

Stability Analysis using Quadratic Constraints for Systems with Neural Network Controllers

He Yin, Peter Seiler and Murat Arcak *Fellow, IEEE*

Abstract—A method is presented to analyze the stability of feedback systems with neural network controllers. Two stability theorems are given to prove asymptotic stability and to compute an ellipsoidal inner-approximation to the region of attraction (ROA). The first theorem addresses linear time-invariant systems, and merges Lyapunov theory with local (sector) quadratic constraints to bound the nonlinear activation functions in the neural network. The second theorem allows the system to include perturbations such as unmodeled dynamics, slope-restricted nonlinearities, and time delay, using integral quadratic constraint (IQCs) to capture their input/output behavior. This in turn allows for off-by-one IQCs to refine the description of activation functions by capturing their slope restrictions. Both results rely on semidefinite programming to approximate the ROA. The method is illustrated on systems with neural networks trained to stabilize a nonlinear inverted pendulum as well as vehicle lateral dynamics with actuator uncertainty.

I. INTRODUCTION

The paradigm of stabilizing dynamical systems with Neural Network (NN) controllers [1] has been revived following recent development in deep NNs, e.g. policy gradient [2]–[5] and behavioral cloning [6]. However, feedback systems with NN controllers suffer from lack of stability and safety certificates due to the complexity of the NN structure. Specifically, NNs have various types of nonlinear activation functions, potentially numerous layers, and a large number of hidden neurons, making it difficult to apply classical analysis methods, e.g. Lyapunov theory [7]. Monte-Carlo simulations can be used to investigate stability but lack formal guarantees, which are important in safety-critical applications.

Several works propose using quadratic constraints (QCs) to bound the nonlinear activation functions. This approach is used to outer-bound the outputs of a (static) NN given a set of inputs in [8] and upper-bound the Lipschitz constant of NNs in [9], [10]. The work [11] uses this idea for finite-time reachability analysis of a system with a NN controller. The work [12] performs stability analysis by constructing QCs from the bounds of partial gradients of NN controllers. Reference [13] assesses global asymptotic stability of dynamic neural network models using QCs and Lyapunov theory.

This paper presents two main stability results for a feedback system with a NN controller. Theorem 1 provides a condition to prove stability and to inner-approximate the region of attraction (ROA) for a linear time-invariant (LTI) plant. It uses Lyapunov theory, and local (sector) QCs to bound the

nonlinear activation functions in the NN. Theorem 2 allows the plant to include perturbations such as unmodeled dynamics, slope-restricted nonlinearities, and time delay, characterizing them with integral quadratic constraints (IQCs) [14], [15]. This in turn allows for the use of off-by-one IQCs [16] to capture the slope restrictions of activation functions. Both results rely on semidefinite programming to approximate the ROA.

The specific contributions of this paper are three-fold. First, our nominal analysis with LTI plants and NN controllers uses *offset* local sector QCs that are centered around the equilibrium inputs to the NN activation functions and allow for analyzing stability around a non-zero equilibrium point. Second, our analysis of uncertain plants and NN controllers provides robustness guarantees for the feedback system. The uncertain plant is modeled as an interconnection of the nominal plant and perturbations that are described by IQCs. The use of IQCs also allows for plants that are not necessarily LTI. Third, the proposed framework allows for local (dynamic) off-by-one IQCs to further sharpen the description of activation functions by capturing their slope restrictions. This differs from [8]–[12], which derive only static QCs for activation functions.

Local (static) sector IQCs have been used in the stability analysis of linear systems with actuator saturation [17], [18], and unbounded nonlinearities [19]. The description of these nonlinearities are refined by incorporating soft (dynamic) IQCs in the stability analysis framework for linear systems [20], and polynomial systems [21]. Compared with these works, this work is specialized to NN-controlled systems: it exploits the specific properties of NNs and uses the Interval Bound Propagation method [22] to derive non-conservative static and dynamic local IQCs to describe NN controllers; and it also allows for the analysis of NN-controlled nonlinear systems by accommodating perturbations.

The paper is organized as follows. Section II presents the problem formulation and the nominal stability analysis when the plant is LTI. Section III addresses uncertain systems using IQCs. Section IV provides numerical examples, including a nonlinear inverted pendulum and an uncertain vehicle model.

Notation: $\mathbb{S}^n$ denotes the set of n -by- n symmetric matrices. $\mathbb{S}_+^n$ and $\mathbb{S}_{++}^n$ denote the sets of n -by- n symmetric, positive semidefinite and positive definite matrices, respectively. $\mathbb{RL}_\infty$ is the set of rational functions with real coefficients and no poles on the unit circle. $\mathbb{RH}_\infty \subset \mathbb{RL}_\infty$ contains functions that are analytic in the closed exterior of the unit disk in the complex plane. $\ell_2^{n_x}$ is the set of sequences $x : \mathbb{N} \rightarrow \mathbb{R}^{n_x}$ with $\|x\|_2 := \sqrt{\sum_{k=0}^{\infty} x(k)^\top x(k)} < \infty$. When applied to vectors, the orders $>, \leq$ are applied elementwise. For $P \in \mathbb{S}_{++}^n$, $x_* \in \mathbb{R}^n$, define the ellipsoid

$$\mathcal{E}(P, x_*) := \{x \in \mathbb{R}^n : (x - x_*)^\top P (x - x_*) \leq 1\}. \quad (1)$$

Funded in part by the Air Force Office of Scientific Research grant FA9550-18-1-0253, the Office of Naval Research grant N00014-18-1-2209, and the U.S. National Science Foundation grant ECCS-1906164.

H. Yin and M. Arcak are with the University of California, Berkeley {he_yin, arcak}@berkeley.edu.

P. Seiler is with the University of Michigan, Ann Arbor pseiler@umich.edu.

II. NOMINAL STABILITY ANALYSIS

A. Problem Formulation

Consider the feedback system consisting of a plant G and state-feedback controller π as shown in Figure 1. As a first step, we assume the plant G is a linear, time-invariant (LTI) system defined by the following discrete-time model:

$$x(k+1) = A_G x(k) + B_G u(k), \quad (2)$$

where $x(k) \in \mathbb{R}^{n_G}$ is the state, $u(k) \in \mathbb{R}^{n_u}$ is the input, $A_G \in \mathbb{R}^{n_G \times n_G}$ and $B_G \in \mathbb{R}^{n_G \times n_u}$. The controller $\pi : \mathbb{R}^{n_G} \rightarrow \mathbb{R}^{n_u}$ is an ℓ -layer, feed-forward neural network (NN) defined as:

$$w^0(k) = x(k), \quad (3a)$$

$$w^i(k) = \phi^i(W^i w^{i-1}(k) + b^i), \quad i = 1, \dots, \ell, \quad (3b)$$

$$u(k) = W^{\ell+1} w^\ell(k) + b^{\ell+1}, \quad (3c)$$

where $w^i \in \mathbb{R}^{n_i}$ are the outputs (activations) from the i^{th} layer and $n_0 = n_G$. The operations for each layer are defined by a weight matrix $W^i \in \mathbb{R}^{n_i \times n_{i-1}}$, bias vector $b^i \in \mathbb{R}^{n_i}$, and activation function $\phi^i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$. The activation function ϕ^i is applied element-wise, i.e.

$$\phi^i(v) := [\varphi(v_1), \dots, \varphi(v_{n_i})]^\top, \quad (4)$$

where $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is the (scalar) activation function selected for the NN. Common choices for the scalar activation function include $\varphi(v) := \tanh(v)$, sigmoid $\varphi(v) := \frac{1}{1+e^{-v}}$, ReLU $\varphi(v) := \max(0, v)$, and leaky ReLU $\varphi(v) := \max(av, v)$ with $a \in (0, 1)$. We assume the activation φ is identical in all layers; this can be relaxed with minor changes to the notation.




Fig. 1: Feedback system with plant G and NN π

The state vector x_* is an equilibrium point of the feedback system with input u_* if the following conditions hold:

$$x_* = A_G x_* + B_G u_*, \quad (5a)$$

$$u_* = \pi(x_*). \quad (5b)$$

Let $\chi(k; x_0)$ denote the solution to the feedback system at time k from initial condition $x(0) = x_0$. Our goal is to analyze asymptotic stability of the equilibrium point and to find the largest estimate of the region of attraction, defined below, using an ellipsoidal inner approximation.

Definition 1. The region of attraction (ROA) of the feedback system with plant G and NN π is defined as

$$\mathcal{R} := \{x_0 \in \mathbb{R}^{n_G} : \lim_{k \rightarrow \infty} \chi(k; x_0) = x_*\}. \quad (6)$$

B. NN Representation: Isolation of Nonlinearities

It is useful to isolate the nonlinear activation functions from the linear operations of the NN as done in [8], [13]. Define v^i as the input to the activation function ϕ^i :

$$v^i(k) := W^i w^{i-1}(k) + b^i, \quad i = 1, \dots, \ell. \quad (7)$$

The nonlinear operation of the i^{th} layer (3b) is thus expressed as $w^i(k) = \phi^i(v^i(k))$. Gather the inputs and outputs of all activation functions:

$$v_\phi := \begin{bmatrix} v^1 \\ \vdots \\ v^\ell \end{bmatrix} \in \mathbb{R}^{n_\phi} \text{ and } w_\phi := \begin{bmatrix} w^1 \\ \vdots \\ w^\ell \end{bmatrix} \in \mathbb{R}^{n_\phi}, \quad (8)$$

where $n_\phi := n_1 + \dots + n_\ell$, and define the combined nonlinearity $\phi : \mathbb{R}^{n_\phi} \rightarrow \mathbb{R}^{n_\phi}$ by stacking the activation functions:

$$\phi(v_\phi) := \begin{bmatrix} \phi^1(v^1) \\ \vdots \\ \phi^\ell(v^\ell) \end{bmatrix}. \quad (9)$$

Thus $w_\phi(k) = \phi(v_\phi(k))$, where the scalar activation function φ is applied element-wise to each entry of v_ϕ . Finally, the NN control policy π defined in (3) can be rewritten as:

$$\begin{bmatrix} u(k) \\ v_\phi(k) \end{bmatrix} = N \begin{bmatrix} x(k) \\ w_\phi(k) \\ 1 \end{bmatrix} \quad (10a)$$

$$w_\phi(k) = \phi(v_\phi(k)). \quad (10b)$$

The matrix N depends on the weights and biases as follows, where the vertical and horizontal bars partition N compatibly with the inputs $(x, w_\phi, 1)$ and outputs (u, v_ϕ) :

$$N := \left[\begin{array}{c|ccc|cc} 0 & 0 & 0 & \dots & W^{\ell+1} & b^{\ell+1} \\ \hline W^1 & 0 & \dots & 0 & 0 & b^1 \\ 0 & W^2 & \dots & 0 & 0 & b^2 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & W^\ell & 0 & b^\ell \end{array} \right] \quad (11a)$$

$$:= \begin{bmatrix} N_{ux} & N_{uw} & N_{ub} \\ N_{vx} & N_{vw} & N_{vb} \end{bmatrix}. \quad (11b)$$

This decomposition of the NN, depicted in Figure 2, isolates the activation functions in preparation for the stability analysis.




Fig. 2: NN representation to isolate the nonlinearities ϕ .

Suppose (x_*, u_*) satisfies (5). Then x_* can be propagated through the NN to obtain equilibrium values v_*^i, w_*^i for the inputs/outputs of each activation function ($i = 1, \dots, \ell$), yielding $(v_\phi, w_\phi) = (v_*, w_*)$. Thus (x_*, u_*, v_*, w_*) is an equilibrium point of (2) and (3) if:

$$x_* = A_G x_* + B_G u_*, \quad (12a)$$

$$\begin{bmatrix} u_* \\ v_* \end{bmatrix} = N \begin{bmatrix} x_* \\ w_* \\ 1 \end{bmatrix}, \quad (12b)$$

$$w_* = \phi(v_*). \quad (12c)$$

C. Quadratic Constraints: Scalar Activation Functions

The stability analysis relies on quadratic constraints (QCs) to bound the activation function. A typical constraint is the sector bound as defined next.

Definition 2. Let $\alpha \leq \beta$ be given. The function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ lies in the (global) sector $[\alpha, \beta]$ if:

$$(\varphi(\nu) - \alpha\nu) \cdot (\beta\nu - \varphi(\nu)) \geq 0 \quad \forall \nu \in \mathbb{R}. \quad (13)$$

The interpretation of the sector $[\alpha, \beta]$ is that φ lies between lines passing through the origin with slope α and β . Many activation functions are bounded in the sector $[0, 1]$, e.g. $\tanh$ and ReLU. Figure 3 illustrates $\varphi(\nu) = \tanh(\nu)$ (blue solid) and the global sector defined by $[0, 1]$ (red solid lines).




Fig. 3: Sector constraints on $\tanh$

The global sector constraint is often too coarse for stability analysis, and a local sector constraint provides tighter bounds.

Definition 3. Let $\alpha, \beta, \underline{\nu}, \bar{\nu} \in \mathbb{R}$ with $\alpha \leq \beta$ and $\underline{\nu} \leq 0 \leq \bar{\nu}$. The function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ satisfies the local sector $[\alpha, \beta]$ if

$$(\varphi(\nu) - \alpha\nu) \cdot (\beta\nu - \varphi(\nu)) \geq 0 \quad \forall \nu \in [\underline{\nu}, \bar{\nu}]. \quad (14)$$

As an example, $\varphi(\nu) := \tanh(\nu)$ restricted to the interval $[-\bar{\nu}, \bar{\nu}]$ satisfies the local sector bound $[\alpha, \beta]$ with $\alpha := \tanh(\bar{\nu})/\bar{\nu} > 0$ and $\beta := 1$. As shown in Figure 3 (green dashed lines), the local sector provides a tighter bound than the global sector. These bounds are valid for a symmetric interval around the origin with $\underline{\nu} = -\bar{\nu}$; non-symmetric intervals ($\underline{\nu} \neq -\bar{\nu}$) can be handled similarly.

The local and global sector constraints above were defined to be centered at the point $(\nu, \varphi(\nu)) = (0, 0)$. The stability analysis will require offset sectors centered around an arbitrary point $(\nu_*, \varphi(\nu_*))$ on the function. For example, $\varphi(\nu) = \tanh(\nu)$ satisfies the global sector bound (red solid) around the point $(\nu_*, \tanh(\nu_*))$ with $[\alpha, \beta] = [0, 1]$, as shown in Figure 4. It satisfies a tighter local sector bound (green dashed) when the input is restricted to $\nu \in [\underline{\nu}, \bar{\nu}]$. An explicit expression for this local sector is $\beta = 1$ and

$$\alpha := \min \left(\frac{\tanh(\bar{\nu}) - \tanh(\nu_*)}{\bar{\nu} - \nu_*}, \frac{\tanh(\nu_*) - \tanh(\underline{\nu})}{\nu_* - \underline{\nu}} \right).$$

The local sector upper bound β can be tightened further. This leads to the following definition of an offset local sector.

Definition 4. Let $\alpha, \beta, \underline{\nu}, \bar{\nu}, \nu_* \in \mathbb{R}$ be given with $\alpha \leq \beta$ and $\underline{\nu} \leq \nu_* \leq \bar{\nu}$. The function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ satisfies the offset local sector $[\alpha, \beta]$ around the point $(\nu_*, \varphi(\nu_*))$ if:

$$(\Delta\varphi(\nu) - \alpha\Delta\nu) \cdot (\beta\Delta\nu - \Delta\varphi(\nu)) \geq 0 \quad \forall \nu \in [\underline{\nu}, \bar{\nu}] \quad (15)$$

where $\Delta\varphi(\nu) := \varphi(\nu) - \varphi(\nu_*)$ and $\Delta\nu := \nu - \nu_*$.




Fig. 4: Offset local sector constraint on $\tanh$

D. Quadratic Constraints: Combined Activation Functions

Offset local sector constraints can also be defined for the combined nonlinearity ϕ , given by (9). Let $\underline{v}, \bar{v}, v_* \in \mathbb{R}^{n_\phi}$ be given with $\underline{v} \leq v_* \leq \bar{v}$. Assume that the activation input $v_\phi \in \mathbb{R}^{n_\phi}$ lies, element-wise, in the interval $[\underline{v}, \bar{v}]$ and the i^{th} input/output pair is $w_{\phi,i} = \varphi(v_{\phi,i})$. Further assume the scalar activation function satisfies the local sector $[\alpha_i, \beta_i]$ around the point $v_{*,i}$ with the input restricted to $v_{\phi,i} \in [\underline{v}_i, \bar{v}_i]$ for $i = 1, \dots, n_\phi$. The local sector bounds can be computed for φ on the given interval either analytically (as above for $\tanh$) or numerically. These local sectors can be stacked into vectors $\alpha_\phi, \beta_\phi \in \mathbb{R}^{n_\phi}$ that provide QCs satisfied by the combined nonlinearity ϕ .

Lemma 1. Let $\alpha_\phi, \beta_\phi, \underline{v}, \bar{v}, v_* \in \mathbb{R}^{n_\phi}$ be given with $\alpha_\phi \leq \beta_\phi$, $\underline{v} \leq v_* \leq \bar{v}$, and $w_* := \phi(v_*)$. Assume ϕ satisfies the offset local sector $[\alpha_\phi, \beta_\phi]$ around the point (v_*, w_*) element-wise for all $v_\phi \in [\underline{v}, \bar{v}]$. If $\lambda \in \mathbb{R}^{n_\phi}$ with $\lambda \geq 0$ then:

$$\begin{bmatrix} v_\phi - v_* \\ w_\phi - w_* \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi - v_* \\ w_\phi - w_* \end{bmatrix} \geq 0$$

$$\forall v_\phi \in [\underline{v}, \bar{v}], w_\phi = \phi(v_\phi),$$

$$\text{where } \Psi_\phi := \begin{bmatrix} \text{diag}(\beta_\phi) & -I_{n_\phi} \\ -\text{diag}(\alpha_\phi) & I_{n_\phi} \end{bmatrix} \quad (16)$$

$$M_\phi(\lambda) := \begin{bmatrix} 0_{n_\phi} & \text{diag}(\lambda) \\ \text{diag}(\lambda) & 0_{n_\phi} \end{bmatrix}. \quad (17)$$

Proof. For any $v_\phi \in \mathbb{R}^{n_\phi}$ and $w_\phi = \phi(v_\phi)$:

$$\begin{aligned} & \begin{bmatrix} v_\phi - v_* \\ w_\phi - w_* \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi - v_* \\ w_\phi - w_* \end{bmatrix} \\ &= \sum_{i=1}^{n_\phi} \lambda_i (\Delta w_i - \alpha_i \Delta v_i) \cdot (\beta_i \Delta v_i - \Delta w_i) \end{aligned}$$

where $\Delta w_i := \varphi(v_{\phi,i}) - \varphi(v_{*,i})$ and $\Delta v_i := v_{\phi,i} - v_{*,i}$. If $v_\phi \in [\underline{v}, \bar{v}]$ then each term in the sum is non-negative by the offset local sector constraints and $\lambda \geq 0$. $\square$

In order to apply the local sector and slope bounds in the stability analysis, we must first compute the bounds $\underline{v}, \bar{v} \in \mathbb{R}^{n_\phi}$ on the activation input v_ϕ . The process to compute the bounds is briefly discussed here with more details provided in [22]. Let v_*^1 be the equilibrium value at the first NN layer. Select $\underline{v}^1, \bar{v}^1 \in \mathbb{R}^{n_1}$ with $\underline{v}^1 \leq v_*^1 \leq \bar{v}^1$. The assumed bounds on v^1 can be used to compute an interval $[\underline{w}^1, \bar{w}^1]$ for the

output $w^1 = \phi^1(v^1)^\dagger$ which can then be used to compute bounds $[\underline{v}^2, \bar{v}^2]$ on the input v^2 to the next activation function.[‡] The intervals computed for w^1 and v^2 will contain their equilibrium value w_*^1 and v_*^2 . This process can be propagated through all layers of the NN to obtain the bounds $\underline{v}, \bar{v} \in \mathbb{R}^{n_\phi}$ for the activation function input v_ϕ . The remainder of the paper will assume the local sector bounds have been computed as briefly summarized in the following property.

Property 1. Let $v_* \in \mathbb{R}^{n_\phi}$ be an equilibrium value of the activation input and $v_*^1 \in \mathbb{R}^{n_1}$ be the corresponding value at the first layer. Let $\underline{v}^1, \bar{v}^1 \in \mathbb{R}^{n_1}$ with $v_*^1 \in [\underline{v}^1, \bar{v}^1]$ and their corresponding activation input bounds $\underline{v}, \bar{v}$ be given. There exist $\alpha_\phi, \beta_\phi \in \mathbb{R}^{n_\phi}$ such that ϕ satisfies the offset local sector around the point $(v_*, \phi(v_*))$ for all $v_\phi \in [\underline{v}, \bar{v}]$.

E. Lyapunov Condition

This section uses a Lyapunov function and the offset local sector to compute an inner approximation for the ROA of the feedback system of G and π . To simplify notation, the interval bound on v^1 is assumed to be symmetrical about v_*^1 , i.e. $\underline{v}^1 = 2v_*^1 - \bar{v}^1$ so that $\bar{v}^1 - v_*^1 = v_*^1 - \underline{v}^1$. This can be relaxed to handle non-symmetrical intervals with minor notational changes.

Theorem 1. Consider the feedback system of plant G in (2) and NN π in (3) with equilibrium point (x_*, u_*, v_*, w_*) satisfying (12). Let $\bar{v}^1 \in \mathbb{R}^{n_1}$, $\underline{v}^1 := 2v_*^1 - \bar{v}^1$, and $\alpha_\phi, \beta_\phi \in \mathbb{R}^{n_\phi}$ be given vectors satisfying Property 1 for the NN. Denote the i^{th} row of the first weight W^1 by W_i^1 and define matrices

$$R_V := \begin{bmatrix} I_{n_G} & 0_{n_G \times n_\phi} \\ N_{ux} & N_{uw} \end{bmatrix}, \text{ and } R_\phi := \begin{bmatrix} N_{vx} & N_{vw} \\ 0_{n_\phi \times n_G} & I_{n_\phi} \end{bmatrix}.$$

If there exists a matrix $P \in \mathbb{S}_{++}^{n_G}$, and vector $\lambda \in \mathbb{R}^{n_\phi}$ with $\lambda \geq 0$ such that

$$R_V^\top \begin{bmatrix} A_G^\top P A_G - P & A_G^\top P B_G \\ B_G^\top P A_G & B_G^\top P B_G \end{bmatrix} R_V + R_\phi^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi R_\phi < 0, \quad (20)$$

$$\begin{bmatrix} (\bar{v}_i^1 - v_{*,i}^1)^2 & W_i^1 \\ W_i^{1\top} & P \end{bmatrix} \geq 0, \quad i = 1, \dots, n_1, \quad (21)$$

then: (i) the feedback system consisting of G and π is locally stable around x_* , and (ii) the set $\mathcal{E}(P, x_*)$, defined by (1), is an inner-approximation to the ROA.

Proof. By Schur complements, (21) is equivalent to:

$$W_i^1 P^{-1} W_i^{1\top} \leq (\bar{v}_i^1 - v_{*,i}^1)^2, \quad i = 1, \dots, n_1. \quad (22)$$

[†]For example, if $\varphi(\nu) = \tanh(\nu)$ then the input bound $\nu \in [-\bar{\nu}, \bar{\nu}]$ implies the output bound $\varphi(\nu) \in [-\tanh(\bar{\nu}), \tanh(\bar{\nu})]$.

[‡]The next activation input is $v^2 := W^2 w^1 + b^2$. The largest value of the i^{th} entry of this vector is obtained by solving the following optimization:

$$\bar{v}_i^2 := \max_{w^1 \leq w^1 \leq \bar{w}^1} y^\top w^1 + b_i^2 \quad (18)$$

where $y^\top$ is the i^{th} row of W^2 . Define $c := \frac{1}{2}(\bar{w}^1 + \underline{w}^1)$ and $r := \frac{1}{2}(\bar{w}^1 - \underline{w}^1)$. The optimization can be rewritten as:

$$\bar{v}_i^2 := (y^\top c + b_i^2) + \max_{-r \leq \delta \leq r} y^\top \delta \quad (19)$$

This has the explicit solution $\bar{v}_i^2 = y^\top c + b_i^2 + \sum_{j=1}^{n_1} |y_j r_j|$. Similarly, the minimal value is $\underline{v}_i^2 = y^\top c + b_i^2 - \sum_{j=1}^{n_1} |y_j r_j|$.

It follows from Lemma 1 in [23] that:

$$\mathcal{E}(P, x_*) \subseteq \{x \in \mathbb{R}^{n_G} : \underline{v}^1 - v_*^1 \leq W^1(x - x_*) \leq \bar{v}^1 - v_*^1\}.$$

Finally, use $v^1 - v_*^1 = W^1(x - x_*)$ to rewrite this as:

$$\mathcal{E}(P, x_*) \subseteq \{x : \underline{v}^1 \leq v^1 \leq \bar{v}^1\}.$$

To summarize, feasibility of (21) verifies that if $x(k) \in \mathcal{E}(P, x_*)$ then $v^1(k) \in [\underline{v}^1, \bar{v}^1]$ and hence the offset local sector conditions are valid.

Next, since the LMI in (20) is strict, there exists $\epsilon > 0$ such that left / right multiplication of the LMI by $[(x(k) - x_*)^\top \quad (w_\phi(k) - w_*)^\top]$ and its transpose yields

$$\begin{bmatrix} \star \end{bmatrix}^\top \begin{bmatrix} A_G^\top P A_G - P & A_G^\top P B_G \\ B_G^\top P A_G & B_G^\top P B_G \end{bmatrix} \begin{bmatrix} x(k) - x_* \\ u(k) - u_* \end{bmatrix} + \begin{bmatrix} \star \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi(k) - v_* \\ w_\phi(k) - w_* \end{bmatrix} \leq -\epsilon \|x(k) - x_*\|^2.$$

where the entries denoted by $\star$ can be inferred from symmetry. Define the Lyapunov function $V(x) := (x - x_*)^\top P(x - x_*)$ and use (2) and (12) to show:

$$V(x(k+1)) - V(x(k)) + \begin{bmatrix} \star \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi(k) - v_* \\ w_\phi(k) - w_* \end{bmatrix} \leq -\epsilon \|x(k) - x_*\|^2. \quad (23)$$

Assume $x(k) \in \mathcal{E}(P, x_*)$ for some $k \geq 0$, i.e., $V(x(k)) \leq 1$. As noted above, $x(k) \in \mathcal{E}(P, x_*)$ implies the offset local sector $[\alpha_\phi, \beta_\phi]$ around v_* . Then, by Lemma 1, the final term on the left side of (23) is ≥ 0 , and thus from (23) we have $V(x(k+1)) \leq 1$, i.e., $x(k+1) \in \mathcal{E}(P, x_*)$. By induction, we have that $\mathcal{E}(P, x_*)$ is forward invariant, i.e., $x(0) \in \mathcal{E}(P, x_*) \implies x(k) \in \mathcal{E}(P, x_*) \forall k \geq 0$. As a result, if $x(0) \in \mathcal{E}(P, x_*)$, then the final term on the left side of (23) is ≥ 0 for all $k \geq 0$, and $V(x(k+1)) - V(x(k)) \leq -\epsilon \|x(k) - x_*\|^2$ for all $k \geq 0$. It follows from a Lyapunov argument, e.g. Theorem 4.1 in [7], that x_* is an asymptotically stable equilibrium point and $\mathcal{E}(P, x_*)$ is an inner approximation of the ROA. $\square$

Remark 1. Note that $\bar{v}^1$ should be chosen with care as it affects the size of ROA inner-approximations directly: decreasing $(\bar{v}^1 - v_*^1)$ gives rise to sharper local sector bounds, which is beneficial on ROA estimation, but also restricts the region where ROA inner-approximations lie in; increasing $(\bar{v}^1 - v_*^1)$ leads to a larger region that contains ROA inner-approximations, but also provides looser local sector bounds. A possible way of choosing $\bar{v}^1$ is to parameterize $(\bar{v}^1 - v_*^1)$ as $\bar{v}^1 - v_*^1 = \delta_v \times 1_{n_1 \times 1}$ with $\delta_v \in \mathbb{R}_{++}$, grid the interval $[0, \bar{\delta}_v]$ [§] where $\bar{\delta}_v$ lies in, inner-approximate the ROA on the grid, and choose δ_v that leads to the largest inner-approximation.

Remark 2. In the paper, the NN controller is assumed to be state-feedback. For the output-feedback case, i.e., $u = \pi(Cx)$, where $C \in \mathbb{R}^{n_y \times n_G}$, the stability analysis can be performed similarly, using a new N_{vx} defined as $N_{vx} := \begin{bmatrix} W^1 C \\ 0_{(n_2 + \dots + n_\ell) \times n_G} \end{bmatrix}$.

[§] $\bar{\delta}_v$ is the largest value such that (20) and (21) stay feasible.

III. ROBUST STABILITY ANALYSIS

A. Problem Formulation

Consider the uncertain feedback system in Figure 5, consisting of an uncertain plant $F_u(G, \Delta)$ and a NN controller π as defined by (3). The uncertain plant $F_u(G, \Delta)$ is an interconnection of a nominal plant G and a perturbation Δ . The nominal plant G is defined by the following equations:

$$x(k+1) = A_G x(k) + B_{G1} q(k) + B_{G2} u(k), \quad (24a)$$

$$p(k) = C_G x(k) + D_{G1} q(k) + D_{G2} u(k), \quad (24b)$$

where $x(k) \in \mathbb{R}^{n_G}$ is the state, $u(k) \in \mathbb{R}^{n_u}$ is the control input, $p(k) \in \mathbb{R}^{n_p}$ and $q(k) \in \mathbb{R}^{n_q}$ are the input and output of Δ , $A_G \in \mathbb{R}^{n_G \times n_G}$, $B_{G1} \in \mathbb{R}^{n_G \times n_q}$, $B_{G2} \in \mathbb{R}^{n_G \times n_u}$, $C_G \in \mathbb{R}^{n_p \times n_G}$, $D_{G1} \in \mathbb{R}^{n_p \times n_q}$, and $D_{G2} \in \mathbb{R}^{n_p \times n_u}$. The perturbation is a bounded, causal operator $\Delta : \ell_{2e}^{n_p} \rightarrow \ell_{2e}^{n_q}$. The nominal plant G and perturbation Δ form the interconnection $F_u(G, \Delta)$ through the constraint

$$q(\cdot) = \Delta(p(\cdot)). \quad (25)$$

Denote the set of perturbations to be considered as $\mathcal{S}$.




Fig. 5: Feedback system with uncertain plant $F_u(G, \Delta)$ and NN controller π

Assumption 1. In this section, we assume (i) the equilibrium point $(x_*, u_*, v_*, w_*, p_*, q_*)$ of the feedback system is at the origin, and (ii) $0 = \Delta(0)$ for all $\Delta \in \mathcal{S}$. Note that Δ is modeled as an operator mapping inputs to outputs. If Δ has an internal state then there is an implicit assumption that it has zero initial condition.

Let $\chi(k; x_0, \Delta)$ denote the solution to the feedback system of $F_u(G, \Delta)$ and π with $\Delta \in \mathcal{S}$ at time k from the initial condition $x(0) = x_0$ [¶]. Define the robust ROA associated with x_* as follows.

Definition 5. The robust ROA of the feedback system with the uncertain plant $F_u(G, \Delta)$ and NN π is defined as:

$$\mathcal{R} := \{x_0 \in \mathbb{R}^{n_G} : \lim_{k \rightarrow \infty} \chi(k; x_0, \Delta) = x_* \forall \Delta \in \mathcal{S}\}. \quad (26)$$

The objective is to prove the uncertain feedback system is asymptotically stable and, if so, to find the largest estimate of the robust ROA using an ellipsoidal inner approximation.

[¶]An input/output model is used for the perturbation Δ so that its internal state and initial condition is not explicitly considered.

B. Integral Quadratic Constraints

The perturbation can represent various types of uncertainty [14], [15], including saturation, time delay, unmodeled dynamics, and slope-restricted nonlinearities. The input-output relationship of Δ is characterized with an integral quadratic constraint (IQC), which consists of a ‘virtual’ filter Ψ_Δ applied to the input p and output q of Δ and a constraint on the output r of Ψ_Δ . The filter Ψ_Δ is an LTI system of the form:

$$\psi(k+1) = A_\Psi \psi(k) + B_{\Psi1} p(k) + B_{\Psi2} q(k) \quad (27a)$$

$$r(k) = C_\Psi \psi(k) + D_{\Psi1} p(k) + D_{\Psi2} q(k) \quad (27b)$$

$$\psi(0) = 0 \quad (27c)$$

where $\psi(k) \in \mathbb{R}^{n_\psi}$ is the state, $r(k) \in \mathbb{R}^{n_r}$ is the output, and A_Ψ is a Schur matrix. The state matrices have compatible dimensions. The dynamics of Ψ_Δ can be compactly denoted by $\begin{bmatrix} A_\Psi & B_{\Psi1} & B_{\Psi2} \\ C_\Psi & D_{\Psi1} & D_{\Psi2} \end{bmatrix}$. By $(p_*, q_*) = 0$ from Assumption 1, the equilibrium state $\psi_* \in \mathbb{R}^{n_\psi}$ of (27) is also zero.

The Lyapunov analysis in the next subsection makes use of time-domain IQCs as defined next:

Definition 6. Let $\Psi_\Delta \in \mathbb{RH}_\infty^{n_r \times (n_p + n_q)}$ and $M_\Delta \in \mathbb{S}^{n_r}$ be given. A bounded, causal operator $\Delta : \ell_{2e}^{n_p} \rightarrow \ell_{2e}^{n_q}$ satisfies the time domain IQC defined by (Ψ_Δ, M_Δ) if the following inequality holds for all $p \in \ell_{2e}^{n_p}$, $q = \Delta(p)$ and for all $N \geq 0$

$$\sum_{k=0}^N r(k)^\top M_\Delta r(k) \geq 0. \quad (28)$$

The notation $\Delta \in \text{IQC}(\Psi_\Delta, M_\Delta)$ indicates that Δ satisfies the IQC defined by Ψ_Δ and M_Δ . Therefore, the precise relation (25), for analysis, is replaced by the constraint (28) on r . The QC proposed in Lemma 1 is a special instance of a time-domain IQCs. Specifically, Lemma 1 defines a QC that holds at each time step k and hence the inequality also holds summing over any finite horizons. This is referred to as the offset local sector IQC.

The time-domain IQCs, as defined here, hold on any finite horizon $N \geq 0$. These are typically called ‘hard IQCs’ [14]. IQCs can also be defined in the frequency domain and equivalently expressed as time-domain constraints over an infinite horizon ($N = \infty$). These are called soft IQCs. Although this paper focuses on the use of hard IQCs, it is possible to also incorporate soft IQCs [20], [21], [24], [25].

C. Lyapunov Condition

Let $\zeta := \begin{bmatrix} x \\ \psi \end{bmatrix} \in \mathbb{R}^{n_\zeta}$ define the extended state vector, $n_\zeta = n_G + n_\psi$, whose dynamics are

$$\zeta(k+1) = \mathcal{A} \zeta(k) + \mathcal{B} \begin{bmatrix} q(k) \\ u(k) \end{bmatrix} \quad (29a)$$

$$r(k) = \mathcal{C} \zeta(k) + \mathcal{D} \begin{bmatrix} q(k) \\ u(k) \end{bmatrix} \quad (29b)$$

$$u(k) = \pi(x(k)) \quad (29c)$$

where the state-space matrices are

$$\mathcal{A} = \begin{bmatrix} A_G & 0 \\ B_{\Psi1} C_G & A_\Psi \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} B_{G1} & B_{G2} \\ B_{\Psi1} D_{G1} + B_{\Psi2} & B_{\Psi1} D_{G2} \end{bmatrix},$$

$$\mathcal{C} = \begin{bmatrix} D_{\Psi1} C_G & C_\Psi \end{bmatrix}, \quad \mathcal{D} = \begin{bmatrix} D_{\Psi1} D_{G1} + D_{\Psi2} & D_{\Psi1} D_{G2} \end{bmatrix}.$$

Let $\zeta_* := \begin{bmatrix} x_* \\ \psi_* \end{bmatrix} = 0$ define the equilibrium point of the extended system (29). Since IQCs implicitly constrain the input p of the extended system (29), the response of the extended system subject to IQCs “covers” the behaviors of the original uncertain feedback system. The following theorem provides a method for inner-approximating the robust ROA by performing analysis on the extended system subject to IQCs.

Theorem 2. Consider the feedback system of an uncertain plant $F_u(G, \Delta)$ in (24)–(25), and the NN π in (3) with zero equilibrium point $(\zeta_*, u_*, v_*, w_*, p_*, q_*)$. Assume $\Delta \in \text{IQC}(\Psi_\Delta, M_\Delta)$ with Ψ_Δ and M_Δ given. Let $\bar{v}^1 \in \mathbb{R}^{n_1}$, $\underline{v}^1 := -\bar{v}^1$, and $\alpha_\phi, \beta_\phi \in \mathbb{R}^{n_\phi}$ be given vectors satisfying Property 1 for the NN, and define matrices

$$R_V = \begin{bmatrix} I_{n_\zeta} & 0 & 0 \\ 0 & 0 & I_{n_q} \\ N_{u\zeta} & N_{uw} & 0 \end{bmatrix}, \quad N_{u\zeta} = [N_{ux}, 0_{n_u \times n_\psi}],$$

$$R_\phi = \begin{bmatrix} N_{v\zeta} & N_{vw} & 0 \\ 0 & I_{n_\phi} & 0 \end{bmatrix}, \quad N_{v\zeta} = [N_{vx}, 0_{n_\phi \times n_\psi}],$$

$$\mathcal{W}_i^1 = \begin{bmatrix} W_i^1 & 0_{1 \times n_\psi} \end{bmatrix}, \quad W_i^1 \text{ is the } i^{\text{th}} \text{ row of } W^1.$$

If there exists a matrix $P \in \mathbb{S}_{++}^{n_\zeta}$, and vector $\lambda \in \mathbb{R}^{n_\phi}$ with $\lambda \geq 0$ such that

$$R_V^\top \begin{bmatrix} \mathcal{A}^\top P \mathcal{A} - P & \mathcal{A}^\top P \mathcal{B} \\ \mathcal{B}^\top P \mathcal{A} & \mathcal{B}^\top P \mathcal{B} \end{bmatrix} R_V + R_\phi^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi R_\phi$$

$$+ R_V^\top [\mathcal{C} \quad \mathcal{D}]^\top M_\Delta [\mathcal{C} \quad \mathcal{D}] R_V < 0, \quad (30a)$$

$$\begin{bmatrix} (\bar{v}_i^1)^2 & \mathcal{W}_i^{1\top} \\ \mathcal{W}_i^{1\top} & P \end{bmatrix} \geq 0, \quad i = 1, \dots, n_1, \quad (30b)$$

then: (i) the feedback system comprising $F_u(G, \Delta)$ and π is locally stable around x_* for any $\Delta \in \text{IQC}(\Psi_\Delta, M_\Delta)$, and (ii) the intersection of $\mathcal{E}(P, \zeta_*)$ with the hyperplane $\psi = 0$, i.e. $\mathcal{E}(P_x, x_*)$ where $P_x \in \mathbb{R}^{n_G \times n_G}$ is the upper left block of P , is an inner-approximation to the robust ROA.

Proof. As in the proof of Theorem 1, feasibility of (30b) implies that if $\zeta(k) \in \mathcal{E}(P, \zeta_*)$ then $v^1(k) \in [\underline{v}^1, \bar{v}^1]$ and hence the offset local sectors conditions are valid. Since the LMI in (30a) is strict, there exists $\epsilon > 0$ such that left/right multiplication of the LMI by $[(\zeta(k) - \zeta_*)^\top, (w_\phi(k) - w_*)^\top, (q(k) - q_*)^\top]$ and its transpose yields:

$$\begin{bmatrix} \star \end{bmatrix}^\top \begin{bmatrix} \mathcal{A}^\top P \mathcal{A} - P & \mathcal{A}^\top P \mathcal{B} \\ \mathcal{B}^\top P \mathcal{A} & \mathcal{B}^\top P \mathcal{B} \end{bmatrix} \begin{bmatrix} \zeta(k) - \zeta_* \\ q(k) - q_* \\ u(k) - u_* \end{bmatrix}$$

$$+ \begin{bmatrix} \star \end{bmatrix}^\top [\mathcal{C} \quad \mathcal{D}]^\top M_\Delta [\mathcal{C} \quad \mathcal{D}] \begin{bmatrix} \zeta(k) - \zeta_* \\ q(k) - q_* \\ u(k) - u_* \end{bmatrix}$$

$$+ \begin{bmatrix} \star \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi(k) - v_* \\ w_\phi(k) - w_* \end{bmatrix} \leq -\epsilon \|\zeta(k) - \zeta_*\|^2,$$

Define the Lyapunov function $V(\zeta) := (\zeta - \zeta_*)^\top P (\zeta - \zeta_*)$, and use (29) to show:

$$V(\zeta(k+1)) - V(\zeta(k)) + r(k)^\top M_\Delta r(k)$$

$$+ \begin{bmatrix} \star \end{bmatrix}^\top \Psi_\phi^\top M_\phi(\lambda) \Psi_\phi \begin{bmatrix} v_\phi(k) - v_* \\ w_\phi(k) - w_* \end{bmatrix} \leq -\epsilon \|\zeta(k) - \zeta_*\|^2.$$

Sum this inequality from $k = 0$ to any finite time $N \geq 0$. The third and fourth term on the left side will be ≥ 0 by the local sector conditions and the IQC. This yields:

$$V(\zeta(N+1)) - V(\zeta(0)) \leq - \sum_{k=0}^N \epsilon \|\zeta(k) - \zeta_*\|^2.$$

Thus if $\zeta(0) \in \mathcal{E}(P, \zeta_*)$ then $\zeta(k) \in \mathcal{E}(P, \zeta_*)$ for all $k \geq 0$. Moreover, this inequality implies that $\zeta(N) \rightarrow \zeta_*$ as $N \rightarrow \infty$. The initial condition for the virtual filter is $\psi(0) = 0$ so that $\zeta(0) \in \mathcal{E}(P, \zeta_*)$ is equivalent to $x(0) \in \mathcal{E}(P_x, x_*)$. Hence $\mathcal{E}(P_x, x_*)$ is an inner approximation for the ROA. $\square$

For a particular perturbation Δ there is typically a class of valid time-domain IQCs defined by a fixed filter Ψ_Δ and a matrix M_Δ drawn from a constraint set $\mathcal{M}_\Delta$. Therefore when formulating an optimization problem, along with P and λ , we can treat $M_\Delta \in \mathcal{M}_\Delta$ as an additional decision variable to reduce conservatism. In this paper, the set $\mathcal{M}_\Delta$ is restricted to one that is described by LMIs [15]. Using $\text{trace}(P_x)$ as the cost function to minimize along with the LMIs developed before, we have the following optimization to compute the “largest” ROA inner-approximation:

$$\min_{P \in \mathbb{S}_{++}^{n_\zeta}, \lambda \geq 0, M_\Delta \in \mathcal{M}_\Delta} \text{trace}(P_x) \text{ s.t. (30a) – (30b) hold,} \quad (31)$$

which is convex in (P, λ, M_Δ) . The strict inequality in (30a) can be enforced by either replacing < 0 with $\leq -\epsilon I$ with $\epsilon = 1 \times 10^{-6}$, or solving (31) with a non-strict inequality ≤ 0 , and checking if the constraint is active afterwards.

D. IQCs for Combined Activation Functions ϕ

Now that we have the general framework that merges Lyapunov theory with IQCs, we will revisit the problem of describing the activation functions ϕ using more general tools. Recall that offset local sector IQCs have been used in Sections II and III to bound activation functions ϕ . However, these local sectors fail to incorporate slope bounds of ϕ . In this subsection, in addition to local sectors, we will use off-by-one IQCs [16] to capture the slope information of ϕ to achieve less conservative ROA inner-approximations.

Besides the local sector bound α_ϕ, β_ϕ , the bounds $\underline{v}, \bar{v}$ on activation input v_ϕ can also be used to compute the local slope bounds $[m_\phi, L_\phi]$ of ϕ , with $m_\phi, L_\phi \in \mathbb{R}^{n_\phi}$. For example, $\phi_i(v_{\phi,i}) := \tanh(v_{\phi,i})$ restricted to the interval $[\underline{v}_i, \bar{v}_i]$ for $i = 1, \dots, n_\phi$ satisfies the local slope bound $[m_{\phi,i}, L_{\phi,i}]$ with $m_{\phi,i} := \min \left(\frac{d \tanh(v_i)}{dv_i} \Big|_{v_i=\underline{v}_i}, \frac{d \tanh(v_i)}{dv_i} \Big|_{v_i=\bar{v}_i} \right)$, and $L_{\phi,i} := 1$. If $w_\phi = \phi(v_\phi)$ with $v_\phi(k) \in [\underline{v}, \bar{v}]$, then ϕ also satisfies the hard IQC defined by $(\Psi_{\text{off}}, M_{\text{off}})$, where

$$\Psi_{\text{off}} := \begin{bmatrix} 0_{n_\phi} & -\text{diag}(L_\phi) & I_{n_\phi} \\ I_{n_\phi} & \text{diag}(L_\phi) & -I_{n_\phi} \\ 0_{n_\phi} & -\text{diag}(m_\phi) & I_{n_\phi} \end{bmatrix},$$

$$M_{\text{off}}(\eta) := \begin{bmatrix} 0_{n_\phi} & \text{diag}(\eta) \\ \text{diag}(\eta) & 0_{n_\phi} \end{bmatrix}, \text{ for all } \eta \in \mathbb{R}^{n_\phi} \text{ with } \eta \geq 0.$$

This is the so-called “off-by-one” IQC [16], which is a special instance of the Zames-Falb IQC [26], [27]. It provides constraints that relate the activation at different time instances, e.g. between $\phi_i(v_{\phi,i}(k))$ and $\phi_i(v_{\phi,i}(k+1))$ for any $i = 1, \dots, n_\phi$.

The analysis on the feedback system of $F_u(G, \Delta)$ and π can be instead performed on the extended system made up by G , Ψ_Δ and Ψ_{off} with additional constraints that $\Delta \in \text{IQC}(\Psi_\Delta, M_\Delta)$, and ϕ satisfies the offset local sector and $\phi \in \text{IQC}(\Psi_{\text{off}}, M_{\text{off}})$. However, since Ψ_{off} introduces a number of n_ϕ states to the extended system, the size of the corresponding Lyapunov matrix P will increase from $\mathbb{S}_{++}^{n_x+n_\psi}$ to $\mathbb{S}_{++}^{n_x+n_\psi+n_\phi}$, which leads to longer computation time. The effectiveness of the off-by-one IQC is demonstrated in Section IV-B.

IV. EXAMPLES

In the following examples, the optimization (31) is solved using MOSEK with CVX. ^{||}

A. Inverted pendulum

Consider the nonlinear inverted pendulum example with mass $m = 0.15$ kg, length $l = 0.5$ m, and friction coefficient $\mu = 0.5$ Nms/rad. The dynamics are:

$$\ddot{\theta}(t) = \frac{mgl \sin(\theta(t)) - \mu \dot{\theta}(t) + \text{sat}(u(t))}{ml^2}, \quad (32)$$

where θ is the angular position (rad) and u is the control input (Nm). The plant state is $x = [\theta, \dot{\theta}]$. The saturation function is defined as $\text{sat}(u) = \text{sgn}(u) \min(|u|, u_{\max})$, with $u_{\max} = 0.7$ Nm. The controller π is obtained through a reinforcement learning process using policy gradient [3]–[5]. During training, the agent decision making process is characterized by a probability: $u(k) \sim \text{Pr}(u(k) = u \mid x(k) = x)$ for all $u \in \mathbb{R}$ and $x \in \mathbb{R}^2$ where the probability is a Gaussian distribution with mean $\pi(x(k))$, and standard deviation σ . After training, the policy mean π is used as the deterministic controller $u(k) = \pi(x(k))$. The controller π is parameterized by a 2-layer, feedforward NN with $n_1 = n_2 = 32$ and tanh as the activation function for both layers. The biases in the NN are set to zero during training to ensure that the equilibrium point is $x_* = 0$ and $u_* = 0$. The dynamics used for training are the discretized version of (32) with sampling time $dt = 0.02$ s.

We rearrange (32) into the form:

$$\ddot{\theta}(t) = \frac{-mglq(t) + mgl\theta(t) - \mu\dot{\theta}(t) + \text{sat}(u(t))}{ml^2}, \quad (33a)$$

$$q(t) = \Delta(\theta(t)) := \theta(t) - \sin(\theta(t)). \quad (33b)$$

The static nonlinearity $\Delta(\theta) = \theta - \sin(\theta)$ is slope-restricted, and sector bounded. If we assume that $\theta(k) \in [\underline{\theta}, \bar{\theta}]$ with $\bar{\theta} = -\underline{\theta} = 0.73$, then the nonlinearity is slope-restricted in $[0, 0.2548]$, and sector bounded in $[0, 0.087]$. We also assume that $v^1 \in [\underline{v}^1, \bar{v}^1]$ with $\bar{v}^1 = -\underline{v}^1 = \delta_v \times 1_{32 \times 1}$ using $\delta_v = 0.1$. Both assumptions are verified using the ROA inner-approximation. Two types of IQCs are used to characterize the nonlinearity $\Delta(\cdot)$: an off-by-one IQC to capture the slope information, and a local sector IQC to express the local sector bound. Only the local sector IQC is used to characterize the activation functions ϕ . The saturation function is static and can also be described using a local sector bound. Let $\bar{u}$ be the largest possible control command from π induced from the

assumption that $v^1 \in [\underline{v}^1, \bar{v}^1]$. Then the saturation function satisfies the local sector $[\alpha, \beta]$, where $\alpha := \frac{u_{\max}}{\bar{u}}$ and $\beta := 1$.

Figure 6 shows the boundaries for the sets $\{x : \underline{v}^1 \leq v^1 \leq \bar{v}^1\}$ and $\{x : \underline{\theta} \leq \theta \leq \bar{\theta}\}$ with orange and brown lines, the ROA inner-approximation with a blue ellipsoid, and the phase portrait of the closed-loop system, with green and red curves representing trajectories inside and outside the ROA.




Fig. 6: A ROA inner-approximation of the inverted pendulum

B. Vehicle lateral control

Consider the vehicle lateral dynamics from [28]:

$$\begin{bmatrix} \dot{e} \\ \dot{e}_\theta \\ \dot{e}_\theta \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & \frac{C_{\alpha f} + C_{\alpha r}}{mU} & -\frac{C_{\alpha f} + C_{\alpha r}}{m} & \frac{aC_{\alpha f} - bC_{\alpha r}}{mU} \\ 0 & 0 & 0 & 1 \\ 0 & \frac{aC_{\alpha f} - bC_{\alpha r}}{I_z U} & -\frac{aC_{\alpha f} - bC_{\alpha r}}{I_z} & \frac{a^2 C_{\alpha f} + b^2 C_{\alpha r}}{I_z U} \end{bmatrix} \begin{bmatrix} e \\ \dot{e} \\ e_\theta \\ \dot{e}_\theta \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{C_{\alpha f}}{m} \\ 0 \\ -\frac{aC_{\alpha f}}{I_z} \end{bmatrix} u + \begin{bmatrix} 0 \\ \frac{aC_{\alpha f} - bC_{\alpha r}}{m} - U^2 \\ 0 \\ \frac{a^2 C_{\alpha f} + b^2 C_{\alpha r}}{I_z} \end{bmatrix} c \quad (34)$$

where e is the perpendicular distance to the lane edge (m), and e_θ is the angle between the tangent to the straight section of the road and the projection of the vehicle's longitudinal axis (rad). Let $x = [e, \dot{e}, e_\theta, \dot{e}_\theta]^\top$ denote the plant state. The control u is the steering angle of the front wheel (rad), the disturbance c is the road curvature (1/m), and the parameters are: longitudinal velocity $U = 28$ m/s, front cornering stiffness $C_{\alpha f} = -1.232 \times 10^5$ N/rad, rear cornering stiffness $C_{\alpha r} = -1.042 \times 10^5$ N/rad, mass $m = 1.67 \times 10^3$ kg, moment of inertia $I_z = 2.1 \times 10^3$ kg/m², distances from vehicle center of gravity to front axle $a = 0.99$ m and rear axle $b = 1.7$ m.

Again, the controller π is obtained using policy gradient, and is parameterized by a 2-layer, feedforward NN, with $n_1 = n_2 = 32$ and tanh as the activation function for both layers. The training process uses a discretized version of (34) with sampling time $dt = 0.02$ s and draws the curvature $c(k)$ at each time step from an interval $[-1/200, 1/200]$. The control command derived from $u(k) = \pi(x(k))$ enters the vehicle dynamics through a saturation function $\text{sat}(\cdot)$ with $u_{\max} = \pi/6$. Let $u_{\text{sat}} := \text{sat}(\pi(x))$ define the saturated control signal.

The analysis is performed for a constant curvature $c \equiv 0$, resulting in a zero equilibrium state $x_* = 0$. In the analysis problem, on top of saturation, we also add a norm-bounded LTI uncertainty $\Delta_{\text{LTI}} \in \mathbb{RH}_\infty$ with $\|\Delta_{\text{LTI}}\|_\infty \leq 0.1$ to the control input. This is used to assess the robustness of the NN

^{||}The code is available at <https://github.com/heyinUCB/Stability-Analysis-using-Quadratic-Constraints-for-Systems-with-Neural-Network-Controllers>.

controller against actuator uncertainty. As shown in Figure 7, the actual input to the vehicle dynamics is

$$u_{\text{pert}}(k) = u_{\text{sat}}(k) + q(k), \text{ and } q(\cdot) = \Delta_{\text{LTI}}(u_{\text{sat}}(\cdot)).$$




Fig. 7: Uncertain vehicle system with actuator uncertainty

It is assumed that $v^1 \in [\underline{v}^1, \bar{v}^1]$, where $\bar{v}^1 = -\underline{v}^1 = \delta_v \times 1_{32 \times 1}$ with $\delta_v = 0.6$. To show effectiveness of the off-by-one IQC, two experiments were carried out: one with only local sector IQC to describe ϕ , and one with both local sector and off-by-one IQCs. The achieved $\text{trace}(P_x)$ for the two experiments are 4.4 and 2.9, respectively. Moreover, the achieved $\det(P_x^{-1})$ (proportional to the volume) for the experiments are 3.2×10^5 , and 1.1×10^6 , respectively. Therefore, with the help of off-by-one IQC to sharpen the description of ϕ , the second experiment achieves a larger ROA inner-approximation. It is also important to note that thanks to the off-by-one IQC, the SDP is able to tolerate looser local sector bounds. The largest value of δ_v such that the SDP is feasible is 0.67 for the first experiment, and 1.4 for the second experiment.

Figure 8 shows slices of the ROA inner-approximation from the second experiment on the $e-\dot{e}$ and $e_\theta-\dot{e}_\theta$ spaces. Specifically, these are intersections of $\mathcal{E}(P_x, x_*)$ with the hyperplanes $(e_\theta, \dot{e}_\theta) = (e_{\theta*}, \dot{e}_{\theta*})$ and $(e, \dot{e}) = (e_*, \dot{e}_*)$, respectively, where $x_* = [e_*, \dot{e}_*, e_{\theta*}, \dot{e}_{\theta*}]^\top$. The slices are shown with blue ellipsoids. The boundary of the polytopic set $\{x : \underline{v}^1 \leq v^1 \leq \bar{v}^1\}$ is shown with the orange lines. The brown crosses represent the zero equilibrium state x_* .




Fig. 8: ROA inner-approximation on the $e-\dot{e}$ and $e_\theta-\dot{e}_\theta$ spaces using both local sector and off-by-one IQCs with $\delta_v = 0.6$.

REFERENCES

- [1] A. U. Levin and K. S. Narendra, "Control of nonlinear dynamical systems using neural networks: controllability and stabilization," *IEEE Transactions on Neural Networks*, vol. 4, no. 2, pp. 192–206, 1993.
- [2] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Mach. Learn.*, vol. 8, no. 3–4, p. 229–256, May 1992.
- [3] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural networks : the official journal of the International Neural Network Society*, vol. 21 4, pp. 682–97, 2008.
- [4] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust Region Policy Optimization," p. arXiv:1502.05477, Feb. 2015.
- [5] S. Kakade, "A natural policy gradient," in *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, ser. NIPS'01. Cambridge, MA, USA: MIT Press, 2001, p. 1531–1538.
- [6] D. A. Pomerleau, *ALVINN: An Autonomous Land Vehicle in a Neural Network*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1989, p. 305–313.
- [7] H. K. Khalil, *Nonlinear systems; 3rd ed.* Upper Saddle River, NJ: Prentice-Hall, 2002.
- [8] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety Verification and Robustness Analysis of Neural Networks via Quadratic Constraints and Semidefinite Programming," p. arXiv:1903.01287, Mar. 2019.
- [9] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of Lipschitz constants for deep neural networks," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 11 427–11 438.
- [10] P. Pauli, A. Koch, J. Berberich, and F. Allgöwer, "Training robust neural networks using Lipschitz bounds," p. arXiv:2005.02929, May 2020.
- [11] H. Hu, M. Fazlyab, M. Morari, and G. J. Pappas, "Reach-SDP: Reachability Analysis of Closed-Loop Systems with Neural Network Controllers via Semidefinite Programming," p. arXiv:2004.07876, Apr. 2020.
- [12] M. Jin and J. Lavaei, "Stability-certified reinforcement learning: A control-theoretic perspective," p. arXiv:1810.11505, Oct. 2018.
- [13] K. K. Kim, E. R. Patrón, and R. D. Braatz, "Standard representation and unified stability analysis for dynamic artificial neural network models," *Neural Networks*, vol. 98, pp. 251–262, 2018.
- [14] A. Megretski and A. Rantzer, "System analysis via integral quadratic constraints," *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 819–830, June 1997.
- [15] J. Veenman, C. W. Scherer, and H. Köroğlu, "Robust stability and performance analysis based on integral quadratic constraints," *European Journal of Control*, vol. 31, pp. 1 – 32, 2016.
- [16] L. Lessard, B. Recht, and A. Packard, "Analysis and design of optimization algorithms via integral quadratic constraints," *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 57–95, 2016.
- [17] H. Fang, Z. Lin, and M. Rotea, "On iqc approach to the analysis and design of linear systems subject to actuator saturation," *Systems & Control Letters*, vol. 57, no. 8, pp. 611 – 619, 2008.
- [18] S. Tarbouriech, *Stability and Stabilization of Linear Systems with Saturating Actuators*. Springer Publishing Company, Incorporated, 2014.
- [19] E. Summers and A. Packard, "L2 gain verification for interconnections of locally stable systems using integral quadratic constraints," in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 1460–1465.
- [20] M. Fetzner, C. W. Scherer, and J. Veenman, "Invariance with dynamic multipliers," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1929–1942, 2018.
- [21] A. Iannelli, P. Seiler, and A. Marcos, "Region of attraction analysis with integral quadratic constraints," *Automatica*, vol. 109, p. 108543, 2019.
- [22] S. Goyal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, R. Arandjelovic, T. Mann, and P. Kohli, "On the Effectiveness of Interval Bound Propagation for Training Verifiably Robust Models," p. arXiv:1810.12715, Oct. 2018.
- [23] H. Hindi and S. Boyd, "Analysis of linear systems with saturation using convex optimization," in *Proceedings of the 37th IEEE Conference on Decision and Control*, vol. 1, 1998, pp. 903–908 vol.1.
- [24] H. Yin, P. Seiler, and M. Arcak, "Backward reachability using integral quadratic constraints for uncertain nonlinear systems," *IEEE Control Systems Letters*, vol. 5, no. 2, pp. 707–712, 2021.
- [25] H. Yin, A. Packard, M. Arcak, and P. Seiler, "Reachability analysis using dissipation inequalities for uncertain nonlinear systems," *Systems & Control Letters*, vol. 142, p. 104736, 2020.
- [26] G. Zames and P. L. Falb, "Stability conditions for systems with monotone and slope-restricted nonlinearities," *SIAM Journal on Control*, vol. 6, no. 1, pp. 89–108, 1968.
- [27] V. V. Kulkarni and M. G. Safonov, "All multipliers for repeated monotone nonlinearities," *IEEE Transactions on Automatic Control*, vol. 47, no. 7, pp. 1209–1212, 2002.
- [28] A. Alleyne, "A comparison of alternative intervention strategies for unintended roadway departure (urd) control," *Vehicle System Dynamics*, vol. 27, no. 3, pp. 157–186, 1997.