

Offset-free setpoint tracking using neural network controllers

Patricia Pauli

PATRICIA.PAULI@IST.UNI-STUTTGART.DE

Institute for Systems Theory and Automatic Control, University of Stuttgart

Johannes Köhler

JKOEHL@ETHZ.CH

Institute for Dynamic Systems and Control, ETH Zurich

Julian Berberich

JULIAN.BERBERICH@IST.UNI-STUTTGART.DE

Anne Koch

ANNE.KOCH@IST.UNI-STUTTGART.DE

Frank Allgöwer

FRANK.ALLGOWER@IST.UNI-STUTTGART.DE

Institute for Systems Theory and Automatic Control, University of Stuttgart

Abstract

In this paper, we present a method to analyze local and global stability in offset-free setpoint tracking using neural network controllers and we provide ellipsoidal inner approximations of the corresponding region of attraction. We consider a feedback interconnection of a linear plant in connection with a neural network controller and an integrator, which allows for offset-free tracking of a desired piecewise constant reference that enters the controller as an external input. Exploiting the fact that activation functions used in neural networks are slope-restricted, we derive linear matrix inequalities to verify stability using Lyapunov theory. After stating a global stability result, we present less conservative local stability conditions (i) for a given reference and (ii) for any reference from a certain set. The latter result even enables guaranteed tracking under setpoint changes using a reference governor which can lead to a significant increase of the region of attraction. Finally, we demonstrate the applicability of our analysis by verifying stability and offset-free tracking of a neural network controller that was trained to stabilize a linearized inverted pendulum.

Keywords: Neural networks, neural network controllers, sector bounds, reference governor

1. Introduction

Alongside with the flourishing success of deep learning methods, in recent times, there has been an increasing interest in feedback interconnections with neural network (NN) controllers. NNs as universal function approximators can learn any desired, nonlinear behavior with high precision and there are efficient ways to train them, e.g., using backpropagation and deep reinforcement learning (RL) (Sutton and Barto, 2018). NNs are useful whenever no precise model is available and in case it is expensive to obtain a model but data can be collected easily (Levin and Narendra, 1993; Psaltis et al., 1988; Ibrahim et al., 2020). NN controllers that are determined from supervised learning can approximate more complicated controllers and, in comparison to the original controller, they might save computation time (Hertneck et al., 2018). Another practically highly relevant application of NN controllers is in terms of deep RL, e.g., using inverse reinforcement learning where RL is used to imitate human behavior (Yin et al., 2020). However, even though NN controllers may work well in practice, they lack safety guarantees (Szegedy et al., 2013) which is the reason why they are typically not used in safety-critical applications such as autonomous driving.

Recently, efforts have been made to give robustness and stability guarantees for NNs and feedback interconnections with NN controllers. [Fazlyab et al. \(2019\)](#) certify Lipschitz bounds that serve as a proxy for robustness of NNs and [Pauli et al. \(2021\)](#); [Revay et al. \(2020\)](#) impose them during training. Alongside with robustness guarantees, there are recent works on stability guarantees, e.g., [Revay and Manchester \(2020\)](#); [Yin et al. \(2021\)](#) present results on stability of recurrent neural networks using incremental \mathcal{L}_2 -bounds and feedback interconnections using Lyapunov theory, respectively. The underlying property that is exploited in the works stated above is *slope-restriction*, i.e., the slopes of activation functions comply with certain bounds, which can be formulated as an incremental quadratic constraint. Lipschitz continuity and closed-loop stability, respectively, are then verified using semidefinite programming.

In this paper, we extend the stability analysis of a linear system in feedback interconnection with an NN controller in ([Yin et al., 2021](#)) to offset-free setpoint tracking with a piecewise constant reference. In comparison to simple linear feedback controllers, NN controllers may provide more flexibility, e.g., they may be used to achieve additional goals beside the mere stabilization of the closed loop using problem-specific objective functions. Note that [Yin et al. \(2021\)](#) require a stability analysis tailored to any possible steady state. In contrast, we are able to derive stability results for a whole set of references, considering a feedback interconnection with an external input and an NN controller in combination with an integrator feedback, that avoids steady-state offsets. We show that via online modification of the reference, a reference governor ensures stability in an extended region of attraction (RoA). In this paper, we give three main theorems that verify stability based on linear matrix inequalities (LMIs). First, we discuss global stability of the suggested feedback interconnection using a Lyapunov function. We proceed with a local stability condition for a fixed reference which is less restrictive than the global stability result, also analyzing the resulting RoA. Finally, we formulate a local stability condition for all references from a specified range and for this case, we show how a reference governor can be used to increase the RoA.

The paper is organized as follows: In [Section 2](#), we present the problem setting and formulate ingredients relevant to the stability analysis. Following, in [Section 3](#), we state LMI conditions for global stability and local stability (i) for a fixed constant reference and (ii) for any reference from a certain set which facilitates tracking of piecewise constant references from that set, respectively. In addition, we present an approach to extend the guaranteed RoA using a reference governor. In [Section 4](#), we provide a numerical example and we summarize our results in [Section 5](#).

2. Setpoint tracking

In this section, we state the problem setting and formulate all ingredients to verify stability in offset-free setpoint tracking using NN controllers. We use $c \in [a, b]$ for vectors c that comply componentwise with the range $[a, b]$ and $|c|$ denotes the componentwise magnitude of the entries of c .

2.1. Problem setting

We consider discrete-time linear time-invariant systems

$$\begin{aligned} \text{G:} \quad x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k \end{aligned}$$




Figure 1: Feedback interconnection of plant G and controller K , consisting of NN controller and integrator feedback k_ξ .

with state $x_k \in \mathbb{R}^{n_x}$, input $u_k \in \mathbb{R}^{n_u}$, and output $y_k \in \mathbb{R}^{n_r}$, of the same input and output dimension ($n_u = n_r$). The control goal is twofold: We aim for (i) stability of the closed-loop system and (ii) offset-free tracking of a constant reference r , i.e., $y_k \rightarrow r$ for $k \rightarrow \infty$.

The considered controller architecture in feedback interconnection with the plant G is illustrated in Fig. 1, where the reference $r \in \mathbb{R}^{n_r}$ enters as an external input. The controller K consists of two parts, (i) an NN controller and (ii) an integrator. In order to track a constant reference r , we use an l -layer feed-forward neural network $\kappa(x, r)$ as a controller whose input-output behavior $\kappa : \mathbb{R}^{n_x} \times \mathbb{R}^{n_r} \rightarrow \mathbb{R}^{n_r}$ reads

$$\begin{aligned} w_k^0 &= H_x^0 x_k + H_r^0 r \\ \text{NN : } \quad w_k^{i+1} &= \phi^{i+1}(W^i w_k^i + b^i) \\ \kappa(x_k, r) &= W^l w_k^l + b^l, \end{aligned}$$

where $w_k^i \in \mathbb{R}^{n_i}$ are the outputs of the neurons, $W^i \in \mathbb{R}^{n_{i+1} \times n_i}$ are weights, $b^i \in \mathbb{R}^{n_{i+1}}$ are biases, and $\phi^i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ is the vector of activation functions of the i -th layer for all $i = 1, \dots, l$ hidden layers. The state x_k and the reference r enter the NN controller through the matrices H_x^0 and H_r^0 . While [Yin et al. \(2021\)](#) consider only state feedback, our description of the NN controller is more general, including the two important special cases of state feedback, i.e., $w^0 = x$, $H_x^0 = I_{n_x}$, $H_r^0 = 0_{n_x \times n_r}$, and output error feedback, i.e., $w^0 = r - y$, $H_x^0 = -C$, $H_r^0 = I_{n_r \times n_r}$. The case of output error feedback is especially relevant as it requires the output measurement only, not the full state or a costly estimate thereof. Yet, state information may still be used during training of the NN to improve performance, cf. [\(Zhang et al., 2016\)](#). By adding an integrator feedback to the controller with dynamics

$$\xi_{k+1} = \xi_k + r - y_k,$$

introducing the additional state ξ_k , an undesired steady-state offset is eliminated independent of the structure of the weights and biases of the NN. The resulting input to the plant G then is $u_k = k_\xi \xi_k + \kappa(x_k, r)$. Here, $k_\xi \in \mathbb{R}^{n_r \times n_r}$ is a constant gain that is assumed to be invertible which is required for offset-free tracking.

Introducing the augmented state $\tilde{x} = [x^\top \quad \xi^\top]^\top$ yields the overall dynamics of the augmented system

$$\begin{aligned} \tilde{G} : \quad \tilde{x}_{k+1} &= \tilde{A} \tilde{x}_k + \tilde{B} \tilde{u}_k + B_r r \\ y_k &= \tilde{C} \tilde{x}_k \end{aligned} \quad (1)$$

with $\tilde{u} = \kappa(x, r)$ and the reference r as inputs, where

$$\tilde{A} = \begin{bmatrix} A & B k_\xi \\ -C & I_{n_r} \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B \\ 0 \end{bmatrix}, \quad \tilde{C} = [C \quad 0], \quad B_r = \begin{bmatrix} 0 \\ I_{n_r} \end{bmatrix}.$$

For a given constant reference r that is tracked with zero steady-state offset according to the defined control goal, the corresponding state and input steady state (x_*, u_*) satisfies

$$\begin{bmatrix} 0 \\ r \end{bmatrix} = \underbrace{\begin{bmatrix} A - I & B \\ C & 0 \end{bmatrix}}_{=: A_a} \begin{bmatrix} x_* \\ u_* \end{bmatrix}. \quad (2)$$

Throughout this paper, we assume that the matrix A_a is square and non-singular, which is a standard condition in tracking (Pannocchia and Rawlings, 2003). For any given reference r , this leads to a unique steady state (x_*, u_*) that achieves the desired output $r = Cx_*$. With the information of the steady state (x_*, u_*) , we can determine the corresponding unique steady state $\xi_* = k_\xi^{-1}(u_* - \kappa(x_*, r))$ and the stationary value w_*^i from a forward propagation of x_* and r through the NN.

With the corresponding steady state $(\tilde{x}_*, \tilde{u}_*)$ of (1), we formulate the incremental dynamics

$$\begin{aligned} \tilde{x}_{k+1} - \tilde{x}_* &= \tilde{A}(\tilde{x}_k - \tilde{x}_*) + \tilde{B}(\tilde{u}_k - \tilde{u}_*), \\ y_k - r &= \tilde{C}(\tilde{x}_k - \tilde{x}_*). \end{aligned}$$

Stated in error coordinates $e_k := \tilde{x}_k - \tilde{x}_*$, the plant dynamics are independent of the reference r , which will be exploited in the analysis. The goal of offset-free setpoint tracking is achieved if $e_k \rightarrow 0$ for $k \rightarrow \infty$.

2.2. Slope restriction

Like it was done by Fazlyab et al. (2019) and many subsequent works in the field, e.g., in (Pauli et al., 2021; Revay and Manchester, 2020; Yin et al., 2021), we exploit that activation functions in NNs are slope-restricted. This is an incremental sector condition on the individual activation functions of each of the neurons of the NN. We define $v^{i+1} = W^i w^i + b^i$ such that $w^i = \phi(v^i)$ and denote the input to the j -th neuron of the i -th layer at time k by $v_k^{i,j}$ and the steady-state input to the same neuron by $v_*^{i,j}$. Throughout this paper, we consider nonlinear activation functions $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ that are slope-restricted on the interval $[\alpha, \beta]$, i.e.,

$$\alpha \leq \frac{\varphi(v^{i,j}) - \varphi(v_*^{i,j})}{v^{i,j} - v_*^{i,j}} \leq \beta \quad \forall i \in 1, \dots, l \quad \forall j \in 1, \dots, n_i. \quad (3)$$

with $0 \leq \alpha < \beta$. This applies to the most common activation functions, e.g., to ReLU and tanh with $\alpha = 0$, $\beta = 1$. The property of slope-restriction can then be stated as an incremental quadratic constraint

$$\begin{bmatrix} v^{i,j} - v_*^{i,j} \\ \varphi(v^{i,j}) - \varphi(v_*^{i,j}) \end{bmatrix}^\top \begin{bmatrix} -2\alpha\beta & (\alpha + \beta) \\ (\alpha + \beta) & -2 \end{bmatrix} \begin{bmatrix} v^{i,j} - v_*^{i,j} \\ \varphi(v^{i,j}) - \varphi(v_*^{i,j}) \end{bmatrix} \geq 0.$$

We stack up all activations $\phi = [\phi^1, \dots, \phi^l]^\top : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with $n = \sum_{i=1}^l n_i$ as well as $v = [v^1, \dots, v^l]^\top$ and $w = [w^1, \dots, w^l]^\top$ and introduce multipliers $\lambda_i \geq 0$ to formulate the incremental quadratic constraint for all neurons

$$\begin{bmatrix} v - v_* \\ \phi(v) - \phi(v_*) \end{bmatrix}^\top \begin{bmatrix} -2\alpha\beta\Lambda & (\alpha + \beta)\Lambda \\ (\alpha + \beta)\Lambda & -2\Lambda \end{bmatrix} \begin{bmatrix} v - v_* \\ \phi(v) - \phi(v_*) \end{bmatrix} \geq 0 \quad (4)$$

with the diagonal matrix $\Lambda \in \mathbb{D}_+^n := \{\lambda_i \in \mathbb{R} \mid \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n), \lambda_i \geq 0\}$. Fazlyab et al. (2019) claimed that (4) holds for a richer class of multipliers Λ . However, as shown by counterexample in (Pauli et al., 2021), this is not true and the restriction to diagonal matrices Λ is indeed necessary.

2.3. Basis transformations

For convenience, we use a compact notation similar to [Yin et al. \(2021\)](#) by defining

$$N_0 := \begin{bmatrix} N_0^1 \\ 0 \end{bmatrix} = \begin{bmatrix} W^0 H_x^0 & 0 \\ 0 & 0 \end{bmatrix}, N_{1:l-1} := \begin{bmatrix} 0 & \dots & 0 & 0 \\ W^1 & \dots & 0 & 0 \\ 0 & \ddots & \vdots & \vdots \\ 0 & \dots & W^{l-1} & 0 \end{bmatrix}, N_l := \begin{bmatrix} 0 & W^l \end{bmatrix}. \quad (5)$$

Using these definitions, the mapping of the NN controller is characterized by

$$\begin{bmatrix} \tilde{x} - \tilde{x}_* \\ \tilde{u} - \tilde{u}_* \end{bmatrix} = \underbrace{\begin{bmatrix} I_{n_{\tilde{x}}} & 0_{n_{\tilde{x}} \times n} \\ 0_{n_u \times n_{\tilde{x}}} & N_l \end{bmatrix}}_{:=R_V} \begin{bmatrix} \tilde{x} - \tilde{x}_* \\ w - w_* \end{bmatrix}, \begin{bmatrix} v - v_* \\ w - w_* \end{bmatrix} = \underbrace{\begin{bmatrix} N_0 & N_{1:l-1} \\ 0_{n \times n_{\tilde{x}}} & I_n \end{bmatrix}}_{:=R_\phi} \begin{bmatrix} \tilde{x} - \tilde{x}_* \\ w - w_* \end{bmatrix}.$$

3. Stability analysis

In this section, we state our main results to verify stability in offset-free setpoint tracking. First, we provide a theorem to show global stability and then state two local stability results with inner approximations of the corresponding RoA. Finally, we show how an extended RoA is accessible via a reference governor.

3.1. Global stability

In this section, we state an LMI condition that verifies global stability of the feedback interconnection of \tilde{G} and $\kappa(x, r)$, with reference r as an external input. Global stability indicates that the system converges to its steady state and achieves offset-free tracking for all references and all initial conditions.

Theorem 1 *Suppose there exist matrices $\Lambda \in \mathbb{D}_+^n$, $P > 0$, such that*

$$R_V^\top \begin{bmatrix} \tilde{A}^\top P \tilde{A} - P & \tilde{A}^\top P \tilde{B} \\ \tilde{B}^\top P \tilde{A} & \tilde{B}^\top P \tilde{B} \end{bmatrix} R_V + R_\phi^\top \begin{bmatrix} -2\alpha\beta\Lambda & (\alpha + \beta)\Lambda \\ (\alpha + \beta)\Lambda & -2\Lambda \end{bmatrix} R_\phi < 0 \quad (6)$$

holds. Then, for any initial condition $\tilde{x}_0 \in \mathbb{R}^{n_{\tilde{x}}}$ and for any reference $r \in \mathbb{R}^{n_r}$, the NN controller $\kappa(x, r)$ ensures exponential stability of the steady state \tilde{x}_ for the closed-loop system (1) and achieves zero steady-state offset, i.e., $y_k \rightarrow r$ for $k \rightarrow \infty$.*

Proof We use the fact that the incremental quadratic constraint (4) holds. Given that (6) is strict, there exists an $\epsilon > 0$, such that right/left multiplying (6) with $[(\tilde{x}_k - \tilde{x}_*)^\top (w_k - w_*)^\top]^\top$ yields

$$\begin{bmatrix} \tilde{x}_k - \tilde{x}_* \\ \tilde{u}_k - \tilde{u}_* \end{bmatrix}^\top \begin{bmatrix} \tilde{A}^\top P \tilde{A} - P & \tilde{A}^\top P \tilde{B} \\ \tilde{B}^\top P \tilde{A} & \tilde{B}^\top P \tilde{B} \end{bmatrix} \begin{bmatrix} \tilde{x}_k - \tilde{x}_* \\ \tilde{u}_k - \tilde{u}_* \end{bmatrix} \leq -\epsilon \|\tilde{x}_k - \tilde{x}_*\|^2. \quad (7)$$

We define the Lyapunov function $V(\tilde{x}) := (\tilde{x} - \tilde{x}_*)^\top P (\tilde{x} - \tilde{x}_*)$ and translate (7) into $V(\tilde{x}_{k+1}) - V(\tilde{x}_k) \leq -\epsilon \|\tilde{x}_k - \tilde{x}_*\|^2$ which shows global exponential stability. \blacksquare

While this result can verify global stability of the considered feedback interconnection, in practice, NN controllers are usually trained on training data from a problem-relevant region resulting

from physical constraints of the environment. Oftentimes, NN controllers hence perform well at and around the training data set, yet show unpredictable behavior on unknown regions. For that reason, (i) it is unlikely to have a globally stabilizing NN controller and (ii) mostly local stability properties corresponding to the relevant data region are sufficient in practice. Imposing local stability is less restrictive and hence allows for a better performance. Therefore, in the next subsection, we present a result to certify local stability.

Remark 2 *Although the focus of this paper is on analysis, the resulting LMI (6) can be incorporated in the training of NN controllers analogous to Pauli et al. (2021); Revay et al. (2021).*

3.2. Local stability for a fixed reference

As considering local sector bounds that are tighter than the global ones can reduce conservatism, we formulate a local stability condition for the feedback interconnection of \tilde{G} and $\kappa(x_k, r)$, with a fixed constant reference r as an external input. In (Yin et al., 2021), a similar local analysis was performed for the reference-free case. First, in this subsection, we consider local stability for a given reference r and in the next subsection, we provide a local stability result for references from a set.

For a given constant reference r , we choose some $d \in \mathbb{R}^{n_1}, d_i > 0$ and form a symmetrical set around the steady state v_*^1 of the inputs to the neurons of the first hidden layer $v^1 \in [v_*^1 - d, v_*^1 + d] = [\underline{v}^1, \bar{v}^1]$. With the bounds on v^1 , we can explicitly determine the bounds on all entries of $v \in [\underline{v}, \bar{v}]$ and $w \in [\underline{w}, \bar{w}]$. Similar to Yin et al. (2021), we define the local incremental sector bounds $\alpha_\phi^{i,j}$ and $\beta_\phi^{i,j}$ of the j -th neuron of the i -th layer of the NN such that the local slope-restriction property

$$\alpha_\phi^{i,j} \leq \frac{\varphi(v^{i,j}) - \varphi(v_*^{i,j})}{v^{i,j} - v_*^{i,j}} \leq \beta_\phi^{i,j} \quad \forall i \in 1, \dots, l \quad \forall j \in 1, \dots, n_i$$

holds for all elements of $v \in [\underline{v}, \bar{v}]$ and such that the bounds are tight. Moreover, we define the diagonal matrices

$$\begin{aligned} \alpha_\phi^i &:= \text{diag}(\alpha_\phi^{i,j}), \quad j = 1, \dots, n_i, & \alpha_\phi &:= \text{blkdiag}(\alpha_\phi^i), \quad j = 1, \dots, l, \\ \beta_\phi^i &:= \text{diag}(\beta_\phi^{i,j}), \quad j = 1, \dots, n_i, & \beta_\phi &:= \text{blkdiag}(\beta_\phi^i), \quad j = 1, \dots, l, \end{aligned}$$

the ellipsoidal set $\mathcal{E}_P(\tilde{x}_*) := \{\tilde{x} \in \mathbb{R}^{n_x} \mid (\tilde{x} - \tilde{x}_*)^\top P(\tilde{x} - \tilde{x}_*) \leq 1\}$, and we have $v^1 - v_*^1 = N_0^1(\tilde{x} - \tilde{x}_*)$ with N_0^1 as defined in (5). The index j in $N_{0,j}^1$ and d_j denotes the j -th row of N_0^1 and respectively, the j -th entry of d . With these definitions, we state an LMI condition for local stability.

Theorem 3 *Suppose there exist matrices $\Lambda \in \mathbb{D}_+^n, P > 0$ such that*

$$R_V^\top \begin{bmatrix} \tilde{A}^\top P \tilde{A} - P & \tilde{A}^\top P \tilde{B} \\ \tilde{B}^\top P \tilde{A} & \tilde{B}^\top P \tilde{B} \end{bmatrix} R_V + R_x^\top \begin{bmatrix} -2\alpha_\phi \beta_\phi \Lambda & (\alpha_\phi + \beta_\phi) \Lambda \\ (\alpha_\phi + \beta_\phi) \Lambda & -2\Lambda \end{bmatrix} R_x < 0, \quad (8a)$$

$$\begin{bmatrix} d_j^2 & N_{0,j}^1 \\ N_{0,j}^{1\top} & P \end{bmatrix} \geq 0, \quad \forall j = 1, \dots, n_1 \quad (8b)$$

holds. Then, for any initial condition $\tilde{x}_0 \in \mathcal{E}_P(x_)$ the NN controller $\kappa(x, r)$ ensures exponential stability of the steady state \tilde{x}_* for the closed-loop system (1) and achieves zero steady-state offset, i.e., $y_k \rightarrow r$ for $k \rightarrow \infty$.*

Proof The proof is a direct extension of Theorem 1 using the derivation in (Yin et al., 2021, Thm. 1). ■

In order to maximize the RoA, instead of solving solely the feasibility problem (8), we can solve the semidefinite program with the objective $\min_{P,\Lambda} \text{trace} P$ and LMI constraints (8). Even though Theorem 3 verifies local stability and therefore, as we argued above, it is potentially less conservative than Theorem 1, this result still is somewhat limiting due to the fact that it is stated for a fixed reference r . Tracking a different constant reference in general necessitates to recalculate the incremental sector bounds α_ϕ and β_ϕ . For example, for state feedback, $v_*^!$ changes with r and it is required to newly determine the bounds on v and w as well as the incremental sector bounds α_ϕ and β_ϕ for a different reference r . Finally, the LMI conditions (8) need to be checked again with the new values of α_ϕ and β_ϕ . For the special case of output error feedback, v_* conveniently remains the same for all references r , given that $v_*^! = W^0(r - Cx_*) + b^0 = b^0 \forall r$ with (2), such that α_ϕ and β_ϕ remain the same for different references r . This means in the case of output error feedback, if Theorem 3 holds, the controller $\kappa(x, r)$ ensures local exponential stability of the closed-loop system (1) for all references $r \in \mathbb{R}^{n_r}$ and all initial conditions $\tilde{x}_0 \in \mathcal{E}_P(\tilde{x}_*)$, i.e., an ellipsoidal set whose size is independent of the reference r and that is centered at the corresponding setpoint \tilde{x}_* . This means that for error output feedback local stability depends only on the tracking error.

3.3. Local stability for a range of references

In this subsection, we state a more general local stability result that guarantees stability for all references from a certain set. According to (2), the steady-state manifold is characterized by

$$\begin{bmatrix} x_* \\ u_* \end{bmatrix} = A_a^{-1} \begin{bmatrix} 0 \\ r \end{bmatrix} = \begin{bmatrix} * & M \\ * & * \end{bmatrix} \begin{bmatrix} 0 \\ r \end{bmatrix}.$$

Thus, we obtain a linear map $x_*(r) = Mr$. Using this map and the fact that we use an offset-free tracking formulation (due to the integrator), we can characterize the set of setpoints $x_*(r)$ for all admissible references. A forward pass of $x_*(r)$ through the NN gives the values of $v_*(r)$ and $w_*(r)$, as well. Carrying out the stability analysis for a range of references necessitates to determine the incremental sector bounds α_ϕ and β_ϕ according to Subsection 3.2 for a given nominal reference $r_{\text{nom}} \in \mathbb{R}^{n_r}$ and a given vector $d \in \mathbb{R}^{n_1}$, $d_i > 0$. This time, we aim to find an underestimate of the RoA formulated in the augmented state \tilde{x} and the desired reference r and therefore, we introduce the ellipsoidal set

$$\mathcal{E}_{P,Q}(r_{\text{nom}}) := \left\{ \tilde{x} \in \mathbb{R}^{n_{\tilde{x}}}, r \in \mathbb{R}^{n_r} \mid (\tilde{x} - \tilde{x}_*(r))^T P (\tilde{x} - \tilde{x}_*(r)) + (r - r_{\text{nom}})^T Q (r - r_{\text{nom}}) \leq 1 \right\}.$$

This choice is natural as it addresses the trade-off between the initial conditions \tilde{x}_0 and the choice of the constant reference r . For $r = r_{\text{nom}}$, we retrieve the ellipsoidal set $\mathcal{E}_P(\tilde{x}_*)$. The ellipse in \tilde{x} corresponding to $r = r_{\text{nom}}$ is the largest possible for the given vector d , whereas for references r that deviate from the nominal reference r_{nom} , the corresponding ellipsoidal set in \tilde{x} shrinks: $(\tilde{x} - \tilde{x}_*(r))^T P (\tilde{x} - \tilde{x}_*(r)) \leq 1 - (r - r_{\text{nom}})^T Q (r - r_{\text{nom}})$. Finally, the union of the ellipses for all $r \in \mathcal{E}_{P,Q}(r_{\text{nom}})$ provides an enlarged guaranteed RoA if the reference is adjusted online, cf. Subsection 3.4.

Theorem 4 *Suppose there exist matrices $\Lambda \in \mathbb{D}_+^n$, $P > 0$, $Q > 0$ such that*

$$R_V^T \begin{bmatrix} \tilde{A}^T P \tilde{A} - P & \tilde{A}^T P \tilde{B} \\ \tilde{B}^T P \tilde{A} & \tilde{B}^T P \tilde{B} \end{bmatrix} R_V + R_x^T \begin{bmatrix} -2\alpha_\phi\beta_\phi\Lambda & (\alpha_\phi + \beta_\phi)\Lambda \\ (\alpha_\phi + \beta_\phi)\Lambda & -2\Lambda \end{bmatrix} R_x < 0, \quad (9a)$$

$$\begin{bmatrix} d_j^2 & [N_0^1 \ N_0^1 M]_j \\ [N_0^1 \ N_0^1 M]_j^\top & \begin{bmatrix} P & 0 \\ 0 & Q \end{bmatrix} \end{bmatrix} \geq 0, \quad \forall j = 1, \dots, n_1 \quad (9b)$$

holds. Then, for all pairs of initial conditions and references $(\tilde{x}_0, r) \in \mathcal{E}_{P,Q}(r_{\text{nom}})$ the NN controller $\kappa(x, r)$ ensures exponential stability of the steady state $\tilde{x}_*(r)$ for the closed-loop system (1) and achieves zero steady-state offset, i.e., $y_k \rightarrow r$ for $k \rightarrow \infty$.

Proof First, we show that for all $(\tilde{x}, r) \in \mathcal{E}_{P,Q}(r_{\text{nom}})$, the input to the first neuron v^1 stays between the given bounds, i.e., $v^1 \in [\underline{v}^1, \bar{v}^1]$. Applying the Schur complement to LMI (9b) yields

$$[N_0^1 \ N_0^1 M]_j \begin{bmatrix} P^{-1} & 0 \\ 0 & Q^{-1} \end{bmatrix} [N_0^1 \ N_0^1 M]_j^\top \leq d_j^2, \quad \forall j = 1, \dots, n_1. \quad (10)$$

According to Lemma 1 in (Hindi and Boyd, 1998), from (10) we get

$$\mathcal{E}_{P,Q}(r_{\text{nom}}) \subseteq \left\{ \tilde{x} \in \mathbb{R}^{n_{\tilde{x}}}, r \in \mathbb{R}^{n_r} \mid \left\| [N_0^1 \ N_0^1 M]_j \begin{bmatrix} \tilde{x} - \tilde{x}_*(r) \\ r - r_{\text{nom}} \end{bmatrix} \right\| \leq d_j, j = 1, \dots, n_1 \right\}$$

and

$$\left\| [N_0^1 \ N_0^1 M] \begin{bmatrix} \tilde{x} - \tilde{x}_*(r) \\ r - r_{\text{nom}} \end{bmatrix} \right\| = |v^1 - v_*^1(r_{\text{nom}})| \leq d,$$

from which we conclude $\mathcal{E}_{P,Q}(r_{\text{nom}}) \subseteq \{\tilde{x} \in \mathbb{R}^{n_{\tilde{x}}}, r \in \mathbb{R}^{n_r} \mid v^1 \in [\underline{v}^1, \bar{v}^1]\}$. Hence, the given positive invariance of $\mathcal{E}_{P,Q}(r_{\text{nom}})$ ensures invariance of the bounds $[\underline{v}^1, \bar{v}^1]$. Closed-loop stability and offset-free tracking then follow from (9a) analogous to the proof of Theorem 1 with α_ϕ and β_ϕ instead of α and β , and $\mathcal{E}_{P,Q}(r_{\text{nom}})$ is an inner approximation of the RoA. \blacksquare

With this result, we can verify stability locally for all combinations of initial conditions \tilde{x}_0 and references r from the ellipsoidal set $\mathcal{E}_{P,Q}(r_{\text{nom}})$ and thus, it allows for tracking of piecewise constant references. In the next subsection, we present an approach to increase the guaranteed RoA.

Remark 5 *Model uncertainties, e.g., formulated as integral quadratic constraints, can be incorporated into the feedback interconnection considered in this paper to analyze robust stability and robust estimates of the RoA, cf. (Yin et al., 2021).*

3.4. Extended guaranteed RoA via a reference governor

A reference governor is an augmentation to a locally stabilizing controller that is often employed to ensure satisfaction of state and input constraints. In particular, a reference governor uses the current state measurement to modify the reference that is fed into the controller (Bemporad, 1998).

In the given setup of this paper, a reference governor provides a tool to extend the RoA. After computing the matrices P and Q via Theorem 4, at all times k , the reference governor checks if the combination of the current state \tilde{x}_k and the reference r lie in the guaranteed RoA $\mathcal{E}_{P,Q}$. If $(\tilde{x}_k, r) \notin \mathcal{E}_{P,Q}$, e.g., due to a change in the reference, then the reference governor adjusts the reference to the closest possible reference \hat{r}_k for which the current state \tilde{x}_k lies in the RoA and the surrogate reference \hat{r}_k is fed into the closed-loop system (1). In this way, even if the desired reference leads

to an unstable closed loop, the reference governor ensures stability. A corresponding surrogate reference \hat{r}_k can be computed according to the following formulation:

$$\hat{r}_k := \arg \min_{\hat{r}} \|r - \hat{r}\|^2 \text{ s.t. } (\tilde{x}_k - \tilde{x}_*(\hat{r}))^\top P(\tilde{x}_k - \tilde{x}_*(\hat{r})) + (\hat{r} - r_{\text{nom}})^\top Q(\hat{r} - r_{\text{nom}}) \leq 1. \quad (11)$$

The surrogate reference \hat{r}_k is updated at every time step until the current state lies in the guaranteed RoA corresponding to the desired reference r (in case this is possible) such that $\hat{r}_k = r$. Even if the reference r is outside of the guaranteed RoA, the reference closest to the desired r that lies in the estimate of the RoA can be tracked. Overall, the reference governor provides stability independent of the desired reference r and guarantees desirable closed-loop properties for any \tilde{x}_0 in the union of the local ellipsoidal sets $\mathcal{E}_{P,Q}$ over all admissible references r , cf. Fig. 4. Note that this reference governor has similarities to the approach in (Donti et al., 2020), where instead of the reference, the control input is adjusted by projection onto a set that guarantees exponential stability.

Remark 6 *Note that in the special case of output error feedback, local stability is independent of r and depends only on the tracking error. Hence, a corresponding reference governor is given by*

$$\hat{r}_k = \arg \min_{\hat{r}} \|r - \hat{r}\| \text{ s.t. } (\tilde{x}_k - \tilde{x}_*(\hat{r}))^\top P(\tilde{x}_k - \tilde{x}_*(\hat{r})) \leq 1.$$

Thus, in the output error feedback case, we can achieve offset-free tracking for any reference $r \in \mathbb{R}^{n_r}$, if the reference governor is feasible for the initial state.

4. Numerical example

In this section, we illustrate our results on an example. The analysis is applicable to all NN controllers within the suggested control architecture, regardless of how they were trained. In the following example, an NN controller approximates a model predictive controller (MPC) which is beneficial in case the MPC is computationally expensive.

Consider the problem of steering an inverted pendulum with mass $m = 0.15$ kg, length $L = 0.5$ m, and friction coefficient $\mu = 0.5$ Nm/rad to a desired reference position. The linearized dynamics are

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ \frac{g}{L} & -\frac{\mu}{mL^2} \end{bmatrix} x + \begin{bmatrix} 0 \\ \frac{g}{L} \end{bmatrix} u, \quad y = \begin{bmatrix} 0 & 1 \end{bmatrix} x \quad (12)$$

with state $x = [\theta^\top \ \dot{\theta}^\top]^\top$, angular position θ , angular velocity $\dot{\theta}$, and control input u . We use supervised learning to determine a stabilizing controller for the dynamics (12) discretized with sampling time $T_s = 0.02$ s. For this purpose, we train an NN controller such that it approximates an MPC trained on the discretized and linearized dynamics of the inverted pendulum with input constraint $|u| \leq 1$, similar to Zhang et al. (2016). For training of the NN, we utilize closed-loop trajectories with initial conditions sampled from a uniform distribution $x_0 \in \mathcal{U}(-0.5, 0.5)$ and we set $k_{\mathcal{E}} = 1$. We consider an NN, with activation function tanh and two layers of 5 neurons each, that serves as a state feedback controller, i.e., $w_0 = x$. While the approach in (Yin et al., 2021) can only handle zero bias to ensure the setpoint is $x_* = 0$, $u_* = 0$, in our setup no structural constraints on the weights and biases of the NN are necessary, as any undesired offset is eliminated by the integrator.

Theorem 1 is not applicable here, meaning that we cannot show global stability of the closed loop. This is not surprising, given that the NN controller was trained on trajectories from a problem-relevant region. To compute a local RoA, we choose $d_i = 0.345$, $i = 1, \dots, n_1$, and then propagate




Figure 2: Steady-state manifold $x_*(r)$, ellipsoidal inner approximations of RoA \mathcal{E}_P and $\mathcal{E}_{P,Q}$, and example trajectory x_{RG} using a reference governor for the linearized inverted pendulum.

the resulting bounds $[\underline{v}^1, \bar{v}^1]$ through the NN to obtain the incremental sector bounds α_ϕ and β_ϕ . With vector d and the resulting bounds α_ϕ and β_ϕ , we apply Theorem 3 with the reference $r = 0$ (computation time 1.14s) and Theorem 4 with the nominal reference $r_{\text{nom}} = 0$ (computation time 0.498s) using numerical SDP solvers (Löfberg, 2004; MOSEK ApS, 2019) in Matlab.

Fig. 4 shows the corresponding ellipses $\mathcal{E}_P(\tilde{x}_*(0))$ (blue, solid) and $\mathcal{E}_{P,Q}(0)$ (red, dashed) for different $r \in \mathcal{E}_{P,Q}(0)$, as well as the steady-state manifold $x_*(r)$ of the problem. Based on the result of Theorem 4, we can track any reference angle $r \in [-0.2, 0.2]$. Further, note that in this example, the extended RoA, i.e., the union of the ellipses $\mathcal{E}_{P,Q}(0)$, is notably larger than $\mathcal{E}_P(\tilde{x}_*(0))$. We also implemented a reference governor to track the reference $r = -1$. Whenever the constraint in (11) is not fulfilled with the current state and the desired reference r , the reference governor determines a surrogate reference \hat{r}_k closest to the desired reference. In Fig. 4, a trajectory x_{RG} for a random initial condition under control with the reference governor is shown. This trajectory tracks the closest point to r within the extended RoA with zero steady-state offset while x_{RG} stays in the extended guaranteed RoA at all times. Note that without the reference governor the trajectory diverges.

5. Conclusion

We proposed a method to analyze local and global stability in offset-free setpoint tracking using NN controllers, also providing ellipsoidal inner approximations of the corresponding RoA. We derived linear matrix inequalities to verify stability using Lyapunov theory, exploiting the fact that activation functions used in NNs are slope-restricted. First, we showed global stability and then stated less restrictive local stability results (i) for a given reference and (ii) for a set of references, which allows for tracking of piecewise constant references. Based on the latter result, we were able to significantly increase the guaranteed RoA using a reference governor. Finally, we tested our stability analysis on the example of a linearized inverted pendulum.

An interesting direction for future research is to extend the guaranteed RoA by optimizing over the range $[\underline{v}^1, \bar{v}^1]$ of the input to the first layer, that in the current setup is chosen by the user.

Acknowledgments

This work was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Patricia Pauli, Julian Berberich, and Anne Koch.

References

- Alberto Bemporad. Reference governor for constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 43(3):415–419, 1998.
- Priya L. Donti, Melrose Roderick, Mahyar Fazlyab, and J. Zico Kolter. Enforcing robust control guarantees within neural network policies. *arXiv preprint arXiv:2011.08105*, 2020.
- Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of Lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, pages 11423–11434, 2019.
- Michael Hertneck, Johannes Köhler, Sebastian Trimpe, and Frank Allgöwer. Learning an approximate model predictive controller with guarantees. *IEEE Control Systems Letters*, 2(3):543–548, 2018.
- Haitham Hindi and Stephen Boyd. Analysis of linear systems with saturation using convex optimization. In *Proc. of the 37th IEEE Conference on Decision and Control*, pages 903–908. IEEE, 1998.
- Mohamed Ibrahim, Christian Kallies, and Rolf Findeisen. Learning-supported approximated optimal control for autonomous vehicles in the presence of state dependent uncertainties. In *Proc. of European Control Conference*, pages 338–343. IEEE, 2020.
- Asriel U Levin and Kumpati S Narendra. Control of nonlinear dynamical systems using neural networks: Controllability and stabilization. *IEEE Transactions on Neural Networks*, 4(2):192–206, 1993.
- J. Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *Proc. of the CACSD Conference*, Taipei, Taiwan, 2004.
- MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019. URL <http://docs.mosek.com/9.0/toolbox/index.html>.
- Gabriele Pannocchia and James B Rawlings. Disturbance models for offset-free model-predictive control. *AIChE journal*, 49(2):426–437, 2003.
- Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural networks using Lipschitz bounds. *IEEE Control Systems Letters*, 2021.
- Demetri Psaltis, Athanasios Sideris, and Alan A Yamamura. A multilayered neural network controller. *IEEE Control Systems Magazine*, 8(2):17–21, 1988.
- Max Revay and Ian Manchester. Contracting implicit recurrent neural networks: Stable models with improved trainability. In *Proc. of Learning for Dynamics and Control*, pages 393–403. PMLR, 2020.
- Max Revay, Ruigang Wang, and Ian R Manchester. Lipschitz bounded equilibrium networks. *arXiv preprint arXiv:2010.01732*, 2020.

- Max Revay, Ruigang Wang, and Ian R. Manchester. A convex parameterization of robust recurrent neural networks. *IEEE Control Systems Letters*, 5(4):1363–1368, 2021. doi: 10.1109/LCSYS.2020.3038221.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- He Yin, Peter Seiler, Ming Jin, and Murat Arcak. Imitation learning with stability and safety guarantees. *arXiv preprint arXiv:2012.09293*, 2020.
- He Yin, Peter Seiler, and Murat Arcak. Stability analysis using quadratic constraints for systems with neural network controllers. *IEEE Transactions on Automatic Control*, 2021.
- Tianhao Zhang, Gregory Kahn, Sergey Levine, and Pieter Abbeel. Learning deep control policies for autonomous aerial vehicles with MPC-guided policy search. In *Proc. of International Conference on Robotics and Automation (ICRA)*, pages 528–535. IEEE, 2016.