mathcal{R}^{\bar{\mathbf{x}}} = \{\mathbf{x} \in \mathcal{X} : \lim_{k \rightarrow \infty} \Phi(\mathbf{x}, k) = \bar{\mathbf{x}}\}$. Topologically speaking, when f is continuous, $\mathcal{R}^{\bar{\mathbf{x}}}$ is an open, invariant set (see Chiang and Alberto, 2015 for exact definitions). The stability boundary $\partial\mathcal{R}^{\bar{\mathbf{x}}}$ is a closed positively invariant set and is of dimension $n - 1$, if $\mathcal{R}^{\bar{\mathbf{x}}}$ is not dense in \mathbb{R}^n .

Lyapunov stability. Let f be a locally Lipschitz continuous dynamics function with an equilibrium point at the origin $\bar{\mathbf{x}} = \mathbf{0}$. Suppose there exists a locally Lipschitz continuous function $V = \mathcal{X} \rightarrow \mathbb{R}$ and a domain $\mathcal{D} \subseteq \mathcal{X}$, such that

$$V(\mathbf{0}) = 0 \quad \text{and} \quad V(\mathbf{x}) > 0 \quad \forall \mathbf{x} \in \mathcal{D} \setminus \{\mathbf{0}\} \quad (2)$$

$$\Delta V(\mathbf{x}) := V(f(\mathbf{x})) - V(\mathbf{x}) < 0 \quad \forall \mathbf{x} \in \mathcal{D} \setminus \{\mathbf{0}\} \quad (3)$$

Then, $\bar{\mathbf{x}}$ is asymptotically stable and V is a *Lyapunov function (lf)*. The domain \mathcal{D} in which (3) is satisfied is called the *Lyapunov decrease region*. It is easy to show that every level set $\mathcal{S}_c(V) = \{\mathbf{x} \in \mathcal{X} : V(\mathbf{x}) < c\}$, for $c \in \mathbb{R}_+$, that is contained within \mathcal{D} is invariant under the dynamics f .

3. Problem Statement

We consider a discrete-time TIA system as in (1), where the control signal is produced by a feedback-controller $\pi(\cdot; \psi) : \mathcal{X} \rightarrow \mathcal{U}$ parameterized by ψ . Therefore, the closed-loop dynamics denoted by f_π is a functional of the controller and is consequently parameterized by ψ as $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \pi(\mathbf{x}_k; \psi)) = f_\pi(\mathbf{x}_k; \psi)$. Without loss of generality, we assume that the equilibrium point of interest is located at the origin $\bar{\mathbf{x}} = \mathbf{0}$. The policy π induces a RoA around the equilibrium point denoted by \mathcal{R}_π .²

Each control task can be broken down into two subtasks: **1)** Controller synthesis and **2)** Closed-loop response evaluation. The controller is designed to optimize some measure of performance. In this work, the performance measure is the size of the stability region.

We endow the state space \mathcal{X} with a measure μ to obtain the measure space $(\mathbb{R}, \mathcal{B}(\mathbb{R}^d), \mu)$ with Borel sigma-algebra \mathcal{B} . To prevent pathological cases, we assume \mathcal{X} to be compact with $\mu(\mathcal{X}) < \mu_\infty < \infty$. Let $\Pi = \{\pi : \mathcal{X} \rightarrow \mathcal{U} : \pi \in C^1 \text{ and bounded}\}$ be the set of all functions from which the policy is chosen. The goal is to find a member π^* of the equivalence class of optimal policies $\Pi^* \subseteq \Pi$, where Π^* is defined as $\Pi^* = \{\pi^* \in \Pi : \mu(\mathcal{R}_{\pi^*}) = \max_{\pi \in \Pi} \mu(\mathcal{R}_\pi)\}$.

The main challenge in this optimization problem is the fact that there is no analytical or straightforward way to infer how changing π changes $\mu(\mathcal{R}_\pi)$. If there exists a differentiable map from π to $\mu(\mathcal{R}_\pi)$, one could locally increase the stability region by perturbing the policy in the direction of $\partial\mu(\mathcal{R}_\pi)/\partial\pi$. However, unless for extremely simple systems, such a map cannot be derived analytically. In this work, we construct a bridge between these two objects by an auxiliary function, which is an evolving Lyapunov function that is learned alongside the policy.

2. We drop $\bar{\mathbf{x}}$ from the superscript of $\mathcal{R}_\pi^{\bar{\mathbf{x}}}$, since we always assume $\bar{\mathbf{x}} = \mathbf{0}$, unless otherwise stated.

Our goal is to construct a sequence of policies (π_1, π_2, \dots) that gives a sequence of RoAs $(\mathcal{R}_{\pi_1}, \mathcal{R}_{\pi_2}, \dots)$, such that $\mathcal{R}_{\pi_n} \xrightarrow{\mu} \bar{\mathcal{R}}$ as $n \rightarrow \infty$, where $\bar{\mathcal{R}} \subseteq \mathcal{D}$ is the largest achievable RoA that is constrained by the physical limitations of the system (see Appendix A for the characterization of $\bar{\mathcal{R}}$ using the concept of control Lyapunov function). To achieve this goal, we need to address two challenges: **1)** Approximating \mathcal{R}_{π_n} for a fixed π_n and **2)** Using (f, \mathcal{R}_{π_n}) to find π_{n+1} . Next section, explains our proposed method to address these two challenges.

4. Proposed Method

For a function g , let's define its sublevel set with level value a as $\mathcal{S}_a(g) := \{\mathbf{x} \in \mathbb{R}^d : g(\mathbf{x}) < a\}$. The index $n \in \mathbb{N} \cup \{0\}$ refers to a phase of the algorithm. At phase n , let \mathcal{R}_{π_n} be the RoA of the closed-loop system (1) that is induced by the state-feedback policy π_n . It can be shown that there exists an optimal Lyapunov function V_{π_n} for this system with a level value c_n , such that $\mathcal{R}_{\pi_n} = \mathcal{S}_{c_n}(V_{\pi_n})$ (Vannelli and Vidyasagar, 1985). Therefore, the information of \mathcal{R}_{π_n} is encoded in (V_{π_n}, c_n) . Starting with a conservative controller (e.g., a quadratic controller for the linearized system), our method inductively constructs a sequence of policies (π_1, π_2, \dots) that eventually converges to a policy with maximal achievable RoA.

Each step of this inductive process is called a *phase* of the algorithm. Each phase consists of two sub-phases: **1)** Learning the Lyapunov function and the RoA corresponding to the policy of that phase **2)** Updating the policy to enlarge the RoA. The RoA estimation sub-phase finds a Lyapunov function V_{π_n} and level value c_n , such that $\mathcal{S}_{c_n}(V_{\pi_n}) = \mathcal{R}_{\pi_n}$. Then, the policy update phase learns a new policy π_{n+1} , such that $\mathcal{S}_{c_n}(V_{\pi_n}) \subseteq \mathcal{R}_{\pi_{n+1}}$. These conditions need to be satisfied for every n that consequently limits the class of treatable systems. Before delving into the algorithmic implementation, Section 4.1 provides necessary theoretical insights into this favorable class of systems along with other theoretical considerations for the proposed multi-phase growing algorithm.

4.1. Theoretical Discussion

We start with the assumptions that are necessary for the practical applicability of the method. Then, we show for which class of systems these assumptions are satisfied.

Assumption 1 *Let $R : \Pi \rightarrow 2^{\mathcal{X}}$ be a set-valued function defined as $R(\pi) = \mathcal{R}_{\pi}$. Let $d_{\Pi} : \Pi \times \Pi \rightarrow \mathbb{R}_+$ and $d_{\mathcal{X}} : 2^{\mathcal{X}} \times 2^{\mathcal{X}} \rightarrow \mathbb{R}_+$ be some specified metrics in the space of policies and the space of all subsets of the state space, respectively. Then, the map R is assumed to be continuous with respect to the topologies induced by the metrics d_{Π} and $d_{\mathcal{X}}$.*

Intuitively, this assumption indicates that a small change in the policy leads to a small change in the RoA that it induces. This property that we refer to as the *persistence* of RoA is formalized as follows: let $(\mathcal{R}_{\pi_{n-1}}, \pi_n)$ comprise the RoA of the previous step and the policy of the current step for which we want to find the RoA. Given Assumption 1, for every $\epsilon > 0$, one can choose small enough $\delta > 0$ such that $d_{\Pi}(\pi_n, \pi_{n-1}) < \delta$ results in $\mu(\mathcal{R}_{\pi_n} \triangle \mathcal{R}_{\pi_{n-1}}) < \epsilon$. We now show in the following theorem that this assumption is in fact valid for hyperbolic systems.

Theorem 2 (Persistence of the stability boundary with variations in the policy) *Consider the closed-loop hyperbolic system $\mathbf{x}_{k+1} = f_{\pi}(\mathbf{x}_k)$ with policy π . Suppose f_{π} is a diffeomorphism and all equilibrium points on $\partial\mathcal{R}_{\bar{\pi}}$ are hyperbolic. Moreover, let the stable and unstable manifolds of*

the equilibrium points on $\partial\mathcal{R}_{\tilde{\pi}}$ intersect transversally³. Finally, assume that every trajectory on $\partial\mathcal{R}_{\tilde{\pi}}$ approaches one of the equilibrium points. For a certain policy $\pi = \tilde{\pi}$, let $(\bar{\mathbf{x}}_{\tilde{\pi}}, \mathcal{R}_{\tilde{\pi}})$ be an asymptotically stable equilibrium point and its corresponding RoA. Then, for every $\epsilon > 0$, there exists $\delta > 0$ such that for every π' with $d_{\Pi}(\tilde{\pi}, \pi') < \delta$, we have $\mu(\mathcal{R}_{\tilde{\pi}} \triangle \mathcal{R}_{\pi'}) < \epsilon$.

Proof The first step to prove this result is to characterize the RoA in terms of the properties of the dynamics function f_{π} . As $\mathcal{R}_{\tilde{\pi}}$ is characterized by its boundary, we focus our attention on the stability boundary $\partial\mathcal{R}_{\tilde{\pi}}$. The critical elements⁴ of the dynamical system determine the structure of the stability boundary. If $\{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ are the hyperbolic equilibrium points on $\partial\mathcal{R}_{\tilde{\pi}}$, $\partial\mathcal{R}_{\tilde{\pi}}$ is completely characterized by

$$\partial\mathcal{R}_{\tilde{\pi}} = \cup_i W^s(\mathbf{x}_i). \quad (4)$$

See Theorems 9-11 in [Chiang and Alberto \(2015\)](#) for the detailed proof. With this characterisation, to show the persistence of $\partial\mathcal{R}_{\tilde{\pi}}$, it is enough to show the persistence of the equilibrium points that live on $\partial\mathcal{R}_{\tilde{\pi}}$ and the persistence of their stability condition. As a result of the continuity of $f(\mathbf{x}, \pi(\mathbf{x}))$ and $\pi(\mathbf{x})$ w.r.t. their arguments, implicit function theorem guarantees that small perturbations to π cause small changes in the hyperbolic equilibrium points ([Krantz and Parks, 2012](#)). If $\mathbf{x}_{\tilde{\pi}}^*$ is a hyperbolic equilibrium point of $\mathbf{x}_{k+1} = f_{\pi}(\mathbf{x}_k)$ for the policy $\pi = \tilde{\pi}$, there exists a $\delta > 0$ and a neighborhood U of $\mathbf{x}_{\tilde{\pi}}^*$ that contains a unique hyperbolic equilibrium point $\mathbf{x}_{\pi'}^*$, for every $\pi' \in \{\pi \in \Pi : d_{\Pi}(\pi, \tilde{\pi}) < \delta\}$. Similarly, the continuity of the eigenvalues of $J_{f_{\pi}}(\mathbf{x}_{\pi}^*)$ w.r.t. π affirms that the perturbed equilibrium point $\mathbf{x}_{\pi'}^*$ has the same stability condition as $\mathbf{x}_{\tilde{\pi}}^*$, i.e., there exists a homeomorphism $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ from $\Phi_{\tilde{\pi}}(\mathbf{x}, t)$ to $\Phi_{\pi'}(\mathbf{x}, t)$. As stated above and given the results of [Chiang and Alberto \(2015\)](#), $\partial\mathcal{R}_{\tilde{\pi}}$ is characterized by the stable manifolds of unstable equilibrium points living on the boundary⁵. It was also shown that the hyperbolic equilibrium points on $\partial\mathcal{R}_{\tilde{\pi}}$ together with their stable and unstable manifolds change continuously with $\tilde{\pi}$. This results in a continuous change of $\partial\mathcal{R}_{\tilde{\pi}}$, and consequently $\mathcal{R}_{\tilde{\pi}}$, w.r.t. small variations in $\tilde{\pi}$. ■

This result suggests restricting the change of the policy at the policy update sub-phase of each growing phase of the algorithm to bound the change of its induced RoA. For parametric policies such as neural networks, this is achieved by cropping parameter values after training.

4.2. Algorithm

The implementation of the algorithm that was outlined in Section 4 comes in the following. Notice that the algorithm is a multi-phase inductive method. Hence, we only present two sub-phases of a single phase. The entire algorithm is iteration over this phase multiple times.

RoA Estimation Sub-Phase: Learn \mathcal{R}_{π_n} from $(\mathcal{R}_{\pi_{n-1}}, \pi_{n-1})$. This sub-phase takes the previous policy π_{n-1} and its associated RoA estimate $\mathcal{S}_{c_{n-1}}(V_{\pi_{n-1}})$ and outputs $\mathcal{S}_{c_n}(V_{\pi_n})$ that approximates \mathcal{R}_{π_n} . For this sub-phase, we improve the growing algorithm of ([Richards et al., 2018](#)) to learn an inner estimate of RoA. Our RoA estimation algorithm takes advantage of the theoretical results of ([Chiang and Thorp, 1989](#)) to improve the stability of training. The proposed loss function is more robust against local minima observed by [Richards et al. \(2018\)](#) as is empirically shown in Section 5.

3. Roughly speaking, the manifolds intersect in a generic way.

4. For theoretical discussion, we focus on critical elements that are equilibrium points. Similar results exist for other types of critical elements such as limit cycles but are left out of this work for brevity

5. Knowing the stable manifolds of the equilibriums that live on the boundary is enough to determine the boundary.

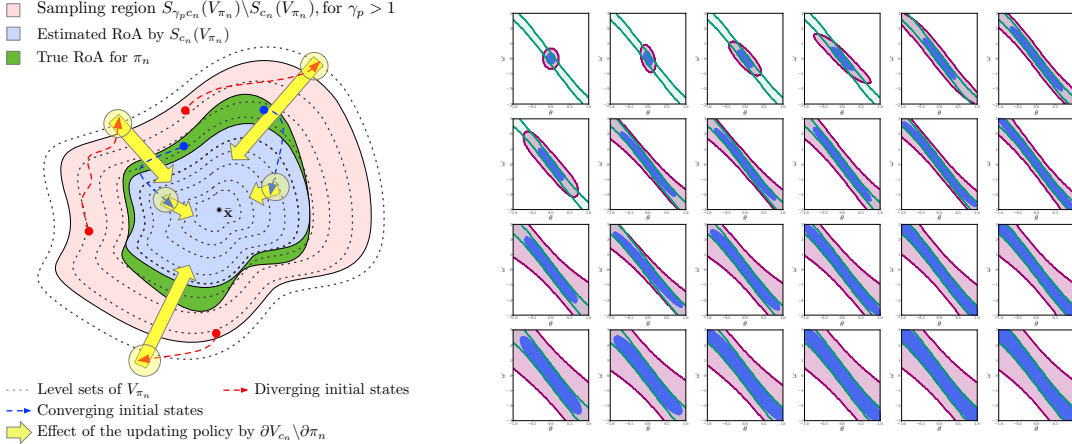


Figure 1: (Left) Illustration of the policy update sub-phase. Given the estimated RoA, the policy update sub-phase (yellow arrows) tries to pull the diverging trajectories towards the level sets of the estimated Lyapunov function that reside inside the RoA. (Right) Visualizing the true RoA which is enlarged by the improved policy and is chased by a learned Lyapunov function. Green boundary: True RoA, Blue: $S_{c_n}(V_{\pi_n})$, Pink: $S_{\gamma_r c_n}(V_{\pi_n})$ for $\gamma_r = 4$. The pink area shows the region from which the samples outside the estimated RoA are taken for both estimating the RoA and updating the policy.

The learning algorithm is verbally described here and the pseudo-code can be seen in Algorithm 1 in the Appendix. It starts with an initial conservative estimate of RoA and grows it during multiple iterations (phases). A neural network, as a universal function approximator, is employed to estimate a Lyapunov function for the system. At each iteration, the Lyapunov function is improved such that one of its level sets gives a better estimate of the RoA. For this purpose, a fixed number of initial states are taken from a gap surrounding the current estimate of the RoA. The initial states are integrated forward in time by the closed-loop dynamics to produce the solution $\Phi(\mathbf{x}, k)$ for each initial state \mathbf{x} . The final state of each trajectory determines the stability of its corresponding initial state. We train the Lyapunov function as a classifier. This allows the method to enjoy the well-developed supervised learning tools that have been exceptionally successful in scaling to higher dimensions especially when implemented as neural networks. The current estimate of the RoA produced by the current Lyapunov function is called the largest stable level set. The initial states are labelled *stable* if their trajectory enters the largest stable level set and are labelled *unstable* otherwise. We denote these two sets of initial states as \mathbb{X}^{IN} and \mathbb{X}^{OUT} respectively and minimize the loss function

$$\mathcal{L}(V) = \sum_{\mathbf{x} \in \mathbb{X}^{\text{IN}}} [V(\mathbf{x}) - \bar{c}] - \sum_{\mathbf{x} \in \mathbb{X}^{\text{OUT}}} [V(\mathbf{x}) - \bar{c}] + \lambda_{\text{RoA}} \sum_{\mathbf{x} \in \mathbb{X}^{\text{IN}}} \Delta V(\mathbf{x}) + \lambda_{\text{monot}} \sum_{\mathbf{x} \in \mathbb{X}^{\text{IN}}} [V(\mathbf{x}) - V_{\pi_{n-1}}(f_{\pi_{n-1}}(\mathbf{x}))]^2 \quad (5)$$

to update the Lyapunov function where \bar{c} is fixed to a constant value (1 throughout this work). The idea is to absorb this degree of freedom in V to ease training. The loss function is minimized by automatic differentiation and Stochastic Gradient Descent (SGD) with respect to the parameters of the neural network that realizes V . Once it is minimized, a line search is carried out on the level value to obtain c_n such that $S_{c_n}(V_{\pi_n})$ does not exceed the true RoA. The same process repeats for a certain number of iterations until a good inner estimate of the RoA is achieved. An important component in this process is the size of the gap around the largest stable level set of each iteration of the algorithm. This gap is produced by $\mathcal{G}_n = S_{\gamma_r c_n}(V_{\pi_n}) \setminus S_{c_n}(V_{\pi_n})$ with $\gamma_r > 1$. The size of this gap is controlled by γ_r . Larger values of γ_r gives a faster convergence but less stable training.

To give an intuitive idea of the terms in Equation (5), the first two terms encourage the Lyapunov function to change in a way that its level set $S_{\bar{c}}(V)$ includes the stable initial states \mathbb{X}^{IN} and

excludes the unstable initial states \mathbb{X}^{OUT} . The third term weighted by λ_{RoA} encourages the negative definiteness of ΔV on the RoA. The last term is inspired by the constructive method of (Chiang and Thorp, 1989) that accelerates capturing the entire RoA (see Appendix C). To satisfy Lyapunov conditions, V needs to be positive definite on its domain. Rather than treating this condition in the loss function, we encode it in the architecture of the neural network using the construction proposed by Richards et al. (2018). Check Appendix G.1 for the detailed description of the architecture.

Policy Update Sub-Phase: Learn π_{n+1} from \mathcal{R}_{π_n} : This sub-phase of the algorithm uses the estimated Lyapunov function V_{π_n} to update the policy so that the new policy gives rise to a larger RoA. The idea is to change the policy in a way that the unstable trajectories starting from around the current RoA enter the RoA and become stable. Let \mathcal{D} be the domain of f around the equilibrium $\bar{\mathbf{x}}$. Given a hypothesis class of feasible policies Π , only a subset of the domain \mathcal{D} is stabilizable. Assume $\bar{\mathcal{B}} \subseteq \mathcal{D}$ is the maximum stabilizable subset of the domain with $\mu(\bar{\mathcal{B}}) = \bar{\mu}$. Therefore, an attempt to improve π_n amounts to appending points from $\bar{\mathcal{B}} \setminus \mathcal{R}_{\pi_n}$ to \mathcal{R}_{π_n} . The set $\bar{\mathcal{B}}$ is not fully known in advance but some of its properties can be derived. Especially, for system (1), if $f, \pi \in C^\infty$, the maximum stabilizable set $\bar{\mathcal{B}}$ whose measure materializes as $\bar{\mu}$ is compact and connected. Using this theoretical result, if $\mathcal{R}_{\pi_n} \subsetneq \bar{\mathcal{B}}$, the stabilizable states can be chosen in a gap around \mathcal{R}_{π_n} .

Because the RoA estimation sub-phase estimates it as a level set of a Lyapunov function, i.e. $\mathcal{R}_{\pi_n} \approx \mathcal{S}_{c_n}(V_{\pi_n})$, the sampling gap is constructed as $\mathcal{G}_n = \mathcal{S}_{\gamma_p c_n}(V_{\pi_n}) \setminus \mathcal{S}_{c_n}(V_{\pi_n})$ for a $\gamma_p > 1$. To make sure the policy does not destabilize the regions that are already stabilized in the previous phases, algorithm also samples initial states from $\mathcal{S}_{c_n}(V_{\pi_n})$. All sampled initial states are integrated forward for L_p steps and the policy is updated via minimizing the loss function

$$\mathcal{L}(\pi) = \sum_{\mathbf{x} \in \mathcal{G}_n \cup \mathcal{S}_{c_n}(V_{\pi_n})} [1_{[V_{\pi_n}(\Phi_\pi(\mathbf{x}, L_p)) < c_n]} + \lambda_u 1_{[V_{\pi_n}(\Phi_\pi(\mathbf{x}, L_p)) > c_n]}] V_{\pi_n}(\Phi_\pi(\mathbf{x}, L_p)). \quad (6)$$

We implement the policy as a differentiable function such as a neural network and use automatic differentiation and SGD for minimization. The parameters of the policy appear in (6) via the end state $\Phi_\pi(\mathbf{x}, L_p)$ of the closed-loop trajectories. It is clear in (6) that minimizing $\mathcal{L}(\pi)$ with respect to π while the Lyapunov function is fixed, pushes the trajectories towards the areas where the Lyapunov function assumes smaller values. This affects both stable and unstable trajectories while its influence on unstable trajectory can be magnified by choosing $\lambda_u > 1$. The detailed pseudo-code of this phase can be found in Algorithm 2 in the Appendix. The quality of the training signal is analyzed next.

Theoretical analysis of the training signal—The effect of the policy π on $\mathcal{L}(\pi)$ passes through the Lyapunov function V as can be seen in (6). Unlike the conventional Lyapunov redesign method in control theory where the Lyapunov function is fixed, here the Lyapunov function itself is learned by the RoA estimation sub-phase of the algorithm. Hence, an ill-conditioned V can harm the policy update phase. To take a closer look at this issue, we expand the training signal analytically. Let the policy π be a function of states parameterized by ψ . Let $T = L_p$ be the time step of the final state of the trajectory. The training signal to update the policy is proportional to $\partial \mathcal{L} / \partial \psi$ expanded as

$$\frac{\partial \mathcal{L}}{\partial \psi} = \frac{\partial \mathcal{L}}{\partial \mathbf{x}_T} \sum_{1 \leq k \leq T} \left(\frac{\partial \mathbf{x}_T}{\partial \mathbf{x}_k} \frac{\partial^+ \mathbf{x}_k}{\partial \psi} \right). \quad (7)$$

by applying the chain rule for differentiation. The term $\partial^+ \mathbf{x}_k / \partial \psi$ is the single-step effect of ψ on \mathbf{x}_k when \mathbf{x}_{k-1} is fixed. We discuss every term in this equation in the following. Observe that $\partial \mathcal{L} / \partial \mathbf{x}_T$ is multiplied the summation, i.e., its small value diminish the entire signal. It can be expanded as

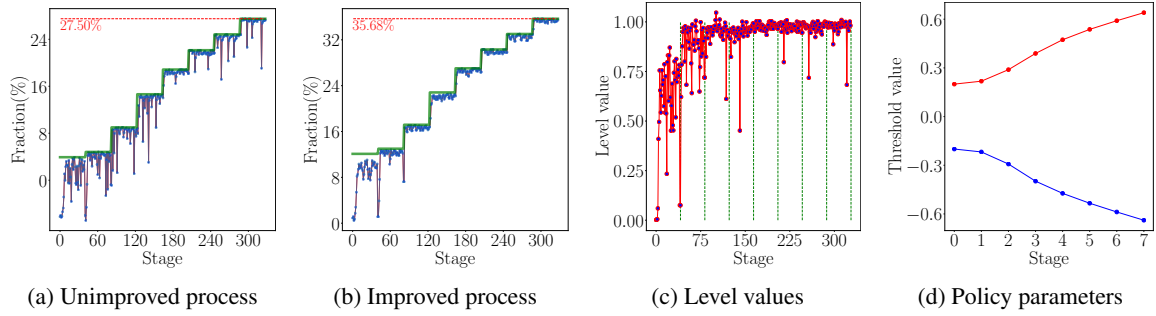


Figure 2: (a), (b) The size of the RoA against the iterative stages (phases) of the Algorithms where (a) uses the RoA estimation method of Richards et al. (2018) while (b) uses ours. The fraction is computed with respect to a rectangular domain around the equilibrium point that is large enough to enclose the RoA. Green: size of the true RoA. Each jump corresponds to a policy update sub-phase that increases the size of the true RoA. Red: The size of level set that the RoA estimation sub-phase learns to approximate the RoA. After each policy update, the RoA estimation sub-phase takes multiple growth iterations to capture the true RoA as close as possible. (c) The trace of the level values corresponding to every iteration of the RoA estimation sub-phase. (d) Red (Blue): The trace of the value of the upper (lower) threshold parameter of the policy during training. Each point corresponds to a policy update iteration.

$$\frac{\partial \mathcal{L}(\mathbf{x})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_T} = \frac{\partial \mathcal{L}(V)}{\partial V} \Big|_{V=V(\mathbf{x}_T)} \frac{\partial V(\mathbf{x})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_T}. \quad (8)$$

The first term on the r.h.s. does not vanish as it is 1 or λ_u for the loss function defined by Equation (6). The second term of the r.h.s depends on the slope of V evaluated at the final state of the trajectory. One potential pathological condition occurs for stable and long trajectories. The reason is that $\nabla_{\mathbf{x}} V(\mathbf{x})$ continuously vanishes at the equilibrium (see Lemma 4 in the Appendix). This means that the derivative of the Lyapunov functions gradually becomes smaller as we get closer to the equilibrium. Therefore, for long stable trajectories that \mathbf{x}_T is too close to the equilibrium, the learning signal to update the policy will be too small. Too long trajectories are detrimental for unstable trajectories as well because the states may grow exponentially and cause damages to the system. Hence, the length of the trajectory is an important design parameter that needs special attention when applying our proposed method in practice. The terms inside the sum depend on the properties of the dynamics. Specifically, the first term shows how long the system keeps the memory of the past states and the second term shows how sensitive the system is with respect to the controller parameters. A more detailed discussion is deferred to Appendix F.

5. Experiments

We consider an inverted pendulum system defined as $\dot{\theta} = \omega$ and $\dot{\omega} = \frac{g}{l} \sin(\theta) + \frac{u}{I} - \mu_f \frac{\omega}{I}$ where the state vector $\mathbf{x} = (\theta, \omega)$ consists of the angle and angular velocity. Moreover ($g = 0.81, l = 0.5, I = \text{mass} \times l^2 = 0.25, \mu_f = 0$) are the acceleration of gravity, length, inertia, and friction coefficient. The scalar u is the input force. The open-loop system (with $u = 0$) has equilibrium points at $(\theta, \omega) = (k\pi, 0)$ with $k \in \mathbb{Z}$. We focus on the equilibrium point $(0, 0)$ in the frictionless setting where the system shows oscillatory behavior and consequently is not asymptotically stable (see Appendix G.2 for system and modeling details). First, an LQR controller K is designed for the linearized system around the origin (see the vector field and the initial RoA




Figure 4: Loose saturation

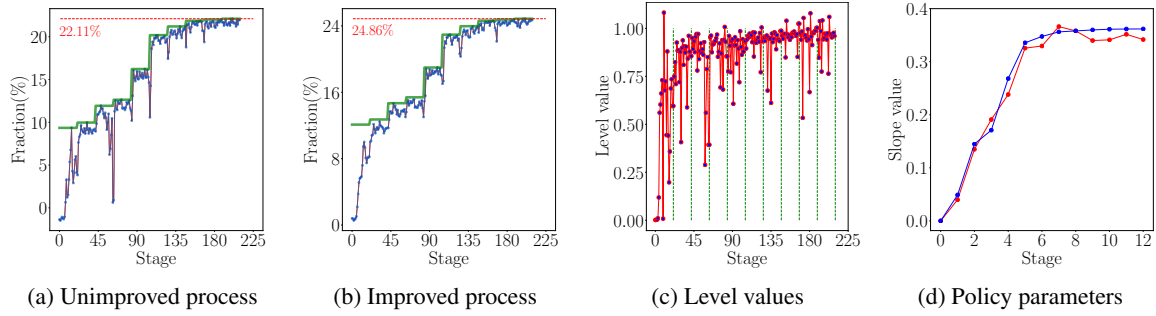


Figure 3: The same plots as Figure 2 with the exception that, here the threshold parameters are kept fixed at $a = 0.2$ and $b = -0.2$ while the slopes m_a and m_b are trainable parameters.

in Figure 5 in the appendix). The control signal passes through a loose saturation function as $u = \pi_0(\mathbf{x}; \psi) = \text{SAT}_\psi([\theta, \omega]^\top K)$. The function SAT is parameterized by $\psi = (a, b, m_a, m_b)$ as illustrated in Figure 4. In the first experiment, the slopes $m_a = 0$ and $m_b = 0$ are kept fixed where a and b are trainable parameters of the policy, i.e., $\psi = (a, b)$. The Lyapunov function $V(\cdot; \theta)$ is realized by a 3-layer neural network parameterized by θ . Each layer has 64 neurons with a special architecture (see Appendix G.1) inspired by (Richards et al., 2018) followed by tanh activation function that imposes the positive definiteness of the entire network as is required by the Lyapunov conditions in (2). The chosen hyper-parameters for the SGD training can be found in Appendix G.3.

The initial controller gives a small RoA since it is designed for the locally linearized system. As the initial policy is LQR designed for the locally linearized model, $V(\cdot; \theta)$ is pre-trained by the quadratic function $0.1\theta^2 + 0.1\omega^2$. After pre-training, sub-phases of the algorithm of Section 4.2 are run alternately to capture the RoA and improve the policy. The green step-like plot in Figure 2a and Figure 2b shows the true size of the RoA. Each jump in the green plots shows one iteration of the policy update algorithm resulting in an increased RoA. The fluctuating red plot in Figure 2a shows the size of the estimated RoA without our improvement over the RoA estimation algorithm of Richards et al. (2018) while Figure 2b shows the outcome of the presence of our proposed additional term in the loss function (5). It shows that the added term results in a less fluctuating estimate of the RoA, and when combined with the policy update phase, gives a faster convergence to a larger RoA (35.68% vs 27.50% fraction of the domain volume after 7 policy updates).

As stated in Section 4.2, $\bar{c} = 1$ is not necessarily equal to c_n . After learning V_n with $\bar{c} = 1$ in (5), the algorithm searches for a value of c_n such that the Lyapunov decrease condition is met for all states within the level set $\mathcal{S}_{c_n}(V_n)$. It can be seen in Figure 2c that these values converge to $\bar{c} = 1$ as a sign of stable training. The trace of the parameters of the policy is shown in Figure 2d. As these policy’s trainable parameters represent the upper and lower limits of the loose threshold function, the policy learning algorithm updates them in the directions that decrease their suppressing effect that is expected from the physics of the system. Graphical visualization of the policy update and RoA estimation phase is shown in Figure 1(Right). Each row shows one phase of the algorithm. The policies are updated along the columns from top to bottom. Within one row, the policy is fixed and RoA is estimated from left to right (See Figure 7 in the appendix for a larger visualization).

In the second experiment, the threshold limits $a = -0.2$ and $b = 0.2$ are kept fixed while the slopes $\psi = (m_a, m_b)$ are trainable parameters. The rest of the training setting remains the same as the previous experiment. Figures 3a and 3b shows that the policy update phase enlarges the RoA (green plot) of the system while the RoA estimation phase (red plot) manages to follow

the new RoA after each policy update. Our improved RoA estimation algorithm results in a more monotonic convergence of the estimated RoA that ultimately learns a controller that induces a larger RoA (24.86% vs 22.11%). Similar to Figure 2c, convergence of c_n values to $\bar{c} = 1$ can be seen in Figure 3c. The trace of the upper and lower slopes are shown in Figure 3d. Almost equal learned values for upper and lower slopes are expected due to the structural symmetry of the saturation function (Figure 4) that appears in the closed-loop system.

Scalability. At each growth step of the algorithm, samples are taken from around the RoA that is an $n - 1$ dimensional surface. As these are the most informative trajectories for RoA maximization, it can be seen as a clever sampling that is hoped to scale better compared with policy update methods with uniform sampling. An empirical investigation of this hypothesis is deferred to future work.

The following remark discusses to what extent the model of the system is needed and how can this requirement be relaxed in the future.

Remark 3 *The RoA estimation phase does not need the model of the system. The system can be launched from sampled initial states and the generated trajectories are all we need in (5). The policy update phase of the algorithm requires a local estimate of the system to be able to compute $\partial V / \partial \pi$. The locality of the model is inversely proportional to the length of the trajectory L_p in (6). A detailed theoretical discussion on this point is deferred to Appendix E.*

6. Related Work and Conclusions

We have proposed a two-player collaborative and iterative algorithm that iterates over two sub-phases that learn the Lyapunov function and use it to learn a controller to enlarge the RoA of the system. The existing approaches that are close to the purpose of this paper are those that simultaneously synthesize a controller and maximize the stability region. Unlike the large body of literature that consider tractable polynomial systems (e.g. Jarvis-Wloszek et al. (2003); Majumdar et al. (2013, 2014)) our work does not put any limit on the class of considered systems except the generic condition of hyperbolicity. Jin et al. (2020) suggest a joint training of the policy and neural certificate functions that model unsafe regions (Barrier function) and reach a target set using Lyapunov-like functions. There are two major differences with our work: Our method proposed a safe sampling using RoA information instead of uniform sampling in the state space that could be harmful to the system in practice. In addition, their objective function does not necessarily enlarge the RoA which is the main concern of our work. Boffi et al. (2020) use methods from statistical learning theory to assess the generalizability of the learned certificate functions such as Lyapunov functions. Robey et al. (2020) learn a control barrier function from safely sampled trajectories with known dynamics. These two methods are related to our work in the sense that they learn certificate functions but, unlike our work, they do not include a specialized sampling (from around the RoA) and also do not learn a policy.

In the context of reinforcement learning, our method can be seen as an actor-critic approach Grondman et al. (2012); Bhasin et al. (2013); Lillicrap et al. (2015) where the actor tries to stabilize the system while the critic estimates the size of the RoA induced by the controller. In control theory, as the title of our work suggests, the proposed algorithm can be seen as an automatic version of the known Lyapunov redesign method where the Lyapunov function and the closed-loop system are being re-designed together iteratively.

In this work, we assume the model of the system is given, even though this condition can be relaxed as suggested by Remark 3. Investigating this relaxation could be a promising future direction.

References

- Steven M Baer, Bingtuan Li, and Hal L Smith. Multiple limit cycles in the standard model of three species competition for three essential resources. *Journal of mathematical biology*, 52(6):745–760, 2006.
- Felix Berkenkamp, Angela P Schoellig, and Andreas Krause. Safe controller optimization for quadrotors with gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 491–496. IEEE, 2016.
- Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in neural information processing systems*, pages 908–918, 2017.
- Shubhendu Bhasin, Rushikesh Kamalapurkar, Marcus Johnson, Kyriakos G Vamvoudakis, Frank L Lewis, and Warren E Dixon. A novel actor–critic–identifier architecture for approximate optimal control of uncertain nonlinear systems. *Automatica*, 49(1):82–92, 2013.
- Nicholas M Boffi, Stephen Tu, Nikolai Matni, Jean-Jacques E Slotine, and Vikas Sindhvani. Learning stability certificates from data. *arXiv preprint arXiv:2008.05952*, 2020.
- Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. In *Advances in Neural Information Processing Systems*, pages 3245–3254, 2019.
- H-D Chiang and James S Thorp. Stability regions of nonlinear dynamical systems: A constructive methodology. *IEEE Transactions on Automatic Control*, 34(12):1229–1241, 1989.
- Hsiao-Dong Chiang and Luís FC Alberto. *Stability regions of nonlinear dynamical systems: theory, estimation, and applications*. Cambridge University Press, 2015.
- Ivo Grondman, Lucian Busoniu, Gabriel AD Lopes, and Robert Babuska. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1291–1307, 2012.
- John J Hopfield. Neurons, dynamics and computation. *Physics Today*, 47(2):40–47, 1994.
- Mehdi Imani Masouleh and David JN Limebeer. Region of attraction analysis for nonlinear vehicle lateral dynamics using sum-of-squares programming. *Vehicle System Dynamics*, 56(7):1118–1138, 2018.
- Alberto Isidori. *Nonlinear Control Systems Design 1989: Selected Papers from the IFAC Symposium, Capri, Italy, 14-16 June 1989*. Elsevier, 2014.
- Zachary Jarvis-Wloszek, Ryan Feeley, Weehong Tan, Kunpeng Sun, and Andrew Packard. Some controls applications of sum of squares programming. In *42nd IEEE international conference on decision and control (IEEE Cat. No. 03CH37475)*, volume 5, pages 4676–4681. IEEE, 2003.
- Wanxin Jin, Zhaoran Wang, Zhuoran Yang, and Shaoshuai Mou. Neural certificates for safe control policies. *arXiv preprint arXiv:2006.08465*, 2020.

- Hassan K Khalil and Jessy W Grizzle. *Nonlinear systems*, volume 3. Prentice hall Upper Saddle River, NJ, 2002.
- Steven G Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2012.
- F Liao and JL Wang. Analysis and synthesis of reliable flight control systems via parameter-dependent lyapunov functions. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 218(6):433–450, 2004.
- Alexandre Liapounoff. Problème général de la stabilité du mouvement. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pages 203–474, 1907.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Anirudha Majumdar, Amir Ali Ahmadi, and Russ Tedrake. Control design along trajectories with sums of squares programming. In *2013 IEEE International Conference on Robotics and Automation*, pages 4054–4061. IEEE, 2013.
- Anirudha Majumdar, Ram Vasudevan, Mark M Tobenkin, and Russ Tedrake. Convex optimization of nonlinear feedback controllers via occupation measures. *The International Journal of Robotics Research*, 33(9):1209–1230, 2014.
- J Jr Palis and Welington De Melo. *Geometric theory of dynamical systems: an introduction*. Springer Science & Business Media, 2012.
- Harry A Pierson and Michael S Gashler. Deep learning in robotics: a review of recent research. *Advanced Robotics*, 31(16):821–835, 2017.
- Spencer M Richards, Felix Berkenkamp, and Andreas Krause. The lyapunov neural network: Adaptive stability certification for safe learning of dynamic systems. *arXiv preprint arXiv:1808.00924*, 2018.
- Alexander Robey, Haimin Hu, Lars Lindemann, Hanwen Zhang, Dimos V Dimarogonas, Stephen Tu, and Nikolai Matni. Learning control barrier functions from expert demonstrations. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3717–3724. IEEE, 2020.
- Anthony Vannelli and Mathukumalli Vidyasagar. Maximal lyapunov functions and domains of attraction for autonomous nonlinear systems. *Automatica*, 21(1):69–80, 1985.
- Zhu Wen-Xing and Zhang Li-Dong. A new car-following model for autonomous vehicles flow with mean expected velocity field. *Physica A: Statistical Mechanics and its Applications*, 492: 2154–2165, 2018.
- Eric R Westervelt, Jessy W Grizzle, and C Canudas De Wit. Switching and pi control of walking motions of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(2):308–312, 2003.
- H Xin, D Gan, J Qiu, and Z Qu. Methods for estimating stability regions with applications to power systems. *European transactions on electrical power*, 17(2):113–133, 2007.