CSMA 2026
17ème Colloque National en Calcul des Structures
18-22 Mai 2026, Presqu'île de Giens (Var)

An analysis of tube radius under approximate model predictive control

Sivalingam S M[1], Ludovic Chamoin[2], Laurent Fribourg[3]

[1 2] *Université Paris-Saclay, CentraleSupélec, ENS Paris-Saclay, CNRS, LMPS-Laboratoire de Mécanique Paris-Saclay, Gif-sur-Yvette 91190, France*
[3] *Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire de Méthodes Formelles, Gif-sur-Yvette, France*

Résumé — Several studies have explored the idea of deep neural networks as surrogates for nonlinear model predictive control. However, the effect of the approximation on the structure and the behavior of a tube model predictive controller has not been thoroughly analyzed. This work investigates this effect using ideas from statistical machine learning and contraction theory, with particular emphasis on systems whose dynamics are contractive. The approximation error is considered as a perturbation on the control and is bounded probabilistically in terms of the network's generalisation error. The study is conducted within the ERC Consolidator Grant project DREAM-ON with a target application to Structural Health Monitoring. The proposed methodology contributes to the feedback loop created between a physical asset and its digital twin. Numerical validation is provided using a simple two-dimensional nonlinear oscillator.
Mots clés — Tube Model Predictive Control, Neural Network, Generalisation Error, Contraction

1 Introduction

Learning-based control of dynamical systems has seen rapid advancements over the last decade, with methods such as reinforcement learning (RL), imitation learning (IL), and model predictive control (MPC) . Among these, IL offers an alternative to RL, especially in settings where expert trajectories are available for supervision [1, 2]. However, transferring these learned policies from simulation to reality remains challenging due to unmodeled physics, disturbances and parametric uncertainty [3].

These arise from process noise or from inaccuracies in physical modelling of real-world systems. Such disturbances are represented in the system dynamics by introducing random noise sampled from an appropriate probability distribution [4]. Control of these systems, therefore, requires the controller to be robust to these stochastic variations. The literature offers several approaches to handle such uncertainties, among which the most widely adopted is tube MPC. Tube MPC achieves robustness by combining a nominal MPC controller with an ancillary feedback controller that constrains the system trajectory to remain within a bounded tube around the reference trajectory [5].

However, solving a nonlinear MPC problem is often resource-constrained. Also, it is often infeasible in real-time, high-dimensional systems. Hence, several works focus on replacing the computationally expensive MPC with a trained neural network. In [6], the authors study the effect of approximation on a nonlinear MPC problem and present a probabilistic bound to guarantee recursive feasibility and stability. Their idea treats the approximation error as a perturbation of the control and quantifies a bound using the Hoeffding inequality. In [7], a bound on the worst-case approximation error of the network is obtained in terms of the generalisation error of the network, and it is proven that the approximate MPC is asymptotically stable. Similarly, [8] treats the network's approximation error as a perturbation and proves stability using Lyapunov theory. It was based on the worst-case approximation error of the neural network, bounded by Lipschitz assumptions on the control policy and the network. To reduce the network's conservatism, the authors also use an ensemble of networks rather than a single network.

Building on the above literature, our study bounds the worst-case approximation error in terms of the network's generalisation gap. In particular, we introduce a new probabilistic bound on the worst-case error of the network in terms of the generalization bound. Our analysis relies entirely on the Lipschitz assumption of the system, which is much simpler to analyze the error propagation than the Lyapunov approach. The proposed approach is tested on a two-dimensional Van der Pol oscillator subject to external disturbances.

1.1 Notations

In this paper, $\mathbb{R}$ and $\mathbb{N}$ denote the sets of real and natural numbers, respectively. Bold letters are used for both matrices and vectors. The symbols $\mathbf{x}$ and $\mathbf{u}$ are used to denote the state and control vectors, respectively. For a function $F : \mathcal{X} \to \mathcal{Y}$, we define $\|F\|_2 = \left(\int_{\mathcal{X}} \|F(x)\|_{\mathcal{Y}}^2 \, dx \right)^{1/2}$, if $\mathcal{X}$ is continuous and $\|F\|_2 = \left(\sum_{x \in \mathcal{X}} \|F(x)\|_{\mathcal{Y}}^2 \right)^{1/2}$, if $\mathcal{X}$ is discrete. Let $\|\cdot\|$ denote the Euclidean norm on $\mathbb{R}^n$ unless another norm is explicitly specified. For a matrix $\mathbf{A}$, $\|\mathbf{A}\|$ denotes the induced operator norm, and $\rho(\mathbf{A})$ denotes its spectral radius. Let $\mathbf{P}$ be a symmetric positive definite matrix that is $\mathbf{P} \succ 0$ then we denote $\|x\|_{\mathbf{P}} = \sqrt{\mathbf{x}^\top \mathbf{P} \mathbf{x}}$ and $\|\mathbf{A}\|_{\mathbf{P}} = \sup_{x \neq 0} \frac{\|\mathbf{A}\mathbf{x}\|_P}{\|\mathbf{x}\|_P}$. The Minkowski sum is defined as $\mathcal{S}_1 \oplus \mathcal{S}_2 = \{s_1 + s_2 : s_1 \in \mathcal{S}_1, s_2 \in \mathcal{S}_2\}$. For a set $\mathcal{S}$ and matrix $\mathbf{A}$ we write $\mathbf{A}\mathcal{S} = \{\mathbf{A}s : s \in \mathcal{S}\}$.

2 Tube MPC

Consider a nonlinear control system with the continuous-time dynamics

$$\dot{\mathbf{x}}(t) = F(\mathbf{x}(t), \mathbf{u}(t)) + \mathbf{w}(t), \tag{1}$$

where $\mathbf{x}(t) \in \mathcal{X} \subset \mathbb{R}^{d_x}$ is the state of the system, $\mathbf{u}(t) \in \mathcal{U} \subset \mathbb{R}^{d_u}$ is the control. $\mathbf{w}(t)$ accounts for the disturbances and process noise. The admissible sets are assumed to be compact $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^{d_x} \mid \mathbf{x}_{\min} \leq \mathbf{x} \leq \mathbf{x}_{\max}\}$, $\quad \mathcal{U} = \{\mathbf{u} \in \mathbb{R}^{d_u} \mid \mathbf{u}_{\min} \leq \mathbf{u} \leq \mathbf{u}_{\max}\}$, and $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^{d_x} \mid \|\mathbf{w}\|_\infty \leq w_{\max}\}$. The objective of this work is to design a neural network-based controller that keeps the true state of the system within a tube. We use the discrete-time version of the original dynamical system governed by

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t, \; t \in \mathbb{N}, \tag{2}$$

where $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \to \mathbb{R}^{d_x}$ parametrizes the dynamics of the discrete system obtained from sampling the continuous time system with time period $T_s > 0$. Let $f(\cdot,\cdot)$ be twice continuously differentiable on a region that contains all trajectories of interest. Further, let it be Lipschitz in $(\mathbf{x}, \mathbf{u})$ with Lipschitz constant $L_{\mathbf{z}}$ and $L_{\mathbf{v}}$. Let there exists an equilibrium pair $(\hat{\mathbf{z}}, \hat{\mathbf{v}})$ with $f(\hat{\mathbf{z}}, \hat{\mathbf{v}}) = \hat{\mathbf{z}}$ and the pair $\mathbf{A}, \mathbf{B}$ defined by $\mathbf{A} = \left.\frac{\partial f(\mathbf{z},\mathbf{v})}{\partial \mathbf{z}}\right|_{(\hat{\mathbf{z}},\hat{\mathbf{v}})}$ and $\mathbf{B} = \left.\frac{\partial f(\mathbf{z},\mathbf{v})}{\partial \mathbf{v}}\right|_{(\hat{\mathbf{z}},\hat{\mathbf{v}})}$ which linearizes f and is stabilizable. Tube MPC is a robust control strategy that breaks the control into two parts. First, a nominal control for a disturbance-free system, and an ancillary feedback to keep the real state close to the nominal state within a tube. The total control using tube MPC is broken as

$$\mathbf{u}_t = \underbrace{\mathbf{v}_t}_{\text{nominal control}} + \underbrace{\mathbf{K}_t(\mathbf{x}_t - \mathbf{z}_t)}_{\text{ancillary control}}, \tag{3}$$

where $\mathbf{v}_t$ is the nominal control and $\mathbf{K}_t$ is the stabilizing feedback gain, $\mathbf{x}_t$ is the state of the true model and $\mathbf{z}_t$ is the nominal state. Let the nominal MPC use a finite horizon N with stage cost $\|\mathbf{z} - \mathbf{x}^r\|_{\mathbf{Q}}^2 + \|\mathbf{v} - \mathbf{u}^r\|_{\mathbf{R}}^2$, $\mathbf{Q} \succ 0$, $\mathbf{R} \succ 0$, and $\exists$ a terminal set $\mathcal{X}_f \subset \mathcal{Z}$ and terminal cost $V_f(\mathbf{z})$ such that the standard MPC stabilizing conditions holds. Further, the terminal ingredients satisfy the usual conditions that make V_N a Lyapunov function candidate for the nominal dynamics as mentioned in [12]. Consider the problem of steering the trajectory of the original system to a reference pair $(\mathbf{x}_t^r, \mathbf{u}_t^r)_{t=0,\cdots,N-1}$ as presented by the authors in [9]. Initially, the disturbance-free nominal system is steered as close as possible to the reference trajectory by using the nominal control $\mathbf{v}_t$. This is addressed by solving the following finite-horizon optimal control problem over the N-step :

$$\mathrm{V}^*(\mathbf{z}) = \min_{\mathbf{v}} \mathrm{V}_N(\mathbf{z}, \mathbf{x}^r, \mathbf{v}) = \sum_{t=0}^{N-1} \left(\|\mathbf{z}_t - \mathbf{x}_t^r\|_{\mathbf{Q}}^2 + \|\mathbf{v}_t - \mathbf{u}_t^r\|_{\mathbf{R}}^2 \right) + V_f(\mathbf{z}_N) \tag{4}$$

subject to $\mathbf{z}_{t+1} = f(\mathbf{z}_t, \mathbf{v}_t),\ \mathbf{z} \in \mathcal{Z},\ \mathbf{v} \in \mathcal{V},\ t = 0, \cdots, N-1,\ \mathbf{z}_N \in \mathcal{X}_f$, where $\mathbf{v} = \{\mathbf{v}_0, \ldots, \mathbf{v}_{N-1}\} \in \mathcal{V}$, $\mathbf{z}_t$ and $\mathbf{v}_t$ are the nominal state and control respectively. $\mathcal{Z}, \mathcal{V}$ are the tightened admissible sets. The solution of the N horizon problem yields a control sequence $\{\mathbf{v}_t^*\}_{t=0}^{N-1}$ and nominal states $\{\mathbf{z}_t^*\}_{t=0}^{N}$. In receding horizon fashion, we apply $\mathbf{v}_0^*$ and re-solve (4) at the next step. Now, let $\mathbf{A}_t, \mathbf{B}_t$ denote the Jacobians of f evaluated at $(\mathbf{z}_t, \mathbf{v}_t)$ $\mathbf{A}_t = \left.\frac{\partial f(\mathbf{z}, \mathbf{v})}{\partial \mathbf{z}}\right|_{(\mathbf{z}_t, \mathbf{v}_t)}, \quad \mathbf{B}_t = \left.\frac{\partial f(\mathbf{z}, \mathbf{v})}{\partial \mathbf{v}}\right|_{(\mathbf{z}_t, \mathbf{v}_t)}$ When using equilibrium linearization, the matrices $\mathbf{A}_t$ and $\mathbf{B}_t$ are constant as in the assumption 2. By using first-order Taylor expansion, for sufficiently small $\|\mathbf{e}\|$ and $\|\mathbf{d}\|$, we obtain $f(\mathbf{z}_t + \mathbf{e}, \mathbf{v}_t + \mathbf{d}) = f(\mathbf{z}_t, \mathbf{v}_t) + \mathbf{A}_t \mathbf{e} + \mathbf{B}_t \mathbf{d} + \mathbf{r}_t(\mathbf{e}, \mathbf{d})$, where the remainder $\mathbf{r}_t(\mathbf{e}, \mathbf{d})$ based on the assumption 2 satisfies the bound $\|\mathbf{r}_t(\mathbf{e}, \mathbf{d})\| \leq c_{\text{rem}}\left(\|\mathbf{e}\|^2 + \|\mathbf{d}\|^2\right)$ for $c_{\text{rem}} \geq 0$ which depends on the sup-norm of second derivatives of f. Initially we describe our ancillary control with the time-varying pair $\mathbf{A}_t, \mathbf{B}_t$. Given a nominal pair $\{\mathbf{z}_t, \mathbf{v}_t\}_{t \geq 0}$, we consider linearization giving a pair of time-varying matrices $\mathbf{A}_t, \mathbf{B}_t$ that are uniformly stabilizable, and that a corresponding time-varying LQR or stabilizing time-varying gain sequence $\{\mathbf{K}_t\}$ is used. Once the nominal control is applied, an error between the nominal state and the true state can be defined by $\mathbf{e}_t = \mathbf{x}_t - \mathbf{z}_t \in \mathcal{E}$. Under the tube controller $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t = f(\mathbf{z}_t, \mathbf{v}_t) + \mathbf{A}_t \mathbf{e}_t + \mathbf{B}_t \mathbf{K}_t \mathbf{e}_t + \mathbf{r}_t(\mathbf{e}_t, \mathbf{K}_t \mathbf{e}_t) + \mathbf{w}_t$. This yields the following recursion for the error dynamics

$$\mathbf{e}_{t+1} \quad = \quad \mathbf{A}_{\mathbf{K}_t} \mathbf{e}_t + \tilde{\mathbf{w}}_t, \tag{5}$$

where $A_{K_t} = A_t + B_t K_t$, and $\tilde{w}_t = r_t(e_t, K_t e_t) + w_t$. Also, $\tilde{\mathbf{w}}_t \in \tilde{\mathcal{W}}$ a compact admissible region. Now, we define the robustly positive-invariant set in which the error resides.

Definition 1. *A set $\mathcal{E} \subset \mathbb{R}^{d_x}$ which is compact is robust positively invariant (RPI) for* (5) *with the disturbance set $\tilde{\mathcal{W}}$ and matrices $\{\mathbf{A}_{\mathbf{K}_t}\}$ if $\forall\ \mathbf{e} \in \mathcal{E}$ and $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}$, $\mathbf{A}_{\mathbf{K}_t}\mathbf{e} + \tilde{\mathbf{w}} \in \mathcal{E}\ \forall\ t$.*

To make computations easier we use a single K. To do this, we make the following two assumptions. There exists a compact region $\mathcal{R} \subseteq \mathcal{Z} \times \mathcal{V} \ni (\mathbf{z}_t, \mathbf{v}_t) \in \mathcal{R}, \forall t \geq 0$, and $\left\|\frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\mathbf{z}, \mathbf{v}) - \mathbf{A}\right\| \leq \delta_{\mathbf{A}}$, and $\left\|\frac{\partial f}{\partial \mathbf{u}}(\mathbf{z}, \mathbf{v}) - \mathbf{B}\right\| \leq \delta_{\mathbf{B}},\ \forall (\mathbf{z}, \mathbf{v}) \in \mathcal{R}$. In simple words, the above assumption is that the Jacobians of the nonlinear system along the nominal trajectory remain sufficiently close to a single linearization and the gain $\mathbf{K}$ stabilizes all such perturbed linearizations. Thus, using a set $\mathfrak{L} = \left\{\left(\mathbf{A} + \Delta_{\mathbf{A}}, \mathbf{B} + \Delta_{\mathbf{B}}\right) \mid \|\Delta_{\mathbf{A}}\| \leq \delta_{\mathbf{A}},\ \|\Delta_{\mathbf{B}}\| \leq \delta_{\mathbf{B}}\right\}$, the gain K satisfies the stability condition $\rho(\mathbf{A}' + \mathbf{B}'\mathbf{K}) < 1,\ \forall (\mathbf{A}', \mathbf{B}') \in \mathfrak{L}$, where $\rho(\cdot)$ denotes the spectral radius. Thus, a RPI set for the time-invariant error dynamics $\mathbf{e}_{t+1} = \mathbf{A}_{\mathbf{K}} \mathbf{e}_t + \tilde{\mathbf{w}}_t$ is the Minkowski sum defined by $\mathcal{E}_{\min} = \bigoplus_{t=0}^{\infty} \mathbf{A}_{\mathbf{K}}^t \tilde{\mathcal{W}}$. To obtain a simple computable bound, the following theorem.

Lemma 1. *Suppose the matrix $\mathbf{A}_{\mathbf{K}}$ satisfies $\rho(\mathbf{A}_{\mathbf{K}}) < 1$. Then $\exists$ a symmetric $\mathbf{P} \succ 0$ such that $\mathbf{A}_{\mathbf{K}}^\top \mathbf{P} \mathbf{A}_{\mathbf{K}} - \mathbf{P} \prec 0$. Also, the induced $\mathbf{P}$-norm satisfies $\|\mathbf{A}_{\mathbf{K}}\|_{\mathbf{P}} < 1$.*

Proposition 1. *Suppose $\mathbf{A}_{\mathbf{K}}$ is Schur, and for $\mathbf{P} \succ 0$ it satisfies $\mathbf{A}_{\mathbf{K}}^\top \mathbf{P} \mathbf{A}_{\mathbf{K}} - \mathbf{P} \preceq -\mathbf{Q}$ for some $\mathbf{Q} \succ 0$ and $\tilde{w}_{\max,\mathbf{P}} = \sup\{\|\tilde{\mathbf{w}}\|_{\mathbf{P}} : \tilde{\mathbf{w}} \in \tilde{\mathcal{W}}\} < \infty$. Then the equation* (5) *satisfies $\sup_{t \in \mathbb{N}} \|\mathbf{e}_t\|_{\mathbf{P}} \leq \frac{\tilde{w}_{\max,\mathbf{P}}}{1 - \|\mathbf{A}_K\|_{\mathbf{P}}}$,*

Now as $\mathbf{z}_t = \mathbf{x}_t - \mathbf{e}_t \in \mathcal{X} \ominus \mathcal{E} = \mathcal{Z}$ and from (3) $\mathbf{v}_t = \mathbf{u}_t - \mathbf{K}\mathbf{e}_t \in \mathcal{U} \ominus \mathbf{K}\mathcal{E} = \mathcal{V}$. The feasibility of the nominal MPC problem follows from the compactness of the feasible inputs, the positive definiteness of $\mathbf{Q}$, $\mathbf{R}$, and the continuity of f. Now, the ancillary controller helps to steer the disturbed trajectory $\mathbf{x}$ as close as possible to the nominal trajectory $\mathbf{z}$ by minimizing $\mathbf{e}_t$ through the stabilizing controller $\mathbf{K}$. This can be mathematically expressed as another control problem [10] as follows :

$$\bar{V}^*(\mathbf{z}) = \min_{\mathbf{v}} \bar{\mathrm{V}}_N(\mathbf{z}, \mathbf{x}^{\mathbf{r}}, \mathbf{u}) = \sum_{t=0}^{N-1} \left(\|\mathbf{x}_t - \mathbf{z}_t^*\|_{\mathbf{Q}} + \|\mathbf{u}_t - \mathbf{v}_t^*\|_{\mathbf{R}}\right) \tag{6}$$

subject to $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t,\ \mathbf{x} \in \mathcal{X}^N,\ \mathbf{u} \in \mathcal{U}^N$. The stabilizing controller $\mathbf{K}$ [11], is computed by linearising the nominal system and solving the following discrete algebraic Riccati equation (DARE)

$$\begin{aligned} \mathbf{P} \quad &= \quad \mathbf{A}^\top \mathbf{P} \mathbf{A} - \mathbf{A}^\top \mathbf{P} \mathbf{B} (\mathbf{R} + \mathbf{B}^\top \mathbf{P} \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{P} \mathbf{A} + \mathbf{Q}, \\ \mathbf{K} \quad &= \quad -(\mathbf{R} + \mathbf{B}^\top \mathbf{P} \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{P} \mathbf{A}, \end{aligned} \tag{7}$$

where $\mathbf{A} \in \mathbb{R}^{d_x \times d_x}$ and $\mathbf{B} \in \mathbb{R}^{d_x \times d_u}$ are formed by decomposing the true system into $\mathbf{z}_{t+1} \approx f(\mathbf{z}_t, \mathbf{v}_t) + \mathbf{A}\mathbf{z}_t + \mathbf{B}\mathbf{v}_t + \mathbf{r}_t(\Delta\mathbf{z}_t, \Delta\mathbf{v}_t)$ through Taylor expansion, $\mathbf{R}$ and $\mathbf{Q}$ are the state and the control penalty matrices used in the optimization problem (4). Further by the assumption (2), $\rho(\mathbf{A}+\mathbf{BK}) < 1$. The error dynamics is given by $\mathbf{e}_{t+1} = (\mathbf{A}+\mathbf{BK})\mathbf{e}_t + r_t + \mathbf{w}_t = (\mathbf{A}+\mathbf{BK})\mathbf{e}_t + \tilde{\mathbf{w}}_t$, Here $\mathbf{r}_t = f(\mathbf{z}_t + \mathbf{e}_t, \mathbf{v}_t + \mathbf{K}\mathbf{e}_t) - f(\mathbf{z}_t, \mathbf{v}_t) - \mathbf{A}\mathbf{e}_t - \mathbf{BK}\mathbf{e}_t$ and $\mathbf{e}_{\max} = \dfrac{\tilde{w}_{\max,\mathbf{P}}}{1 - \|\mathbf{A}_K\|_{\mathbf{P}}}$. Thus, conservatively choosing the $\mathbf{e}_{\max}$ as radius, the true state $\mathbf{x}_t$ and the reference $\mathbf{x}_t^r$ can be made to lie within the tube with center at $\mathbf{z}_t$. The overall stability of the MPC problem has been demonstrated by using the Lyapunov theory in [12]. The nominal finite horizon cost $\mathrm{V}_N(\mathbf{z}, \mathbf{x}^r, \mathbf{v})$ serves as a Lyapunov function candidate. With appropriate terminal cost and terminal set, it is known that V_N decreases along nominal closed-loop trajectories. Combined with the bounded error tube, this guarantees that the real system state $\mathbf{x}_t = \mathbf{z}_t + \mathbf{e}_t$ remains close to the reference trajectory $\mathbf{x}_t^r$. Such guarantees have been established in several studies, such as [13]. Although tube MPC provides robustness guarantees, repeatedly solving the OCP (4) online makes it computationally expensive. Thus imitation of the MPC by a neural network is explored as a viable option to reduce the computational cost.

2.1 Neural imitation of tube MPC signals

This section elaborates on the methodology for using imitation learning to approximate the control signals of Tube MPC. The approach reduces the computational burden of online MPC by leveraging a neural network to mimic the nominal control $\mathbf{v}_t$, while maintaining stability guarantees through bounded generalization error. We focus on an underparameterized neural network inspired by the work [14]. The neural network used is a single-layer feed-forward neural network defined as $\pi_\theta(\mathbf{x}; \mathbf{W}, \mathbf{b}) = \dfrac{1}{\sqrt{m}} \sum_{i=1}^{m} \mathbf{b}_i \sigma(\mathbf{w}_i^\top \mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^n$ is the input, σ is an activation (Lipschitz) function, $\mathbf{W} = \{w_i \mid \in \mathbb{R}^d, i = 1, \cdots, m\}$ are the weights in the first layer, $\mathbf{b} = \{b_i \mid \in \mathbb{R}^m, i = 1, \cdots, m\}$ are the weights of the second layer, and the collection is denoted by $\theta = (\{w_i\}_{i=1}^m, \{b_i\}_{i=1}^m)$. This work primarily focuses on approximating the nominal control $\mathbf{v}(t)$ using a neural network. Thus given an initial states of the nominal system $\mathbf{z}_t$ and reference to follow $\mathbf{x}_{t+1}^r$ the neural network provides an approximate control signal which drives the nominal system $\mathbf{z}_{t+1}$ towards the reference $\mathbf{x}_{t+1}^r$. For convenience, we denote the imitated control and the corresponding state by $\mathbf{v}_{NN,t}$ $(= \pi_{\theta,t})$ and $\mathbf{z}_{NN,t+1}$ respectively. Thus $\mathbf{v}_{NN,t} : \mathcal{Z} \times \mathcal{X} \to \mathcal{V}$ and is defined by

$$\mathbf{v}_{NN,t}(\mathbf{z}_{NN,t}, \mathbf{x}_{t+1}^r; \mathbf{W}, \mathbf{b}) = \frac{1}{\sqrt{m}} \sum_{i=1}^{m} \mathbf{b}_i \sigma\left(\mathbf{w}_\mathbf{i}^\mathbf{T} \begin{bmatrix} \mathbf{z}_{NN,t} \\ \mathbf{x}_{t+1}^r \end{bmatrix}\right). \tag{8}$$

Let $\mathcal{F}$ denote the family of such networks defined as follows, which is used in your analysis

$$\mathcal{F} = \{\mathbf{v}_{NN}(\mathbf{W}, \mathbf{b}) : \|\mathbf{w}_r\| \leq S,\ \forall r \in [m], \|\mathbf{b}\|_2 \leq R\}, \tag{9}$$

This is used to formulate the corresponding nominal state $\mathbf{z}_{NN,t+1} = f(\mathbf{z}_{NN,t}, \mathbf{v}_{NN,t})$, where $\mathbf{z}_{NN,0}$ is the same initial condition as used for the OCP (4). This basically reduces the computational burden of the tube MPC by eliminating the need to solve equation (4) at each step. The predictor $v_{NN,t}$ induces an error and let it be $\mathbf{v}_t = \mathbf{v}_{NN,t} + \varepsilon_t$, where $\varepsilon = \max_{t \in \mathbb{N}} \|\varepsilon_t\|$. By construction, of the network $\mathbf{v}_{NN,t} \in \mathcal{V}$, and we assume ε is finite. This holds because $\mathcal{V}$ is compact, $\mathcal{F}$ approximates the MPC map to finite accuracy on the dataset, and $\mathcal{F}$ generalizes reasonably. Now let $\delta_t = \mathbf{z}_{NN,t} - \mathbf{z}_t$ then $\|\delta_{t+1}\| \leq L_z\|\delta_t\| + L_v\varepsilon$. Recursively, when $L_z \neq 1$ $\|\delta_t\| \leq L_v\varepsilon \sum_{k=0}^{t-1} L_z^k + L_z^t\|\delta_0\|$. Since they have the same initial condition for both the nominal trajectories, $\delta_0 = 0$. This helps to simplify the bound as $\|\delta_t\| \leq S_t\varepsilon$, where $S_t = L_v \sum_{k=0}^{t-1} L_z^k = L_v \dfrac{1 - L_z^t}{1 - L_z}$. This is true when the Lipschitz constant L_z is less than one. Hence, the overall error in the nominal dynamics is bounded within a ball of radius $S_t\varepsilon$. This provides a quantified guarantee on the neural network's tracking performance for the nominal system. Suppose the true control signal from the approximate signal lies within a ball of radius ε around the expert signal. Then, the maximum deviation between the approximate nominal system $\mathbf{z}_{NN,t}$ and the reference trajectory $\mathbf{x}_t^r$ can be bounded using the triangle inequality as $\|\mathbf{z}_{NN,t+1} - \mathbf{x}_{t+1}^r\| \leq S_t\varepsilon + \mathbf{e}_t^\mathrm{r}$, where $S_t\varepsilon$ is the propagated error due to the neural network approximation of the nominal control, and $\mathbf{e}_t^\mathrm{r} = \|\mathbf{z}_{t+1} - \mathbf{x}_{t+1}^r\|$ is the tracking error of the expert signal. This shows that the overall deviation of the neural-network-driven nominal system from the reference is bounded by the sum of the nominal MPC tracking error and the neural network's approximation error. Consequently, by ensuring that ε and the nominal MPC error e_t are sufficiently small, the neural-network controller can guarantee bounded tracking performance on the training dataset. Now, under the influence of the disturbance $\mathbf{w}_t$, the true system trajectory becomes as $\mathbf{x}_{NN,t+1} = f(\mathbf{x}_{NN,t}, \mathbf{u}_{NN,t}) + \mathbf{w}_t$, where $\mathbf{x}_{NN,t+1}$ is the evolution of the true system under the control $\mathbf{u}_{NN,t} = \mathbf{v}_{NN,t} + \mathbf{K}(\mathbf{z}_{NN,t} - \mathbf{x}_{NN,t})$. Let us denote the error between the nominal and the true trajectory generated by the neural network-based control as $\mathbf{e}_{NN,t}$. Thus $\mathbf{e}_{NN,t+1} = \mathbf{x}_{NN,t+1} - \mathbf{z}_{NN,t+1} = \mathbf{A_K}\mathbf{e}_{NN,t} + \mathbf{r}_{NN,t} + \mathbf{w}_t$, where $\|\mathbf{r}_{NN,t}\| \leq c_{rem}(\|\mathbf{e}_{NN,t}\|^2 + \|\mathbf{K}\mathbf{e}_{NN,t}\|^2)$. Since by the assumption on the original tube MPC, $\|\mathbf{A}+\mathbf{BK}\|$ is less than 1. Hence, the maximum of the error as $t \to \infty$ becomes $\mathbf{e}_{\max} = \max_{t\in\mathbb{N}} \|\mathbf{A_K}\mathbf{e}_t + \mathbf{r}_{NN,t} + \mathbf{w}_t\|$. Now since the network $\mathbf{v}_{NN,t}$ maps to $\mathcal{V}$ and again $\mathbf{z}_{NN,t}$ belongs to $\mathcal{Z}$, the maximum value of $\mathbf{r}_{NN,t}$ is the same as that of $\mathbf{r}_t$ that is $\tilde{r}$. Thus, $\mathbf{r}_{NN,t} + \mathbf{w}_t$ belongs to $\tilde{\mathcal{W}}$. Thus $e_{\max} = \dfrac{\tilde{w}_{max,\mathbf{P}}}{1-\|\mathbf{A_K}\|_{\mathbf{P}}}$. Now, the approximation of $\mathbf{v}_{NN,t}$ induces a shift effect on the tube structure due to the error in the nominal system $S_t\varepsilon$. Thus, it is important to enlarge the approximated tube to accommodate the reference. The center of the original and the approximated tube has the following relation $\mathbf{z}_{NN} - S_t\varepsilon \leq \mathbf{z} \leq \mathbf{z}_{NN} + S_t\varepsilon$. The original tube with boundaries $\mathbf{z}-\mathbf{e}_{\max}, \mathbf{z}+\mathbf{e}_{\max}$ is capable to accommodate the distrubed dynamics $\mathbf{x}$ and reference $\mathbf{x}^r$. Similarly, the approximated tube with boundaries $[\mathbf{z}_{NN} - \mathbf{e}_{\max}, \mathbf{z}_{NN} + \mathbf{e}_{\max}]$ is capable to accommodating the disturbed dynamics $\mathbf{x}_{NN}$ as noticed from the $\mathbf{e}_{\max}$. However, if the error ε_v is high, the approximated tube does not accommodate the reference $\mathbf{x}^r$ or the true state x. Thus, for the approximated tube to accommodate the trajectories $\mathbf{x}$, $\mathbf{x}^r$, and $\mathbf{x}_{NN}$, it should be enlarged by the $S_t\varepsilon$. Thus, the boundaries used are $\mathbf{z}_{NN} - C(S_t\varepsilon) - \mathbf{e}_{max}, \mathbf{z}_{NN} + C(S_t\varepsilon) + \mathbf{e}_{max}$. Due to conservatism we use a scaling factor C. Further, as the network approximation capability becomes better, that is, when the generalization error of the network improves, the value of ε reduces to zero. This makes the tube size be influenced by e_{max}. Now, we connect the deterministic propagation bounds above with statistical bounds on ε obtained from training and generalization theory. That is, we intend to link the tube radius to the worst-case error using the network's generalization error, thereby improving performance when the test dataset follows the distribution of the training dataset. Before this, we present the training procedure and the metrics used. The training is done so that the network is capable of producing control signals $\mathbf{v}_{NN}$ that steer all state-reference pairs $(\mathbf{z},\mathbf{v})$ in the space $\mathcal{Z}\times\mathcal{X}$. Suppose $\mathbf{v}$, $\mathbf{z}$ and $\mathbf{x^r}$ follows a distribution $\mathcal{D}$, then the task is to minimize $\mathcal{L}_{\mathcal{D}}(\mathbf{v}_{NN}) = \mathbb{E}_{(\mathbf{z},\mathbf{x}^r,\mathbf{v})\sim\mathcal{D}}[\ell(\mathbf{v}_{NN}(\mathbf{z},\mathbf{x}^r),\mathbf{v})]$, where $\ell : \mathcal{Z}\times\mathcal{X}\times\mathcal{V} \to \mathbb{R}_{\geq 0}$ and the common choice is $\ell(\mathbf{v}_{NN}(\mathbf{z},\mathbf{x}^r),\mathbf{v}) = \frac{1}{2}\|\mathbf{v}_{NN}(\mathbf{z},\mathbf{x}^r) - \mathbf{v}\|_2^2$. Due to the computational intractability of this population loss function, in this work we focus on the empirical risk minimization problem using the training dataset with n samples $S = \{\mathbf{z}_t, \mathbf{x}_t^r, \mathbf{u}_t\}_{t=1}^n$ drawn iid from the underlying data distribution $\mathcal{D}$ over $\mathcal{Z}\times\mathcal{X}\times\mathcal{V}$. The training of the network is performed by minimizing $\mathcal{L}_S = \dfrac{1}{n}\sum_{t=1}^{n} \ell(\mathbf{v}_{NN}(\mathbf{x}_t, \mathbf{x}_{t+1}^r), \mathbf{v}_t)$, over the space $\mathcal{F}$. The performance of the network outside the samples in S in relative to its performance in S is defined as generalization error of the network mathematically described as $\mathcal{L}_{\mathcal{D}}(\mathbf{v}_{NN}) - \mathcal{L}_S(\mathbf{v}_{NN})$ for network $\mathbf{v}_{NN}$ from the class of functions $\mathcal{F}$, trained over a given set of samples S. The Rademacher complexity bounds this error in ML theory.

Theorem 1. *[14] Suppose the squared error $\ell(\mathbf{v}_{NN},\mathbf{v}) = \frac{1}{2}(\mathbf{v}_{NN}-\mathbf{v})^2$ is bounded in $[0,c]$ and is ρ-Lipschitz in the first argument. Then with a proability of $1-\delta_1$ over the sample S of size n :*

$$\sup_{\mathbf{v}_{NN}\in\mathcal{F}} |\mathcal{L}_{\mathcal{D}}(\mathbf{v}_{NN}) - \mathcal{L}_S(\mathbf{v}_{NN})| \leq 2\rho\mathcal{R}_S(\mathcal{F}) + 3c\sqrt{\frac{\log\left(\frac{2}{\delta_1}\right)}{2n}}.$$

The following theorem gives the generalisation bound on the single-layered neural network $\mathbf{v}_{NN} \in \mathcal{F}$ considered in the study :

Theorem 2. *Given $R > 0$, with probability at least $1-\delta_1$ for every $S > 0$, the function class $\mathcal{F}$ with empirical Rademacher complexity $\mathcal{R}_S(\mathcal{F}) \leq \dfrac{\|b\|_2 RSR_x}{\sqrt{n}}$.*

Theorem 3. *Consider the squared loss ℓ, and suppose ℓ is bounded in $[0,c]$ and is ρ-Lipschitz in its* *first argument. Then with probability at least* $1-\delta_1$ *over sample S* $\sup_{\mathbf{v}_{NN}\in\mathcal{F}}\{L_D(\mathbf{v}_{NN})-L_S(\mathbf{v}_{NN})\}\leq\Gamma$ *with*

$$\Gamma=\frac{2\rho\|b\|_2 LSR_x}{\sqrt{n}}+3c\sqrt{\frac{\log\left(\frac{2}{\delta_1}\right)}{2n}}.$$

The generalization error provides an estimate of the neural network's expected error. Further, in an under-parametrized setting with the considered single-layer neural network, it is easier to numerically evaluate it. The following theory estimates the worst-case error in terms of generalisation .

Theorem 4. *Let* $\mathbf{v}_{NN}$ *be a fixed predictor after full training,* ℓ *denote the squared sample loss, Let* $(\mathbf{z},\mathbf{x}^r,\mathbf{v})\sim\mathcal{D}$*,* $S=\{(\mathbf{z}_t,\mathbf{x}_t^r,\mathbf{v}_t)\}_{t=1}^n$ *denote the training sample and let* $\{(\mathbf{z}_t^{\text{test}},\mathbf{x}_t^{\text{test},r},\mathbf{v}_t^{\text{test}})\}_{t=1}^{n_{\text{test}}}$ *denote an i.i.d. test sample drawn from the data distribution* $\mathcal{D}$*. Then, the centered random variable* $Z=\ell(\mathbf{v}_{NN}(\mathbf{z},\mathbf{x}^r,\mathbf{v})-\mathcal{L}_{\mathcal{D}}(\mathbf{v}_{NN})$ *is sub-gaussian for* $(\mathbf{z},\mathbf{x}^r,\mathbf{v})\sim\mathcal{D}$ *and there exists* $c>0$ *such that with probability of at least* $1-\delta$ $\max_{1\leq t\leq n_{\text{test}}}(Z_t)\leq c\sqrt{2\log\left(\frac{n_{\text{test}}}{\delta}\right)}$*. Further, with probability of at least* $1-(\delta_1+\delta)$ *jointly over the training and test draws, the following worst-case approximation bound holds :*

$$\max_{1\leq t\leq n_{\text{test}}}\ell(\mathbf{v}_{NN}(\mathbf{z}_t^{\text{test}},\mathbf{x}_t^{\text{test}}),\mathbf{y}_t^{\text{test}})\leq\mathcal{L}_S(\mathbf{v}_{NN})+\Gamma+c\sqrt{2\log\left(\frac{n_{\text{test}}}{\delta}\right)}.$$

Using the above theorem, we obtain a high-probability bound on the maximum squared loss at test points. But we require the worst case approximation error. If $\ell(\mathbf{v}_{NN},\mathbf{v})=\frac{1}{2}\|\mathbf{v}_{NN}-\mathbf{v}\|^2\leq\eta$. Then we have $\varepsilon_{test}\leq\sqrt{2\eta}$. Therefore, with probability at least $1-(\delta_1+\delta)$ $\varepsilon_{\text{test}}=\max_{1\leq j\leq n_{\text{test}}}\|v_{NN}(\mathbf{z}_j^{\text{test}},\mathbf{x}_j^{\text{test}})-\mathbf{v}_j^{\text{test}}\|\leq\sqrt{2\left(\mathcal{L}_S(\mathbf{v}_{NN})+\Gamma+c\sqrt{2\log\left(\frac{n_{\text{test}}}{\delta}\right)}\right)}$. Thus, with probability at least $1-(\delta_1+\delta)$, the actual trajectory under the approximate tube controller and the reference trajectory both lie in the conservative tube centered at $z_{NN,t}$:

$$\|\mathbf{x}_{NN,t}-\mathbf{z}_{NN,t}\|\leq\mathbf{e}_{\max},\quad\|\mathbf{z}_{NN,t}-\mathbf{x}_t^r\|\leq S_t\varepsilon_{\text{test}}+\|\mathbf{z}_t-\mathbf{x}_t^r\|,\tag{10}$$

hence the conservative tube that covers both disturbed true trajectories and references is $\Big\{x\mid\|\mathbf{x}-\mathbf{z}_{NN,t}\|\leq\mathbf{e}_{\max}+C(S_t\varepsilon_{\text{test}}+\|\mathbf{z}_t-\mathbf{x}_t^r\|)\Big\}$, where $C>0$ is chosen to make the tube bounds less conservative and realistic, it helps to overcome the traditional bound. Thus with probability of atleast $1-\delta_1+\delta$, the tube $\left\{x\mid\|\mathbf{x}-\mathbf{z}_{NN,t}\|\leq\mathbf{e}_{\max}+S_t\left(\mathcal{L}_S(\mathbf{v}_{NN})+\Gamma+c\sqrt{2\log\left(\frac{n_{\text{test}}}{\delta}\right)}\right)+\|\mathbf{z}_t-\mathbf{x}_t^r\|\right\}$ contains the reference $\mathbf{x}_t^r$ and the disturbed trajectory $\mathbf{x}_{NN,t}$, $\mathbf{x}_t$.

3 Numerical Validation

This section validates the proposed theoretical results on a nonlinear two-dimensional oscillator of the following form :

$$\begin{aligned}\dot{x}_1&=x_2,\\\dot{x}_2&=\mu(1-x_1^2)x_2-x_1+u.\end{aligned}\tag{11}$$

This is a Van der Pol oscillator. The input and state constraints are given by $|u|\leq 1$, $|x_1|\leq 1$, and $|x_2|\leq 1$. The parameter is fixed to $\mu=0.01$. This value is chosen to ensure that the local Lipschitz constant L_z is approximately 1 in the region of operation. The continuous system is discretised using a sampling time of $T_s=0.5$ s and expressed as $\mathbf{z}_{t+1}=A\mathbf{z}_t+B\mathbf{v}_t+f(\mathbf{z}_t,\mathbf{v}_t)$, where $A=\begin{bmatrix}1&0.5\\-0.5&1.5\end{bmatrix}$, and $B=\begin{bmatrix}0\\0.5\end{bmatrix}$. These are obtained by linearization around the the equilibrium point $\hat{\mathbf{z}}=(0,0)$, and $\hat{\mathbf{v}}=0$. The LQR feedback gain is computed using the DARE and obtained as $K=\begin{bmatrix}0.0979&-1.4397\end{bmatrix}$. For which we chose $Q=\begin{bmatrix}1&0\\0&1\end{bmatrix}$, $R=1$, Using Monte-Carlo sampling, the $\delta_A=0.0105$ and $\delta_B=0$ were estimated. Using these, the value of e_{max}, that is the invariant error bound, was obtained as $e_{\max}=\begin{bmatrix}0.0279\\0.0293\end{bmatrix}$, Thus, the tightened state constraint set was $\mathcal{Z} = [-0.9720, 0.9720] \times [-0.9706, 0.9706]$ and the control set was obtained as $\mathcal{V} = [-0.0.9604, 0.9604]$. The reference for the first state x_1 was generated as a piecewise-constant signal sampled uniformly within $\mathcal{Z}$, with changes or jumps occurring every 50 time steps. The expert tube-MPC controller was implemented using the CasADi optimization framework in Python. It achieved a mean tracking error of 0.0564.

Now to imitate the behaviour with a neural network, the dataset $\mathcal{S} = \{(z_t, x_t^r, v_t)\}_{t=1}^{1000}$ was collected from the expert controller. We use a single-hidden-layer neural network with 128 neurons and an input dimension of 3 (z_1, z_2, x_1^r) that was trained to approximate the nominal control law v_t. The network employed a ReLU activation in the hidden layer and a sigmoid output scaled to the control bounds. The network was trained for 3000 epochs using the Adam optimizer on 900 samples. The remaining 100 samples were used for testing. The final training loss achieved was $\mathcal{L}_S = 4.22 \times 10^{-3}$, and the estimated generalisation gap by using the test dataset was $|\mathcal{L}_\mathcal{D} - \mathcal{L}_S| = 2 \times 10^{-3}$. The theoretical bound on $|\mathcal{L}_\mathcal{D} - \mathcal{L}_S|$ was estimated as $\Gamma = 0.1560$. To maintain practical relevance, this bound was conservatively scaled by a factor of $C = 10^{-4}$. The inflation factor from the states $S_t = 64.66$, and the updated tube radius became $e_{\max} = \begin{bmatrix} 0.054 \\ 0.055 \end{bmatrix}$. Using the original tube, the true state was found to be out of the invariant set at 36 time instants. With the new calculated inflated tube, the number of violations was reduced to 10. Furthermore, this could be improved by either refining the scaling factor. The closed-loop performance and tube containment are shown in Figure 1.

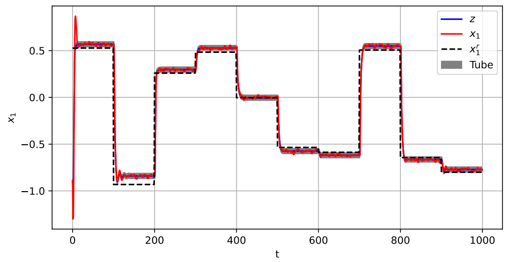

(a) The tube, reference, nominal, and disturbed system trajectory with the same tube structure used by the expert MPC.

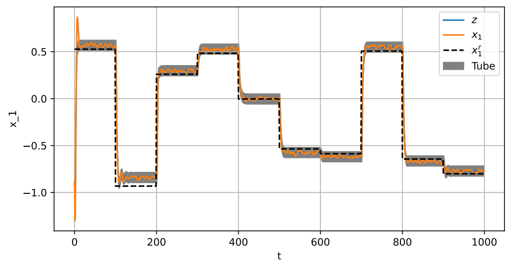

(b) The tube, reference, nominal, and disturbed system trajectory under the inflated tube.

FIGURE 1 – Comparison of the inflated tube radius and the original tube radius when used for the approximate controller.

4 Conclusion

This work analyzed the inflation of the tube size when a neural network approximation replaces the tube MPC expert signal. Furthermore, it introduces a new probabilistic bound on the worst-case approximation error of the neural network, based on the network's generalisation gap. This was utilized to inflate the tube radius with a scaling factor. The proposed theory relies on Lipschitz continuity and applies to systems that have contractive dynamics. Numerical validations were performed to validate the proposed results on a Van der Pol oscillator. Using the new tube radius, it was observed that the number of instances in which the disturbed system violates the constraints decreased. However, the conservatism persists due to the worst-case nature of the bound. It could be mitigated by choosing an appropriate scaling factor, which, again, becomes a hyperparameter. Further, the proposed structure works only when the inflation factor from the states is bounded or when the system is contractive. Future work could extend this analysis to a more general system without the Lipschitzian assumption or contractiveness. This could benefit the use of approximate MPC within the tube MPC framework.

Acknowledgement

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme-ERC-2020-COG-ERC Consolidator Grant

(Grant Agreement No. 101002857, DREAM-ON Project).

Références

[1] Hussein, A., Gaber, M. M., Elyan, E., & Jayne, C. (2017). Imitation learning : A survey of learning methods. ACM Computing Surveys (CSUR), 50(2), 1-35.

[2] Peters, J., & Schaal, S. (2008). Learning to control in operational space. The International Journal of Robotics Research, 27(2), 197-212.

[3] Ross, S., Gordon, G., & Bagnell, D. (2011, June). A reduction of imitation learning and structured prediction to no-regret online learning. In Proceedings of the fourteenth international conference on artificial intelligence and statistics (pp. 627-635). JMLR Workshop and Conference Proceedings.

[4] Paulson, J. A., & Mesbah, A. (2018). Nonlinear model predictive control with explicit backoffs for stochastic systems under arbitrary uncertainty. IFAC-PapersOnLine, 51(20), 523-534.

[5] Fleming, J., Kouvaritakis, B., & Cannon, M. (2014). Robust tube MPC for linear systems with multiplicative uncertainty. IEEE Transactions on Automatic Control, 60(4), 1087-1092.

[6] Hertneck, M., Köhler, J., Trimpe, S., & Allgöwer, F. (2018). Learning an approximate model predictive controller with guarantees. IEEE Control Systems Letters, 2(3), 543-548.

[7] Wang, R., Li, H., Liang, B., Shi, Y., & Xu, D. (2023). Policy learning for nonlinear model predictive control with application to USVs. IEEE Transactions on Industrial Electronics, 71(4), 4089-4097.

[8] Tong, J., Du, S., & Fan, W. (2025). Ensemble Neural Network-Based Approximate Model Predictive Control With Strict Guarantees. International Journal of Robust and Nonlinear Control.

[9] Köhler, J., Müller, M. A., & Allgöwer, F. (2018). Nonlinear reference tracking : An economic model predictive control perspective. IEEE Transactions on Automatic Control, 64(1), 254-269.

[10] Mayne, D. Q., & Kerrigan, E. C. (2007). Tube-based robust nonlinear model predictive control1. IFAC Proceedings Volumes, 40(12), 36-41.

[11] Jezierski, A., Mozaryn, J., & Suski, D. (2017, June). A comparison of LQR and MPC control algorithms of an inverted pendulum. In Polish Control Conference (pp. 65-76). Cham : Springer International Publishing.

[12] Rawlings, J. B., Mayne, D. Q., & Diehl, M. (2020). Model predictive control : theory, computation, and design (Vol. 2). Madison, WI : Nob Hill Publishing.

[13] Mayne, D. Q., Kerrigan, E. C., Van Wyk, E. J., & Falugi, P. (2011). Tube-based robust nonlinear model predictive control. International journal of robust and nonlinear control, 21(11), 1341-1353.

[14] Xavier, D. M., Chamoin, L., & Fribourg, L. (2023). Training and Generalization Errors for Underparameterized Neural Networks. IEEE Control Systems Letters, 7, 3926-3931.

[15] Giné, E., & Nickl, R. (2021). Mathematical foundations of infinite-dimensional statistical models. Cambridge University Press.