

Lipschitz-bounded 1D convolutional neural networks using the Cayley transform and the controllability Gramian

Patricia Pauli¹, Ruigang Wang², Ian R. Manchester² and Frank Allgöwer¹

Abstract— We establish a layer-wise parameterization for 1D convolutional neural networks (CNNs) with built-in end-to-end robustness guarantees. In doing so, we use the Lipschitz constant of the input-output mapping characterized by a CNN as a robustness measure. We base our parameterization on the Cayley transform that parameterizes orthogonal matrices and the controllability Gramian of the state space representation of the convolutional layers. The proposed parameterization by design fulfills linear matrix inequalities that are sufficient for Lipschitz continuity of the CNN, which further enables unconstrained training of Lipschitz-bounded 1D CNNs. Finally, we train Lipschitz-bounded 1D CNNs for the classification of heart arrhythmia data and show their improved robustness.

I. INTRODUCTION

Robustness of neural networks (NNs) has lately been a topic of increasing importance, for which the Lipschitz constant of the NN's input-output mapping has become a common metric [1]. For a low Lipschitz constant, a slightly perturbed input admits only small changes in the output which is desirable in robust NNs. Finding an accurate upper bound on an NN's Lipschitz constant has been broadly tackled, e.g. using relaxations by quadratic constraints [2], [3], average operators [4] and polynomial optimization [5]. In addition, the training of provably Lipschitz-bounded NNs was proposed by including constraints [6], [7] and regularization techniques [8]. While effective, one drawback of these methods is the computational overhead coming from constraints and projections in the optimization problem [7].

To overcome this, [9], [10], [11] suggest so-called direct parameterizations for equilibrium networks, recurrent equilibrium networks, and feedforward neural networks, respectively, with guaranteed Lipschitz bounds. From a set of unconstrained variables [9], [10], [11] formulate the NNs in such a way that they by design satisfy linear matrix inequalities (LMIs). These LMIs in turn are sufficient conditions for Lipschitz continuity such that, this way, one can parameterize the class of Lipschitz-bounded NNs with a Lipschitz upper bound predefined by the user. The underlying training problem boils down to an unconstrained optimization problem that can be solved using gradient methods. In

this work, we take the same approach as in [9], [10], [11] to parameterize Lipschitz-bounded 1D convolutional neural networks (CNNs).

CNNs play an important role in deep learning and have been tremendously successful in image and audio processing tasks [12], [13], [14]. We distinguish between 1D and 2D CNNs: By 1D CNNs, we mean CNNs with inputs with one propagation dimension such as time signals, whereas 2D CNNs take inputs with two propagation dimensions, e.g., pictures. This paper focuses on 1D CNNs, that typically consist of convolutions (= finite impulse response (FIR) filters), nonlinear activation functions, that are slope-restricted, pooling layers, and linear layers that are concatenated in a feedforward structure. 1D CNNs parameterize a rich class of nonlinear systems and are utilized as classifiers for time signals. An extension to 2D CNNs based on [15] is left for future work.

While numerous methods exist for enforcing Lipschitz continuity and orthogonality in fully connected layers [16], the design of Lipschitz-bounded convolutional layers and CNNs is less studied and often restricted to special convolutions [17]. Recently, this has been approached via parameterization of convolutional layers in the Fourier domain [11], [17], however, parameterizations in [11] are built for a particular input size, whereas our new parameterization in state space is independent of the input dimension, and can even be applied causally to a signal on $[0, \infty)$. Another feature of our approach is that we impose Lipschitz continuity directly onto the input-output mapping rather than on all individual layers, like it is done in many other works [18], using that the product of the Lipschitz bounds of the layers yields a Lipschitz bound for the overall NN. In this way, our approach shows reduced conservatism in the compliance with the Lipschitz bound, i.e., our CNNs have higher expressivity for the same Lipschitz bound. In addition our approach accounts for standard pooling layers, which were not addressed in other recent Lipschitz-bounded parameterizations of CNNs [11].

Our main contribution is a scalable and expressive layer-wise parameterization of Lipschitz-bounded 1D CNNs that makes use of the Cayley transform to parameterize orthogonal matrices. Beside the Cayley transform, a tool that was used for NN parameterization before, we newly propose to utilize the controllability Gramian in the context of parameterizing convolutional layers of Lipschitz-bounded CNNs. In particular, we reformulate parts of the underlying LMI, that enforces dissipativity onto convolutional layers, as a Lyapunov equation whose unique analytical solution is the

*This work was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016 and under grant 468094890. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Patricia Pauli.

¹Patricia Pauli and Frank Allgöwer are with the Institute for Systems Theory and Automatic Control, University of Stuttgart, 70569 Stuttgart, Germany patricia.pauli@ist.uni-stuttgart.de

²Ruigang Wang and Ian R. Manchester are with the Australian Centre for Robotics and School of Aerospace, Mechanical and Mechatronic Engineering, The University of Sydney, Australia.

controllability Gramian. Using our parameterization, we then train Lipschitz-bounded 1D CNNs solving an unconstrained optimization problem.

The remainder of the paper is structured as follows: Section II first introduces 1D CNNs and formally states the training problem. In Section III, we discuss preliminaries, including state space representations for 1D convolutions and Lipschitz constant estimation for 1D CNNs. In Section IV, we present our direct parameterization for Lipschitz-bounded 1D CNNs and in Section V, we train Lipschitz-bounded 1D CNNs on the MIT-BIH arrhythmia database [19], a well-known benchmark dataset for 1D CNNs. Finally, in Section VI, we conclude the paper.

Notation: By \mathbb{D}^n (\mathbb{D}_+^n) and \mathbb{S}^n (\mathbb{S}_+^n), we denote the set of n -dimensional (positive definite) diagonal and symmetric matrices, respectively, and by \mathbb{N}_+ the natural numbers without zero. \mathcal{I} is a set of indices with elements $i \in \mathbb{N}_+$, and $|\mathcal{I}|$ gives the number of elements in the index set \mathcal{I} .

II. PROBLEM STATEMENT

We consider 1D CNNs that are a concatenation of convolutional layers $\mathcal{C}_i: \mathbb{R}^{c_{i-1} \times N_{i-1}} \rightarrow \mathbb{R}^{c_i \times N_i}$ with indices $i \in \mathcal{I}_C$, and fully connected layers $\mathcal{L}_i: \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ with indices $i \in \mathcal{I}_F$

$$\text{CNN}_\theta = \mathcal{L}_1 \circ \dots \circ \mathcal{L}_{p+1} \circ F \circ \mathcal{C}_p \circ \dots \circ \mathcal{C}_1, \quad (1)$$

adding up to a total number of $l = |\mathcal{I}_C| + |\mathcal{I}_F|$ layers. Here, N_i denotes the signal length, c_i the channel size, and n_i the layer dimension of the respective i -th layer. To transition from the fully convolutional part of the CNN to the fully connected part, we necessitate a flattening operation $F: \mathbb{R}^{c_p \times N_p} \rightarrow \mathbb{R}^{n_p}$ that operates on the output of the p -th (last) convolutional layer with $n_p = c_p N_p$.

A *convolutional* layer consists of two to three stages, a convolution operation, a nonlinear activation, and possibly a pooling operation. The first two stages are

$$\tilde{\mathcal{C}}_i: w_k^j = \phi_i \left(b_i + \sum_{j=0}^{\ell_i-1} K_j^i w_{k-j}^{i-1} \right), \quad k=0, \dots, N_i-1 \quad \forall i \in \mathcal{I}_C, \quad (2)$$

with convolution kernel $K_j^i \in \mathbb{R}^{c_i \times c_{i-1}}$, $j=0, \dots, \ell_i-1$, kernel size ℓ_i , and bias $b_i \in \mathbb{R}^{c_i}$. First, a convolution on the signal $w^{i-1} \in \mathbb{R}^{c_{i-1} \times N_{i-1}}$ is applied and subsequently, the nonlinear activation function $\phi_i: \mathbb{R}^{c_i} \rightarrow \mathbb{R}^{c_i}$ is evaluated element-wise to obtain the output $w^i \in \mathbb{R}^{c_i \times N_{i-1}}$. Oftentimes, a convolutional layer additionally contains pooling layers $\mathcal{P}_i: \mathbb{R}^{c_i \times N_{i-1}} \rightarrow \mathbb{R}^{c_i \times N_i}$ to downsample the signal w^i . We consider maximum pooling

$$\mathcal{P}_i^{\max}: \tilde{w}_k^i = \max_{j=1, \dots, \ell_i} w_{\ell_i(k-1)+j}^i, \quad k=0, \dots, N_i-1, \quad \forall i \in \mathcal{I}_P^{\max},$$

and average pooling

$$\mathcal{P}_i^{\text{av}}: \tilde{w}_k^i = \frac{1}{\ell_i} \sum_{j=1}^{\ell_i} w_{\ell_i(k-1)+j}^i, \quad k=0, \dots, N_i-1, \quad \forall i \in \mathcal{I}_P^{\text{av}},$$

where $\mathcal{I}_P^{\text{av}} \cup \mathcal{I}_P^{\max} \subseteq \mathcal{I}_C$. As a result, the convolutional layer becomes $\mathcal{C}_i = \mathcal{P}_i \circ \tilde{\mathcal{C}}_i$ in case a pooling layer is added or $\mathcal{C}_i =$

$\tilde{\mathcal{C}}_i$ otherwise. Finally, a CNN typically holds *fully connected* layers, which we define as mappings

$$\begin{aligned} \mathcal{L}_i: w^i &= \phi_i(W_i w^{i-1} + b_i) \quad \forall i \in \mathcal{I}_F \setminus \{l\}, \\ \mathcal{L}_l: w^l &= W_l w^{l-1} + b_l \end{aligned} \quad (3)$$

with weights $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$, biases $b_i \in \mathbb{R}^{n_i}$ and activation functions $\phi_i: \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ that are applied element-wise.

The 1D CNN $f_\theta(w^0) = w^l$ is hence characterized by $\theta = \{(K^i, b_i)_{i=1}^p, (W_i, b_i)_{i=p+1}^l\}$ and the chosen activation and pooling operations. In this work, we present a direct parameterization for Lipschitz-bounded 1D CNNs (1).

Problem 1: Find a parameterization $\kappa \mapsto \theta$ of f_θ for a predefined Lipschitz bound $\rho > 0$ such that all 1D CNNs parameterized by κ are ρ -Lipschitz continuous with respect to the ℓ_2 norm, i. e., they satisfy

$$\|f_\theta(x) - f_\theta(y)\|_2 \leq \rho \|x - y\|_2 \quad \forall x, y \in \mathbb{R}^n. \quad (4)$$

In the case of multiple channels c , $n = cN$ denotes the stacked up version of the input. Note that $\|\cdot\|_2$ in (4) can either be interpreted as the Euclidean norm of a vector-valued input x or as the ℓ_2 norm of a signal x .

To train a Lipschitz-bounded CNN, we minimize a learning objective $\mathcal{L}(\theta)$, e. g., the mean squared error, the cross-entropy loss or, to encourage robustness through the learning objective a tailored loss, e.g. the hinge loss [20], while at the same time enforcing Lipschitz-boundedness onto the CNN. Rather than solving a training problem subject to a Lipschitz constraint, which can get computationally involved, i. e.,

$$\min_{\theta} \mathcal{L}(\theta) \quad \text{s.t.} \quad f_\theta \text{ is Lipschitz-bounded,}$$

the suggested parameterization $\kappa \mapsto \theta$ allows to solve an unconstrained training problem over κ using gradient methods

$$\min_{\kappa} \mathcal{L}(\theta(\kappa)).$$

III. PRELIMINARIES

Before we state the parameterization of Lipschitz-bounded 1D CNNs in Section IV, we introduce a compact formulation of convolutions in state space and state LMI conditions that certify Lipschitz boundedness and that can be used to estimate the Lipschitz constant for 1D CNNs [3]. In addition, we introduce the Cayley transform used to parameterize orthogonal matrices.

A. State space representation for convolutions

To formulate LMI conditions for convolutional layers, we can either reformulate the convolutional operation as a fully connected layer characterized by a sparse and redundant Toeplitz matrix [7] that scales with the input dimension or, as suggested in [3], we can compactly state the convolution, i. e., an FIR filter, in state space, completely independent of the input signal length. A possible discrete-time state space representation of the i -th convolutional layer (2) with state $x_k^i \in \mathbb{R}^{n_{x_i}}$ and state dimension $n_{x_i} = (\ell_i - 1)c_{i-1}$ is

$$\begin{aligned} x_{k+1}^i &= A_i x_k^i + B_i w_k^{i-1}, \\ y_k^i &= C_i x_k^i + D_i w_k^{i-1} + b_i, \\ w_k^i &= \phi(y_k^i), \end{aligned} \quad (5)$$

where

$$A_i = \begin{bmatrix} 0 & I & & 0 \\ 0 & 0 & \ddots & \\ \vdots & & \ddots & I \\ 0 & \dots & & 0 \end{bmatrix}, \quad B_i = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ I \end{bmatrix}, \quad (6a)$$

$$C_i = [K_{\ell_i-1}^i \quad \dots \quad K_1^i], \quad D_i = K_0^i. \quad (6b)$$

Note that 2D convolutions also admit a state space realization, namely as a 2D system [15], based on which our parameterization can potentially be extended to end-to-end Lipschitz-bounded 2D CNNs.

B. Lipschitz constant estimation

The Lipschitz constant is a sensitivity measure to changes in the input, which is commonly used to verify robustness for NNs [1]. Since, however, the calculation of the true Lipschitz constant is an NP-hard problem, an accurate upper bound is sought instead. For this purpose, we over-approximate the nonlinear activation functions by their slope-restriction cone [2], [6]. Commonly used activation functions $\varphi: \mathbb{R} \rightarrow \mathbb{R}$, such as ReLU and tanh, are slope-restricted in $[0, 1]$, i. e.,

$$0 \leq \frac{\varphi(x) - \varphi(y)}{x - y} \leq 1 \quad \forall x, y \in \mathbb{R}.$$

Based on this element-wise property, we formulate an incremental quadratic constraint

$$\begin{bmatrix} \phi(x) - \phi(y) \\ x - y \end{bmatrix}^\top \begin{bmatrix} -2\Lambda & \Lambda \\ \Lambda & 0 \end{bmatrix} \begin{bmatrix} \phi(x) - \phi(y) \\ x - y \end{bmatrix} \geq 0 \quad \forall x, y \in \mathbb{R}^n, \quad (7)$$

which, by the multipliers $\Lambda \in \mathbb{D}_+^n$, is a conic combination of the element-wise slope-restriction constraint and a suitable over-approximation of the nonlinearities in all n neurons. The following theorem states a set of l LMI conditions that serve as a sufficient condition for Lipschitz continuity for 1D CNNs based on the relaxation (7) [3].

Theorem 2 ([3]): Let CNN_θ and $\rho > 0$ be given and let all activation functions be slope-restricted in $[0, 1]$. If there exist

- (i) $Q_i \in \mathbb{S}^{c_i}$ ($Q_i \in \mathbb{D}^{c_i}$ if a convolutional layer contains a maximum pooling layer), $P_i \in \mathbb{S}_+^{m_{x_i}}$, and $\Lambda_i \in \mathbb{D}_+^{c_i}$ such that $\forall i \in \mathcal{I}_C$

$$\left[\begin{array}{cc|c} P_i - A_i^\top P_i A_i & -A_i^\top P_i B_i & -C_i^\top \Lambda_i \\ -B_i^\top P_i A_i & Q_{i-1} - B_i^\top P_i B_i & -D_i^\top \Lambda_i \\ \hline -\Lambda_i C_i & -\Lambda_i D_i & 2\Lambda_i - Q_i \end{array} \right] \succeq 0, \quad (8)$$

where $Q_0 = \tilde{\rho}^2 I$,

- (ii) $Q_i \in \mathbb{S}^{n_i}$ and $\Lambda_i \in \mathbb{D}_+^{n_i}$ such that $\forall i = \mathcal{I}_F \setminus \{l\}$

$$\left[\begin{array}{cc} Q_{i-1} & -W_i^\top \Lambda_i \\ -\Lambda_i W_i & 2\Lambda_i - Q_i \end{array} \right] \succeq 0 \quad \text{and} \quad \left[\begin{array}{cc} Q_{l-1} & -W_l^\top \\ -W_l & I \end{array} \right] \succeq 0, \quad (9)$$

where $Q_p := I_{N_p} \otimes Q_p$,

then the CNN_θ is ρ -Lipschitz continuous with $\rho = \tilde{\rho} \prod_{s \in \mathcal{I}_p} \mu_s$, where μ_s are the Lipschitz constants of the average pooling layers.

The proof of Theorem 2 is based on the dissipativity of the individual layers of CNN_θ that are connected in a feedforward fashion, see [3] for details and the proof. Note that the matrix Q_{i-1} , which is a gain matrix, links the i -th layer to the previous layer and by this interconnection we can finally analyze Lipschitz continuity of the input-output mapping $w^l = \text{CNN}_\theta(w^0)$ [3].

Based on Theorem 2, we can determine an upper bound on the Lipschitz constant for a given CNN solving a semidefinite program

$$\min_{\rho^2, \Lambda, P, Q} \rho^2 \quad \text{s.t.} \quad (8), (9), \quad (10)$$

where $\Lambda = \{\Lambda_i\}_{i \in \mathcal{I}_C \cup \mathcal{I}_F \setminus \{l\}}$, $Q = \{Q_i\}_{i \in \mathcal{I}_C \cup \mathcal{I}_F \setminus \{l\}}$, $P = \{P_i\}_{i \in \mathcal{I}_C}$ serve as decision variables together with ρ^2 .

C. Cayley transform

Typically, the Cayley transform maps skew-symmetric matrices to orthogonal matrices and its extended version parameterizes the Stiefel manifold from non-square matrices, which can be useful in designing NNs [11], [17], [21].

Lemma 3 (Cayley transform [22]): For all $Y \in \mathbb{R}^{n \times n}$ and $Z \in \mathbb{R}^{m \times n}$ the Cayley transform

$$\text{Cayley} \left(\begin{bmatrix} Y \\ Z \end{bmatrix} \right) = \begin{bmatrix} U \\ V \end{bmatrix} = \begin{bmatrix} (I+M)^{-1}(I-M) \\ 2Z(I+M)^{-1} \end{bmatrix},$$

where $M = Y - Y^\top + Z^\top Z$, yields matrices $U \in \mathbb{R}^{n \times n}$ and $V \in \mathbb{R}^{m \times n}$ that satisfy $U^\top U + V^\top V = I$.

Note that $I + M$ is nonsingular since $1 \leq \lambda_{\min}(I + Z^\top Z) \leq \text{Re}(\lambda_{\min}(I + M))$.

IV. DIRECT PARAMETERIZATION

While (10) analyzes Lipschitz continuity for given 1D CNNs, it also is desirable to train robust CNNs, i. e., ρ -Lipschitz bounded CNNs where the robustness level ρ is chosen by the user. In this section, we introduce a layer-wise parameterization for 1D CNNs (1) that renders the input-output mapping Lipschitz continuous. We first discuss a parameterization for fully connected layers that satisfy (9) by design, using a similar construction to [11]. Our key contribution then is the parameterization of convolutional layers, which is carried out in two steps. In a first step, we establish a parameterization of P_i that renders the left upper block in (8) positive definite using the controllability Gramian and afterwards, we introduce the parameterization for convolutional layers that by design satisfy (8).

A. Fully connected layers

In the following, we present a mapping $\kappa_i \mapsto (W_i, b_i)$ from unconstrained variables κ_i that renders (9) feasible by design.

Theorem 4: Fully connected layers (3) parameterized by

$$W_i = \sqrt{2} \Gamma_i^{-1} V_i^\top L_{i-1}, \quad b_i \in \mathbb{R}^{n_i}, \quad \forall i \in \mathcal{I}_F \setminus \{l\}, \quad (11a)$$

$$W_l = V_l^\top L_{l-1}, \quad b_l \in \mathbb{R}^{n_l}, \quad (11b)$$

wherein

$$\Gamma_i = \text{diag}(\gamma_i), \quad L_i = \sqrt{2} U_i \Gamma_i, \quad \begin{bmatrix} U_i \\ V_i \end{bmatrix} = \text{Cayley} \left(\begin{bmatrix} Y_i \\ Z_i \end{bmatrix} \right),$$

satisfy (9). This yields the mappings $(Y_i, Z_i, \gamma_i, b_i) \mapsto (W_i, b_i)$, $i \in \mathcal{S}_F \setminus \{l\}$, and $(Y_l, Z_l, b_l) \mapsto (W_l, b_l)$, respectively, where $Y_i \in \mathbb{R}^{n_i \times n_i}$, $Z_i \in \mathbb{R}^{n_{i-1} \times n_i}$, $b_i \in \mathbb{R}^{n_i}$, $i \in \mathcal{S}_F$, $\gamma_i \in \mathbb{R}^{n_i}$, $i \in \mathcal{S}_F \setminus \{l\}$ are free variables.

Proof: According to Lemma 3, U_i and V_i satisfy $U_i^\top U_i + V_i^\top V_i = I$, wherein we insert the parameterization (11a) and $U_i = \frac{1}{\sqrt{2}} L_i \Gamma_i^{-1}$ to obtain

$$\frac{1}{2}(\Gamma_i^{-1} L_i^\top L_i \Gamma_i^{-1} + \Gamma_i W_i^\top L_{i-1}^{-1} L_{i-1}^{-\top} W_i \Gamma_i) = I.$$

With $Q_i = L_i^\top L_i$ and $\Lambda_i = \Gamma_i^\top \Gamma_i$, we further obtain

$$Q_i + \Lambda_i W_i Q_{i-1}^{-1} W_i^\top \Lambda_i = 2\Lambda_i,$$

which implies $2\Lambda_i - Q_i - \Lambda_i W_i Q_{i-1}^{-1} W_i^\top \Lambda_i \succeq 0$. Next, we apply the Schur complement, which yields the left inequality in (9). The last fully connected layer is a special case that does not contain an activation function. Inserting the parameterization (11b) gives

$$U_l^\top U_l + V_l^\top V_l = U_l^\top U_l + W_l Q_{l-1}^{-1} W_l^\top = I,$$

which implies $I - W_l Q_{l-1}^{-1} W_l^\top = U_l^\top U_l \succeq 0$, which by the Schur complement satisfies the right inequality in (9). ■

Note that the connection between the auxiliary matrices L_i in (11) and Q_i in (9) is $L_i^\top L_i = Q_i$ and the relation between the multiplier matrices Λ_i in (7) / (9) and Γ_i in (11) is $\Gamma_i^\top \Gamma_i = \Lambda_i$.

Remark 5: Throughout the paper, we assume that Γ_i and L_i are nonsingular. In our experiments, this was always the case. However, there also are tricks to enforce this property, e.g., by choosing $\Gamma_i = \text{diag}(e^{\gamma_i})$ [11].

Remark 6: Our parameterization (11) is equivalent to the one established in [11], where they show that it is necessary and sufficient, i.e., the fully connected layers (3) satisfy (9) if and only if the weights can be parameterized by (11).

B. Parameterization by the controllability Gramian

In this section, we make use of the controllability Gramian of (5) to parameterize convolutional layers, which to the best knowledge of the authors, has thus far not appeared in the context of parameterizing NNs. For that purpose, we introduce

$$F_i := \begin{bmatrix} P_i - A_i^\top P_i A_i & -A_i^\top P_i B_i \\ -B_i^\top P_i A_i & Q_{i-1} - B_i^\top P_i B_i \end{bmatrix} \succ 0, \quad (12)$$

which is the left upper block in (8) and further, we introduce $\widehat{C}_i := [C_i \ D_i]$, which simplifies the notation of (8) to

$$\begin{bmatrix} F_i & -\widehat{C}_i^\top \Lambda_i \\ -\Lambda_i \widehat{C}_i & 2\Lambda_i - Q_i \end{bmatrix} \succeq 0. \quad (13)$$

We note that the LMI (13) and the left LMI in (9) share a similar structure. The right lower block is the same in both LMIs. Beside the biases b_i , the parameters to be trained in the CNN layers are collected in \widehat{C}_i , cmp. (6), whereas the parameters W_i characterize the fully connected layers. In the off-diagonal blocks of the respective LMIs (13) and (9) \widehat{C}_i and W_i appear respectively multiplied by Λ_i . The only difference appears in the left upper blocks of the LMIs. While in LMI (9) for fully connected layers, we here have

$Q_{i-1} = L_{i-1}^\top L_{i-1} \succ 0$ for nonsingular L_{i-1} , LMI (13) for convolutional layers here contains F_i , that depends on Q_{i-1} . To render F_i positive definite, we parameterize P_i as follows, using the controllability Gramian.

Lemma 7: For some $\varepsilon > 0$ and all $H_i \in \mathbb{R}^{n_{x_i} \times n_{x_i}}$, the matrix $P_i = X_i^{-1}$ with

$$X_i = \sum_{k=0}^{n_{x_i} - c_{i-1}} A_i^k (B_i Q_{i-1}^{-1} B_i^\top + H_i^\top H_i + \varepsilon I) (A_i^\top)^k, \quad (14)$$

renders (12) feasible.

Proof: The matrix A_i is a nilpotent matrix, i.e., $A_i^{n_{x_i} - c_{i-1} + k} = 0 \ \forall k \geq 1$, such that (14) can be written as

$$X_i = \sum_{k=0}^{\infty} A_i^k (B_i Q_{i-1}^{-1} B_i^\top + H_i^\top H_i + \varepsilon I) (A_i^\top)^k,$$

which corresponds to the controllability Gramian of the linear time-invariant system characterized by (A_i, B_i) as defined in (6a), i.e., the unique solution $X_i \succ 0$ to the Lyapunov equation

$$X_i - A_i X_i A_i^\top - B_i Q_{i-1}^{-1} B_i^\top = H_i^\top H_i + \varepsilon I \succ 0. \quad (15)$$

Note that X_i is positive definite by design, given that $Q_{i-1} = L_{i-1}^\top L_{i-1} \succ 0$ such that $B_i Q_{i-1}^{-1} B_i^\top + H_i^\top H_i + \varepsilon I$ is positive definite. Next, we apply the Schur complement to (15) to obtain

$$\begin{bmatrix} X_i^{-1} & 0 & A_i^\top \\ 0 & Q_{i-1} & B_i^\top \\ A_i & B_i & X_i \end{bmatrix} \succ 0.$$

Now inserting $P_i = X_i^{-1}$ and again applying the Schur complement yields (12). ■

C. Convolutional layers

In this subsection, we present a direct parameterization for convolutional layers such that they satisfy (13) by design. Our parameterization of convolution kernels $K^i \in \mathbb{R}^{c_i \times c_{i-1} \times \ell_i}$, or equivalently $\widehat{C}_i \in \mathbb{R}^{c_i \times \ell_i c_{i-1}}$, is independent of the input dimension N_i whereas other approaches design a parameterization for Lipschitz-bounded convolutions and CNNs in the Fourier domain which involves the inversion of N_i matrices [17], [11].

Theorem 8: Convolutional layers (2) that contain an average pooling layer or no pooling layer parameterized by

$$\widehat{C}_i = \sqrt{2} \Gamma_i^{-1} V_i^\top L_i^F, \quad b_i \in \mathbb{R}^{c_i}, \quad \forall i \in \mathcal{S}_C \setminus \mathcal{S}_C^{\max}, \quad (16)$$

wherein

$$\Gamma_i = \text{diag}(\gamma_i), \quad \begin{bmatrix} U_i \\ V_i \end{bmatrix} = \text{Cayley} \left(\begin{bmatrix} Y_i \\ Z_i \end{bmatrix} \right), \quad L_i^F = \text{chol}(F_i),$$

satisfy (8). Here, $\text{chol}(\cdot)$ denotes the Cholesky decomposition, $Q_i = L_i^\top L_i$, $L_0 = \rho I$, $L_i = \sqrt{2} U_i \Gamma_i$, F_i is given by (12) with P_i parameterized from Q_{i-1} and H_i using (14). The free variables beside b_i are $Y_i \in \mathbb{R}^{c_i \times c_i}$, $Z_i \in \mathbb{R}^{\ell_i c_{i-1} \times c_i}$, $H_i \in \mathbb{R}^{n_{x_i} \times n_{x_i}}$, $\gamma_i \in \mathbb{R}^{c_i}$, $i \in \mathcal{S}_C \setminus \mathcal{S}_C^{\max}$, which yields the mapping $(Y_i, Z_i, H_i, \gamma_i, b_i) \mapsto (K^i, b_i)$.

Proof: The matrices U_i and V_i satisfy $U_i^\top U_i + V_i^\top V_i = I$, wherein we insert the parameterization (16) to obtain

$$\frac{1}{2}(\Gamma_i^{-1} L_i^\top L_i \Gamma_i^{-1} + \Gamma_i W_i^\top L_{i-1}^{F^{-1}} L_{i-1}^{F^{-1} \top} W_i \Gamma_i) = I.$$

Lemma 7 ensures positive definiteness of F_i , i.e., its Cholesky decomposition exists, and we insert $Q_i = L_i^\top L_i$, $F_i = L_i^{F^{-1}} L_i^F$, $\Lambda_i = \Gamma_i^\top \Gamma_i$, to further obtain

$$Q_i + \Lambda_i \widehat{C}_i F_i^{-1} \widehat{C}_i^\top \Lambda_i = 2\Lambda_i,$$

which implies $2\Lambda_i - Q_i - \Lambda_i \widehat{C}_i F_i^{-1} \widehat{C}_i^\top \Lambda_i \succeq 0$. Next, we apply the Schur complement and obtain (13) which corresponds to (8). ■

To account for average pooling layers present in the CNN, we rescale the desired Lipschitz bound with the product of the Lipschitz bounds of the average pooling layers, i.e., $\tilde{\rho} = \rho / \prod_{s \in \mathcal{S}_p^{\text{av}}} \text{cmp}$. Theorem 2, and then we use (16) to parameterize the convolutional layer.

D. Maximum pooling layers

In case there is a maximum pooling layer in \mathcal{C}_i , i.e. $\mathcal{C}_i = \mathcal{P}_i^{\text{max}} \circ \mathcal{C}_i$, we need to parameterize Q_i as a diagonal matrix, cmp. Theorem 2, which also affects the parameterization. This constraint comes from the incremental quadratic constraint used to include maximum pooling layers [3]. While average pooling is a linear operation and allows for full matrices Q_i , the nonlinearity of the maximum pooling layer is the reason for the additional diagonality constraint.

Corollary 9: Convolutional layers (3) that contain a maximum pooling layer parameterized by

$$\widehat{C}_i = \Lambda_i^{-1} \tilde{\Gamma}_i \tilde{U}_i^\top L_i^F, \quad b_i \in \mathbb{R}^{c_i}, \quad \forall \mathcal{C}_i^{\text{max}} \quad (17)$$

satisfy (8), wherein

$$\Lambda_i = \frac{1}{2} \left(\tilde{\Gamma}_i^\top \tilde{\Gamma}_i + Q_i \right), \quad \tilde{\Gamma}_i = \text{diag}(\tilde{\gamma}_i), \quad \tilde{U}_i = \text{Cayley}(\tilde{Y}_i),$$

$Q_i = L_i^\top L_i$, $L_0 = \rho I$, $L_i = \text{diag}(l_i)$, $L_i^F = \text{chol}(F_i)$, where F_i is given by (12) with P_i parameterized from Q_{i-1} and H_i using (14). The free variables $\tilde{Y}_i \in \mathbb{R}^{l_i c_{i-1} \times c_i}$, $H_i \in \mathbb{R}^{n_{x_i} \times n_{x_i}}$, $\tilde{\gamma}_i, l_i \in \mathbb{R}^{c_i}$, $i \in \mathcal{S}_C^{\text{max}}$ compose the mapping $(\tilde{Y}_i, H_i, \tilde{\gamma}_i, l_i, b_i) \mapsto (K^i, b_i)$.

Proof: Using that \tilde{U}_i satisfies $\tilde{U}_i^\top \tilde{U}_i = I$, we insert (17) and we replace $F_i = L_i^{F^{-1}} L_i^F$, whose Cholesky decomposition exists by Lemma 7, to obtain

$$\Lambda_i \widehat{C}_i F_i^{-1} \widehat{C}_i^\top \Lambda_i = \tilde{\Gamma}_i^\top \tilde{\Gamma}_i = 2\Lambda_i - Q_i,$$

which implies $2\Lambda_i - Q_i - \Lambda_i \widehat{C}_i F_i^{-1} \widehat{C}_i^\top \Lambda_i \succeq 0$, which by the application of the Schur complement satisfies (13). ■

E. ρ -Lipschitz layers

In this work, we enforce Lipschitz continuity onto the input-output mapping of the CNN, which is more general and less conservative than approaches that include Lipschitz guarantees for the individual layers, using that the product of the Lipschitz bounds gives the Lipschitz bound for the NN [18]. In this section, we adapt our approach to the special




Fig. 1: Different heart wave classes.

case of parameterizing ρ -Lipschitz linear layers, i.e., by choosing $Q_{i-1} = \rho^2 I$, $Q_i = I$, and $\Lambda_i = I$.

Corollary 10: The i -th linear fully connected layer

$$v^i = W_i w^{i-1} + b_i \quad \text{with} \quad W_i = \rho \tilde{U}_i^\top, \quad b_i \in \mathbb{R}^{n_i},$$

is ρ -Lipschitz continuous, where $\tilde{U}_i = \text{Cayley}(\tilde{Y}_i)$, and $\tilde{Y}_i \in \mathbb{R}^{n_{i-1} \times n_i}$ is a free variable, which yields the mapping $(\tilde{Y}_i, b_i) \mapsto (W_i, b_i)$.

Corollary 11: The i -th 1D convolution

$$v_i = b_i + \sum_{j=0}^{\ell_i-1} K_j^i w_{k-j}^{i-1} \quad \text{with} \quad \widehat{C}_i = \tilde{U}_i^\top L_i^F, \quad b_i \in \mathbb{R}^{c_i},$$

where K^i is recovered according to (6b), is ρ -Lipschitz continuous. Herein, $\tilde{U}_i = \text{Cayley}(\tilde{Y}_i)$, $L_i^F = \text{chol}(F_i)$, where F_i is given by (12) with P_i parameterized from $Q_{i-1} = \rho^2 I$ and H_i using (14). Beside b_i the free variables are $\tilde{Y}_i \in \mathbb{R}^{l_i c_{i-1} \times c_i}$, $H_i \in \mathbb{R}^{n_{x_i} \times n_{x_i}}$, yielding the mapping $(\tilde{Y}_i, H_i, b_i) \mapsto (K^i, b_i)$.


The proofs of Corollaries 10 and 11 follow along the lines of the proof of Corollary 9.

V. SIMULATION RESULTS


In this section, we train Lipschitz-bounded 1D CNNs (LipCNNs) to classify heart arrhythmia data from the MIT-BIH database, an ECG database [19], in its preprocessed form [23] that assigns heart wave signals to one of the five classes depicted in Fig. 1: N (normal beat), L (left bundle branch block), R (right bundle branch block), A (atrial premature contraction), V (ventricular premature contraction). In particular, we use 26,490 samples, 13,245 training data points and 13,245 test data points. Fig. 2a shows the architecture of the 1D CNN used to train on the ECG dataset, and for training we choose the cross-entropy loss as the objective. For training, we use Flux, JuMP, and RobustNeuralNetworks.jl [24] in Julia in combination with MOSEK [25] on a standard i7 notebook¹.

We compare our approach to vanilla CNNs and L2 regularized CNNs with different weighting factors γ . To evaluate the robustness of the CNNs, we compute the test accuracy on adversarial examples from the L2 projected gradient descent attack for different perturbation strengths, which is shown in Fig. 2b. Note that the two shown LipCNNs indicate comparably high test accuracies to the shown vanilla and L2 regularized CNNs on unperturbed data while maintaining higher accuracies than their counterparts as the perturbation strength increases. Further, Table I shows the averaged test accuracies and upper and lower Lipschitz bounds for different 1D CNNs and Fig. 3 illustrates the tradeoff between

¹Code is available at www.github.com/ppauli/1D-LipCNNs.



(a) 1D CNN architecture.



(b) Adversarial attack.

Fig. 2: (a) CNN used for ECG dataset. (b) Vanilla, L2 regularized ($\gamma = 0.1$) and Lipschitz ($\rho = 10, 50$) CNN on adversarial examples with different perturbation strengths ϵ created using the L2 projected gradient descent attack [26].

TABLE I: Test accuracy and Lipschitz upper bound [3] and empirical Lipschitz lower bound of trained CNNs, averaged over 5 CNNs.

	vanilla	L2 $_{\gamma=0.05}$	L2 $_{\gamma=0.1}$	Lip5	Lip10	Lip50
Test acc.	94.9%	92.4%	90.6%	84.8%	90.0%	94.7%
Lip. UB	147	45.7	33.9	4.99	9.85	45.3
Emp. LB	28.3	13.2	10.2	2.20	4.06	15.1

accuracy and robustness. Low Lipschitz bounds mean high robustness, yet we observe lower test accuracies and vice versa. We note that the vanilla CNN has significantly larger upper Lipschitz bounds than LipCNN and further, LipCNN maintains high test accuracies as the Lipschitz bound decreases in Fig. 3. With even larger weighting parameters γ in L2 regularized training the training failed, whereas LipCNNs allows for training with very low Lipschitz bounds.

VI. CONCLUSION

In this paper, we introduced a parameterization for Lipschitz-bounded 1D CNNs using Cayley transforms and controllability Gramians. Using our parameterization we can train Lipschitz-bounded 1D CNNs in an unconstrained training problem which we illustrated in the classification of ECG data from the MIT-BIH database. Future research includes the extension of our parameterization to 2D CNNs using a 2D systems approach as suggested in [15].




Fig. 3: Accuracy robustness tradeoff. Test accuracy over Lipschitz constant from SDP [3], averaged over 5 CNNs.

REFERENCES

- [1] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv:1312.6199*, 2013.
- [2] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of Lipschitz constants for deep neural networks," *Adv. Neural Inf. Process.*, vol. 32, 2019.
- [3] P. Pauli, D. Gramlich, and F. Allgöwer, "Lipschitz constant estimation for 1D convolutional neural networks," in *Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 1321–1332.
- [4] P. L. Combettes and J.-C. Pesquet, "Deep neural network structures solving variational inequalities," *Set-Valued and Variational Analysis*, vol. 28, no. 3, pp. 491–518, 2020.
- [5] F. Latorre, P. Rolland, and V. Cevher, "Lipschitz constant estimation of neural networks via sparse polynomial optimization," *arXiv:2004.08688*, 2020.
- [6] P. Pauli, A. Koch, J. Berberich, P. Kohler, and F. Allgöwer, "Training robust neural networks using Lipschitz bounds," *IEEE Control Systems Letters*, vol. 6, pp. 121–126, 2022.
- [7] P. Pauli, N. Funcke, D. Gramlich, M. A. Msalmi, and F. Allgöwer, "Neural network training under semidefinite constraints," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2022, pp. 2731–2736.
- [8] H. Gouk, E. Frank, B. Pfahringer, and M. J. Cree, "Regularisation of neural networks by enforcing Lipschitz continuity," *Machine Learning*, vol. 110, no. 2, pp. 393–416, 2021.
- [9] M. Revay, R. Wang, and I. R. Manchester, "Lipschitz bounded equilibrium networks," *arXiv:2010.01732*, 2020.
- [10] —, "Recurrent equilibrium networks: Flexible dynamic models with guaranteed stability and robustness," *IEEE Trans. Automat. Contr.*, 2023.
- [11] R. Wang and I. Manchester, "Direct parameterization of lipschitz-bounded deep networks," in *International Conference on Machine Learning*. PMLR, 2023, pp. 36093–36110.
- [12] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai *et al.*, "Recent advances in convolutional neural networks," *Pattern recognition*, vol. 77, pp. 354–377, 2018.
- [13] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, "1D convolutional neural networks and applications: A survey," *Mechanical systems and signal processing*, vol. 151, p. 107398, 2021.
- [14] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv:1609.03499*, 2016.
- [15] D. Gramlich, P. Pauli, C. W. Scherer, F. Allgöwer, and C. Ebenbauer, "Convolutional neural networks as 2-D systems," *arXiv:2303.03042*, 2023.
- [16] C. Anil, J. Lucas, and R. Grosse, "Sorting out lipschitz function approximation," in *International Conference on Machine Learning*. PMLR, 2019, pp. 291–301.
- [17] A. Trockman and J. Z. Kolter, "Orthogonalizing convolutional layers with the Cayley transform," *arXiv:2104.07167*, 2021.
- [18] A. Araujo, A. J. Havens, B. Delattre, A. Allauzen, and B. Hu, "A unified algebraic perspective on Lipschitz neural networks," in *11th Int. Conf. Learn. Represent.*, 2023.
- [19] <http://www.physionet.org/physiobank/database/mitdb/>.
- [20] L. Béthune, T. Boissin, M. Serrurier, F. Mamalet, C. Friedrich, and A. Gonzalez Sanz, "Pay attention to your loss: understanding misconceptions about lipschitz neural networks," *Adv. Neural Inf. Process.*, vol. 35, pp. 20077–20091, 2022.
- [21] K. Helfrich, D. Willmott, and Q. Ye, "Orthogonal recurrent neural networks with scaled Cayley transform," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1969–1978.
- [22] M. Jauch, P. D. Hoff, and D. B. Dunson, "Random orthogonal matrices and the Cayley transform," *Bernoulli*, vol. 26, no. 2, pp. 1560–1586, 2020.
- [23] S. Abuadba, "split-learning-1d," <https://github.com/SharifAbuadba/split-learning-1D>, 2020.
- [24] N. H. Barbara, M. Revay, R. Wang, J. Cheng, and I. R. Manchester, "Robustneuralnetworks.jl: a package for machine learning and data-driven control with certified robustness," *arXiv:2306.12612*, 2023.
- [25] M. ApS, *The MOSEK optimization toolbox. Version 10.0.*, 2022. [Online]. Available: <http://docs.mosek.com/10.0/toolbox/index.html>
- [26] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv:1706.06083*, 2017.