
Robust Implicit Networks via Non-Euclidean Contractions

Saber Jafarpour^{1,*}, Alexander Davydov^{1,*}, Anton V. Proskurnikov^{2,3}, and Francesco Bullo¹

¹ Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara, 93106-5070, USA, {saber, davydov, bullo}@ucsb.edu.

² Department of Electronics and Telecommunications, Politecnico di Torino, Turin, Italy;

³ Institute for Problems in Mechanical Engineering, Russian Academy of Sciences, St. Petersburg, Russia, anton.p.1982@ieee.org

Abstract

Implicit neural networks, a.k.a., deep equilibrium networks, are a class of implicit-depth learning models where function evaluation is performed by solving a fixed point equation. They generalize classic feedforward models and are equivalent to infinite-depth weight-tied feedforward networks. While implicit models show improved accuracy and significant reduction in memory consumption, they can suffer from ill-posedness and convergence instability.

This paper provides a new framework, which we call Non-Euclidean Monotone Operator Network (NEMON), to design well-posed and robust implicit neural networks based upon contraction theory for the non-Euclidean norm ℓ_∞ . Our framework includes (i) a novel condition for well-posedness based on one-sided Lipschitz constants, (ii) an average iteration for computing fixed-points, and (iii) explicit estimates on input-output Lipschitz constants. Additionally, we design a training problem with the well-posedness condition and the average iteration as constraints and, to achieve robust models, with the input-output Lipschitz constant as a regularizer. Our ℓ_∞ well-posedness condition leads to a larger polytopic training search space than existing conditions and our average iteration enjoys accelerated convergence. Finally, we evaluate our framework in image classification through the MNIST and the CIFAR-10 datasets. Our numerical results demonstrate improved accuracy and robustness of the implicit models with smaller input-output Lipschitz bounds. Code is available at https://github.com/davydovalexander/Non-Euclidean_Mon_Op_Net.

1 Introduction

Implicit neural networks are infinite-depth learning models with layers defined implicitly through a fixed-point equation. Examples of implicit neural networks include deep equilibrium models [Bai et al., 2019] and implicit deep learning models [El Ghaoui et al., 2021]. Implicit networks can be considered as generalizations of feedforward neural networks with input-injected weight tying, i.e., training parameters are transferable between layers. Indeed, in implicit networks, function evaluation is executed by solving a fixed-point equation and backpropagation is implemented by computing gradients using implicit differentiation. Due to these unique features, implicit models enjoy more flexibility and improved memory efficiency compared to traditional neural networks. At the same time, implicit networks can suffer from instability in their training due to the nonlinear nature of their fixed-point equations and can show brittle input-output behaviors due to their model flexibility.

*These authors contributed equally.

It is known that implicit neural networks require careful tuning and initialization to avoid ill-posed training procedures. Indeed, without additional assumptions, their fixed-point equation may not have a unique solution and the numerical algorithms for finding their solutions might not converge. Several recent works in the literature have focused on studying well-posedness and convergence of the fixed-point equations of implicit networks using frameworks such as monotone operator theory [Winston and Kolter, 2020], contraction theory [El Ghaoui et al., 2021], and a mixture of both [Revay et al., 2020]. Despite several insightful results, important questions about conditions for well-posedness of implicit networks and efficient algorithms that converge to their solutions are still open.

One of the key features of implicit neural networks is their flexibility, which might come at the cost of low input-output robustness. As first noted in [Szegedy et al., 2014], the input-output behavior of deep neural networks can be vulnerable to perturbations; close enough input data can lead to completely different outputs. This lack of robustness can lead to unreliable performance of neural networks in safety-critical applications. Among several notions of robustness, the Lipschitz constant of a neural network is a coarse but rigorous measure which can be used to estimate input-output sensitivity of the network [Szegedy et al., 2014]. For this reason, there has been a growing interest in estimating the input-output Lipschitz constant of deep neural networks with respect to the ℓ_2 -norm [Fazlyab et al., 2019, Combettes and Pesquet, 2020]. However, it turns out that in some applications, the input-output Lipschitz constants with respect to non-Euclidean norms are more informative measures for studying robustness. One such application appears in the robustness analysis of neural networks with large-scale inputs under widely-distributed adversarial perturbations (examples of these adversarial perturbations can be found in [Szegedy et al., 2014]). For these examples, the input-output ℓ_2 -Lipschitz constant does not provide complete information about robustness of the network; a neural network with small input-output ℓ_2 -Lipschitz constant can be very sensitive to widespread entrywise-small perturbations of the input signal. On the other hand, the input-output ℓ_∞ -Lipschitz constant provides a different metric which appears to be well-suited for the analysis of widespread distributed perturbations. Another application is the estimation of input signal confidence intervals from output deviations, where the input-output ℓ_∞ -Lipschitz constant of the network provides more scalable bounds than its ℓ_2 counterpart.

Related works

Implicit learning models. Numerous works in learning theory have shown the power of deep learning models with implicit layers. In these learning models, the notion of layers are replaced by a composition rule, which can be either a fixed-point iteration or a solution to a differential equation. Well-known frameworks for deep learning using implicit infinite-depth layers include deep equilibrium networks [Bai et al., 2019], implicit deep learning [El Ghaoui et al., 2021], and Neural ODEs [Chen et al., 2018]. In [Kag et al., 2020], a class of implicit recurrent neural networks is considered and it is demonstrated that, with this architecture, the models do not suffer from vanishing nor exploding gradients. Implicit layers have also been used to study convex optimization problems [Agrawal et al., 2019] and to design control strategies [Amos et al., 2018]. Convergence to global minima of certain classes of implicit networks is studied in [Kawaguchi, 2021].

Well-posedness and numerical algorithms for fixed-point equations. There has been a recent interest in studying well-posedness and numerical stability of implicit-depth learning models. [El Ghaoui et al., 2021] proposes a sufficient spectral condition for well-posedness and for convergence of the Picard iterations associated with the fixed-point equation of implicit networks. In [Winston and Kolter, 2020, Revay et al., 2020], using monotone operator theory, a suitable parametrization of the weight matrix is proposed which guarantees the stable convergence of suitable fixed-point iterations. A recent influential survey on monotone operators is [Ryu and Boyd, 2016]. A recent survey on fixed point strategies in data science is given by [Combettes and Pesquet, 2021].

Robustness of learning models It is known that neural networks can be vulnerable to adversarial input perturbations [Szegedy et al., 2014]. A large body of literature is devoted to improve robustness of neural networks using various defense strategies against adversarial examples [Goodfellow et al., 2015, Papernot et al., 2016]. While these strategies are effective in many scenarios, they do not provide formal guarantees for robustness [Carlini and Wagner, 2017]. However, there has been a recent interest in designing classifiers that are provably robust with respect to adversarial perturbations [Madry et al., 2018, Wong and Kolter, 2018]. The input-output Lipschitz constant of a neural

network is a rigorous metric for its worst-case sensitivity with respect to input perturbations. Several recent works have focused on estimating the Lipschitz constant and enforcing its boundedness. For example, [Fazlyab et al., 2019, 2020] propose a convex optimization framework using quadratic constraints and semidefinite programming to obtain upper bounds on Lipschitz constants of deep neural networks. In [Pauli et al., 2021], a training algorithm is designed to ensure boundedness of the Lipschitz constant of the neural network via a semidefinite program. Other methods for estimating the Lipschitz constant of deep neural networks include [Krishnan et al., 2020, Revay et al., 2021, Combettes and Pesquet, 2020]. For implicit neural networks, a sensitivity-based robustness analysis is proposed in [El Ghaoui et al., 2021]. Lipschitz constants of deep equilibrium networks have also been studied in [Pabbaraju et al., 2021, Revay et al., 2020] using monotone operator theory.

Contributions

In this paper, using non-Euclidean contraction theory with respect to the ℓ_∞ -norm, we propose our novel framework, Non-Euclidean Monotone Operator Network (NEMON), to design implicit neural networks and study their well-posedness, stability, and robustness. First, we develop elements of a novel non-Euclidean monotone operator theory akin to the frameworks in [Bauschke and Combettes, 2017, Ryu and Boyd, 2016]. Using the concept of matrix measure, we introduce the essential notion of one-sided Lipschitz constant of a map. Based upon this notion, we prove a general fixed-point theorem with weaker requirements than classical results on Picard and Krasnosel'skii–Mann iterations. For maps with one-sided Lipschitz constant less than unity, we show that an average iteration converges for sufficiently small step sizes and optimize its rate of convergence. For the special case of the weighted ℓ_∞ -norm, we show that this average iteration can be accelerated by choosing a larger step size. Additionally, we study perturbed fixed-point equations and establish a bound on the distance between perturbed and nominal equilibrium points as a function of one-sided Lipschitz condition. Second, for implicit neural networks, we use our new fixed-point theorem to (i) establish ℓ_∞ -norm conditions for their well-posedness, (ii) design accelerated numerical algorithms for computing their solutions, and (iii) provide upper bounds on their input-output ℓ_∞ -Lipschitz constants. Third, we propose a parametrization for matrices with appropriate bound on their one-sided Lipschitz constants and use this parametrization with the average iteration to design a training optimization problem. Finally, we perform several numerical experiments illustrating improved performance of NEMON in image classification on the MNIST and the CIFAR-10 datasets compared to the state-of-the-art models in [El Ghaoui et al., 2021, Winston and Kolter, 2020]. Additionally, by adding the input-output Lipschitz constant as regularizer in the training problem, we observe improved robustness to some classes of adversarial perturbations. We include all relevant proofs in Appendix C.

2 Review material

Matrix measures Let $\|\cdot\|$ be a norm on \mathbb{R}^n and its induced norm on $\mathbb{R}^{n \times n}$. The matrix measure of $A \in \mathbb{R}^{n \times n}$ with respect to $\|\cdot\|$ is defined by $\mu(A) := \lim_{h \rightarrow 0^+} \frac{\|I_n + hA\| - 1}{h}$, that is, the one-sided directional derivative of the induced norm in direction of A , evaluated at I_n . Remarkably, the matrix measure is a tighter upper bound on the spectral abscissa of A than $\|A\|$ and the set of matrices $A \in \mathbb{R}^{n \times n}$ satisfying $\mu(A) \leq 1$ is an unbounded subset of $\mathbb{R}^{n \times n}$ strictly containing the compact ball $\|A\| \leq 1$. We refer to [Desoer and Haneda, 1972] for a list of properties enjoyed by matrix measures.

We will be specifically interested in diagonally weighted ℓ_∞ norms defined by

$$\|x\|_{\infty, [\eta]^{-1}} = \max_i \frac{1}{\eta_i} |x_i|, \quad (1)$$

where, given a positive vector $\eta \in \mathbb{R}_{>0}^n$, we use $[\eta]$ to denote the diagonal matrix with diagonal entries η . The corresponding matrix norm and measure are

$$\|A\|_{\infty, [\eta]^{-1}} = \max_{i \in \{1, \dots, n\}} \sum_{j=1}^n \frac{\eta_j}{\eta_i} |a_{ij}|, \quad \mu_{\infty, [\eta]^{-1}}(A) = \max_{i \in \{1, \dots, n\}} \left(a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| \frac{\eta_j}{\eta_i} \right). \quad (2)$$

Lipschitz maps Given a norm $\|\cdot\|$ with induced matrix measure $\mu(\cdot)$, a differentiable map $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous with constant $\text{Lip}(F) \in \mathbb{R}_{\geq 0}$ if

$$\|DF(x)\| \leq \text{Lip}(F) \quad \text{for all } x \in \mathbb{R}^n. \quad (3)$$

For example, for an affine $F(x) = Ax + b$, the (smallest) Lipschitz constant is $\text{Lip}(F) = \|A\|$.

One-sided Lipschitz maps Given a norm $\|\cdot\|$, a differentiable map $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is one-sided Lipschitz continuous with constant $\text{osL}(F) \in \mathbb{R}$ if

$$\mu(DF(x)) \leq \text{osL}(F) \quad \text{for all } x \in \mathbb{R}^n. \quad (4)$$

For example, for an affine $F(x) = Ax + b$, the (smallest) one-sided Lipschitz constant is $\text{osL}(F) = \mu(A)$. Note that (i) the one-sided Lipschitz constant is upper bounded by the Lipschitz constant, (ii) a Lipschitz continuous map is always one-sided Lipschitz continuous, and (iii) the one-sided Lipschitz constant may be negative. For a more in-depth review we refer to Appendix A. The notion of one-sided Lipschitz continuity unifies several important concepts in dynamical systems and optimization theory. In operator theory, the map F is called a monotone operator if it is one-sided Lipschitz continuous with respect to the ℓ_2 -norm with the constant $-\text{osL}(-F) > 0$ [Ryu and Boyd, 2016, Bauschke and Combettes, 2017]. In control theory, the vector field F is called strongly infinitesimally contracting if it is one-sided Lipschitz continuous with the constant $\text{osL}(F) < 0$ [Desoer and Haneda, 1972, Lohmiller and Slotine, 1998, Pavlov et al., 2004]. In what follows, we let $\text{osL}_{\infty, [\eta]^{-1}}(F) \in \mathbb{R}$ denote the one-sided Lipschitz constants with respect to the weighted ℓ_∞ -norm.

3 Fixed-point equations and one-sided Lipschitz constants

In this section, we show that the notion of one-sided Lipschitz constant can be used to study solvability of fixed-point equation:

$$x = F(x), \quad (5)$$

where $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a differentiable map. Let $\|\cdot\|$ be a norm on \mathbb{R}^n , then in view of the Banach fixed-point theorem, a simple sufficient condition for existence of a unique solution for the fixed-point equation (5) is $\text{Lip}(F) < 1$. We note that the sufficient condition $\text{Lip}(F) < 1$ depends on the specific form of the fixed-point equation (5) and can be relaxed by a suitable rewriting of this fixed-point equation. Given an averaging parameter $\alpha \in (0, 1]$ we define the average map $F_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}^n$ by $F_\alpha := (1 - \alpha)I + \alpha F$, where I is the identity map. Using this notion, an equivalent reformulation of the fixed-point equation (5) is:

$$x = (1 - \alpha)x + \alpha F(x) = F_\alpha(x). \quad (6)$$

For $\alpha = 1$, we have $F_\alpha(x) = F(x)$ and equation (6) coincides with equation (5). For every $\alpha \in (0, 1)$, the map F_α is different from F but equations (5) and (6) are equivalent. Hence, if $\text{Lip}(F_\alpha) < 1$, then by the Banach fixed-point theorem, the fixed point equation (6) (and therefore the fixed point equation (5)) has a unique solution x^* and the sequence $\{y_k\}_{k=1}^\infty$ defined by

$$y_{k+1} = (1 - \alpha)y_k + \alpha F(y_k), \quad \text{for all } k \in \mathbb{Z}_{\geq 0} \quad (7)$$

converges geometrically to x^* with rate $\text{Lip}(F_\alpha)$. As a result of the parametrization (6), the condition $\text{Lip}(F) < 1$ for existence and uniqueness of the fixed-point can be relaxed to sufficient conditions

$$\text{Lip}(F_\alpha) < 1, \quad (8)$$

parametrized by $\alpha \in (0, 1]$. Additionally, if condition (8) is satisfied, then algorithm (7) computes the fixed point x^* . It can be shown that the condition (8) becomes less conservative as α decreases. The next theorem shows that in the limit as $\alpha \rightarrow 0^+$, condition (8) approaches the condition $\text{osL}(F) < 1$.

Theorem 1 (Fixed points via one-sided Lipschitz conditions). *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be differentiable and Lipschitz with constant $\ell > 0$ with respect to a norm $\|\cdot\|$. Define the average map $F_\alpha = (1 - \alpha)I + \alpha F$ and, for $c > 0$, the function $\gamma_{\ell, c} :]0, \frac{c}{(c+\ell+1)(\ell+1)}[\rightarrow \mathbb{R}$ by:*

$$\gamma_{\ell, c}(\alpha) := \left(1 + \alpha c - \frac{\alpha^2(\ell + 1)^2}{1 - \alpha(\ell + 1)}\right)^{-1}.$$

Then the following statements are equivalent:

- (i) $\text{osL}(F) < 1 - c$,
- (ii) $\text{Lip}(F_\alpha) = \gamma_{\ell, c}(\alpha)$, for $0 < \alpha < \frac{c}{(c+\ell+1)(\ell+1)}$.

Moreover, if the equivalent conditions (i) or (ii) hold, then, for condition number $\kappa = \frac{1 + \text{Lip}(F)}{1 - \text{osL}(F)}$,

- (iii) F has a unique fixed point x^* ;
- (iv) for $0 < \alpha < \frac{1}{\kappa(\kappa+1)}$, F_α is a contraction mapping with contraction factor $\gamma_{\ell,c}(\alpha) < 1$;
- (v) the minimum contraction factor $\gamma_{\ell,c}^* = 1 - \frac{1}{4\kappa^2} + \frac{1}{8\kappa^3} + \mathcal{O}\left(\frac{1}{\kappa^4}\right)$ and the minimizing averaging parameter α^* of F_α is

$$\alpha^* = \frac{\kappa}{1 - \text{osL}(F)} \left(1 - \frac{1}{\sqrt{1 + 1/\kappa}} \right) = \frac{1}{1 - \text{osL}(F)} \left(\frac{1}{2\kappa^2} - \frac{3}{8\kappa^3} + \mathcal{O}\left(\frac{1}{\kappa^4}\right) \right).$$

The average iteration (6) is often referred to as the Krasnosel'skii–Mann iteration or the damped iteration [Bauschke and Combettes, 2017]. Compared to [Bauschke and Combettes, 2017, Theorem 5.15], Theorem 1(iv) studies convergence of the Krasnosel'skii–Mann iteration for arbitrary norms, proposes a weaker convergence condition of the form $\text{osL}(F) < 1$ (hence, F need not be non-expansive). However, it ensures convergence for only sufficiently small $\alpha > 0$ and assumes that F is differentiable (as will be shown, however, the latter assumption can be relaxed).

3.1 Accelerated convergence for weighted ℓ_∞ norms

For diagonally weighted ℓ_∞ norms, one can strengthen Theorem 1(iv) to prove the convergence of the average iteration (6) on a larger domain of the parameter α .

Theorem 2 (Accelerated fixed point algorithm for ℓ_∞ norms). *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be differentiable and Lipschitz with respect to the weighted non-Euclidean norm $\|\cdot\|_{\infty, [\eta]}^{-1}$. Define the average map $F_\alpha = (1 - \alpha)I + \alpha F$ and pick $\text{diagL}(F) \in [-\text{Lip}(F), \text{osL}(F)]$ to satisfy*

$$\text{diagL}(F) \leq \min_{i \in \{1, \dots, n\}} \inf_{x \in \mathbb{R}^n} DF_{ii}(x). \quad (9)$$

If $\text{osL}(F) < 1$, then F has a unique fixed-point x^ and*

- (i) for $0 < \alpha \leq \frac{1}{1 - \text{diagL}(F)}$, F_α is a contraction mapping with the contraction factor $1 - \alpha(1 - \text{osL}(F)) < 1$;
 - (ii) the minimum contraction factor and minimizing averaging parameter of F_α are, respectively,
- $$\text{Lip}(F_{\alpha^*}) = 1 - \frac{1 - \text{osL}(F)}{1 - \text{diagL}(F)} = 1 - \frac{1}{\kappa_\infty}, \quad \text{for } \kappa_\infty = \frac{1 - \text{diagL}(F)}{1 - \text{osL}(F)} \leq \frac{1 + \text{Lip}(F)}{1 - \text{osL}(F)},$$
- $$\alpha^* = \frac{1}{1 - \text{diagL}(F)}.$$

Note that $\text{diagL}(F)$ is well-defined because of the Lipschitz continuity assumption. Specifically, one can show that $\text{diagL}(F)$ is the minus the minimum over $i \in \{1, \dots, n\}$ of the one-sided Lipschitz constants of the maps $x_i \mapsto -F(x_i, x_{-i})$ at $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ fixed.

It is instructive to compare the minimum contraction factor in the general Theorem 1 with the minimum contraction factor for ℓ_∞ norms in Theorem 2 and how they depend upon the corresponding condition numbers κ and κ_∞ . We note that (i) the relevant condition number diminishes $\kappa \geq \kappa_\infty$, and (ii) the minimum contraction factor $\text{Lip}(F_{\alpha^*}) = 1 - \frac{1}{4\kappa^2} + \mathcal{O}(1/\kappa^4)$ improves to $\text{Lip}(F_{\alpha^*}) = 1 - \frac{1}{\kappa_\infty}$. This acceleration justifies the title of this section.

3.2 Perturbed fixed-point problems

In this subsection, we focus on solvability of the perturbed fixed-point equation:

$$x = F(x, u), \quad (10)$$

where $F : \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}^n$ is differentiable in x . We define $F_u(x) = F(x, u)$ and $F_x(u) = F(x, u)$. Given a norm $\|\cdot\|_{\mathcal{X}}$ in \mathbb{R}^n and $\|\cdot\|_{\mathcal{U}}$ in \mathbb{R}^r , F is Lipschitz in its first argument with constant $\text{Lip}_x(F) \in \mathbb{R}_{\geq 0}$ if

$$\|F(x_1, u) - F(x_2, u)\|_{\mathcal{X}} \leq \text{Lip}_x(F) \|x_1 - x_2\|_{\mathcal{X}} \quad \text{for all } x_1, x_2 \in \mathbb{R}^n \text{ and } u \in \mathbb{R}^r,$$

and it is Lipschitz in its second argument with constant $\text{Lip}_u(\mathbf{F}) \in \mathbb{R}_{\geq 0}$ if

$$\|\mathbf{F}(x, u_1) - \mathbf{F}(x, u_2)\|_{\mathcal{X}} \leq \text{Lip}_u(\mathbf{F}) \|u_1 - u_2\|_{\mathcal{U}} \quad \text{for all } x \in \mathbb{R}^n \text{ and } u_1, u_2 \in \mathbb{R}^r,$$

and it is one-sided Lipschitz in its first argument with constant $\text{osL}_x(\mathbf{F}) \in \mathbb{R}$ if

$$\mu(D_x \mathbf{F}(x, u)) \leq \text{osL}_x(\mathbf{F}) \quad \text{for all } x_1, x_2 \in \mathbb{R}^n \text{ and } u \in \mathbb{R}^r.$$

The following result, which is in the spirit of Lim's Lemma [Lim, 1985], provides an upper bound on the distance between fixed-points of the perturbed equation (10).

Theorem 3 (Perturbed fixed-points). *Given a norm $\|\cdot\|_{\mathcal{X}}$ in \mathbb{R}^n and a norm $\|\cdot\|_{\mathcal{U}}$ in \mathbb{R}^r , consider a map $\mathbf{F} : \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}^n$ differentiable in the first argument and Lipschitz in both arguments. If \mathbf{F} is one-sided Lipschitz with constant $\text{osL}_x(\mathbf{F}) < 1$, then*

(i) *for every $u \in \mathbb{R}^m$, the map \mathbf{F}_u has a unique fixed point x_u^* ;*

(ii) *for every $u, v \in \mathbb{R}^m$, $\|x_u^* - x_v^*\|_{\mathcal{X}} \leq \frac{\text{Lip}_u(\mathbf{F})}{1 - \text{osL}_x(\mathbf{F})} \|u - v\|_{\mathcal{U}}$.*

Finally, Theorems 1, 2, and 3 are not directly applicable to activation function that are not differentiable. In Appendix C.3, we show that for specific form of the fixed-point equation (5), where $\mathbf{F} = \Phi \circ \mathbf{H}$ and $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a weakly increasing, non-expansive, diagonal activation function and $\mathbf{H} : \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}^n$ is a differentiable function, all of the conclusions of Theorems 1, 2, and 3 hold by requiring equation (9) to be true almost everywhere.

4 Contraction analysis of implicit neural networks

In this section, we use contraction theory to lay the foundation for our Non-Euclidean Monotone Operator Network (NEMON) model of implicit neural networks. Given $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times r}$, $C \in \mathbb{R}^{q \times n}$, and $D \in \mathbb{R}^{q \times r}$, we consider the implicit neural network

$$x = \Phi(Ax + Bu) := \mathbf{N}(x, u), \quad y = Cx + Du, \quad (11)$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^r$, $y \in \mathbb{R}^q$, and $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined by $\Phi(x) = (\phi_1(x_1), \dots, \phi_n(x_n))$. For every $i \in \{1, \dots, n\}$, we assume the activation function $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ is weakly increasing, i.e., $\phi_i(x_i) \geq \phi_i(z_i)$ for $x_i \geq z_i$, and non-expansive, i.e., $|\phi_i(x_i) - \phi_i(z_i)| \leq |x_i - z_i|$ for all x_i and z_i ; if ϕ_i is differentiable, these conditions are equivalent to $0 \leq \phi'_i(x_i) \leq 1$ for all $x_i \in \mathbb{R}$.

We are able to provide the following estimates on all relevant Lipschitz constants.

Theorem 4 (Lipschitz and one-sided Lipschitz constants for the implicit neural network). *Consider the implicit neural network in equation (11) with weakly increasing and non-expansive activation functions Φ . With respect to $\|\cdot\|_{\infty, [\eta]}^{-1}$, $\eta \in \mathbb{R}_{>0}^n$, on \mathbb{R}^n and $\|\cdot\|_{\mathcal{U}}$ on the input space \mathbb{R}^r , the map $\mathbf{N} : \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}^n$ is one-sided Lipschitz continuous in the first variable and Lipschitz continuous in both variables with constants:*

$$\text{osL}_x(\mathbf{N}) = \mu_{\infty, [\eta]}^{-1}(A)_+, \quad \text{Lip}_x(\mathbf{N}) = \|A\|_{\infty, [\eta]}^{-1}, \quad (12)$$

$$\text{Lip}_u(\mathbf{N}) = \|B\|_{(\infty, [\eta]^{-1}), \mathcal{U}}, \quad \text{diagL}(\mathbf{N}) = \min_{i \in \{1, \dots, n\}} (A_{ii})_-, \quad (13)$$

where $(z)_+ = z$ if $z \geq 0$ and $(z)_+ = 0$ if $z < 0$; and $(z)_- = 0$ if $z \geq 0$ and $(z)_- = z$ if $z < 0$.

We now use these estimates to establish multiple properties of the implicit neural network.

Corollary 5 (Well posedness, input-state Lipschitz constant, and computation). *Consider the model (11), with parameters (A, B, C, D) and with weakly increasing and non-expansive activation functions Φ . Define the average map $\mathbf{N}_\alpha := (1 - \alpha)\mathbf{I} + \alpha\mathbf{N}$ and consider the norms $\|\cdot\|_{\infty, [\eta]}^{-1}$, $\eta \in \mathbb{R}_{>0}^n$, on \mathbb{R}^n , $\|\cdot\|_{\mathcal{U}}$ on the input space \mathbb{R}^r and $\|\cdot\|_{\mathcal{Y}}$ on the output space \mathbb{R}^q . Then*

(i) *if $\mu_{\infty, [\eta]}^{-1}(A) < 1$, then (11) is well posed, i.e., there exists a unique fixed point,*

(ii) *the map \mathbf{N}_α is a contraction mapping for $0 < \alpha \leq \alpha^* := (1 - \min_{i \in \{1, \dots, n\}} (A_{ii})_-)^{-1}$ with minimum contraction factor $\text{Lip}(\mathbf{N}_{\alpha^*}) = 1 - \frac{1 - \mu_{\infty, [\eta]}^{-1}(A)_+}{1 - \min_{i \in \{1, \dots, n\}} (A_{ii})_-}$.*

(iii) the Lipschitz constants from input u to fixed point x_u^* and to the output $y = Cx_u^* + Du$ are

$$\text{Lip}_{u \rightarrow x^*} := \frac{\text{Lip}_u(\mathbf{N})}{1 - \text{osL}_x(\mathbf{N})} = \frac{\|B\|_{(\infty, [\eta]^{-1}), \mathcal{U}}}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+}, \quad (14)$$

$$\text{Lip}_{u \rightarrow y} := \frac{\|B\|_{(\infty, [\eta]^{-1}), \mathcal{U}} \|C\|_{\mathcal{Y}, (\infty, [\eta]^{-1})}}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+} + \|D\|_{\mathcal{Y}, \mathcal{U}}. \quad (15)$$

5 Training implicit neural networks

Problem setup Given an input data matrix $U = [u_1, \dots, u_m] \in \mathbb{R}^{r \times m}$ and a corresponding output data matrix $Y = [y_1, \dots, y_m] \in \mathbb{R}^{q \times m}$, we aim to learn matrices A, B, C, D so that the neural network (11) approximates the input-output relationship. We rewrite the model for matrix inputs as $\hat{Y} = CX + DU$, where $X = \Phi(AX + BU)$. From Corollary 5(i), if each ϕ_i is weakly increasing and non-expansive, the fixed point problem is well-posed when $\mu_{\infty, [\eta]^{-1}}(A) < 1$ for some $\eta \in \mathbb{R}_{>0}^n$. We consider a training problem of the form

$$\min_{A, B, C, D, X} \mathcal{L}(Y, CX + DU) + \mathcal{P}(A, B, C, D) \quad (16)$$

$$X = \Phi(AX + BU), \quad \mu_{\infty, [\eta]^{-1}}(A) \leq \gamma,$$

where \mathcal{L} is a loss function, \mathcal{P} is a penalty function, and $\gamma < 1$ is a hyperparameter ensuring the fixed point problem is well-posed. For $\eta = \mathbf{1}_n$, we can remove the constraint $\mu_{\infty}(A) \leq \gamma$ in the training optimization problem (16) using the following parametrization of weight matrix A :

$$A = T - \text{diag}(|T| \mathbf{1}_n) + \gamma I_n. \quad (17)$$

In Appendix B, we show that parametrization (17) characterizes the set of matrices in $\mathbb{R}^{n \times n}$ satisfying $\mu_{\infty}(A) \leq \gamma$. Using the parametrization (17) in the training problem not only improves the computational efficiency of the optimization but also allows for the design of implicit neural networks with additional structure such as convolutions. Suppose $u \in \mathbb{R}^{rs^2}$ is a r -channel input of size $s \times s$ and $x \in \mathbb{R}^{ns^2}$ is an n -channel hidden layer. To define our implicit CNN, we select the weight matrix $A \in \mathbb{R}^{ns^2 \times ns^2}$ as the matrix form of a 2D convolutional operator. If we consider a circular convolution operator, then A is a circulant matrix. Using the parametrization (17), A is circulant if and only if T is circulant. Therefore, the training problem for implicit CNNs can be cast as an unconstrained optimization problem using the above parametrization with a circulant T .

Improving robustness via Lipschitz regularization We now focus on learning robust implicit neural networks with bounded Lipschitz constants via a regularization strategy. Setting both $\|\cdot\|_{\mathcal{U}}$ and $\|\cdot\|_{\mathcal{Y}}$ as $\|\cdot\|_{\infty}$ in the input-output Lipschitz bound (15), we get

$$\begin{aligned} \text{Lip}_{u \rightarrow y} &= \frac{\|B\|_{(\infty, [\eta]^{-1}), (\infty)} \|C\|_{(\infty), (\infty, [\eta]^{-1})}}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+} + \|D\|_{\infty, \infty} \\ &\leq \frac{\frac{1}{2} \|B\|_{(\infty, [\eta]^{-1}), (\infty)}^2 + \|C\|_{(\infty), (\infty, [\eta]^{-1})}^2}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+} + \|D\|_{\infty, \infty}, \end{aligned}$$

where the inequality provides a convex upper bound for the input-output Lipschitz constant. Therefore, using the hyperparameter $\lambda > 0$, the regularized optimization problem is written as

$$\min_{A, B, C, D, X} \mathcal{L}(Y, CX + DU) + \lambda \left(\frac{\frac{1}{2} \|B\|_{(\infty, [\eta]^{-1}), (\infty)}^2 + \|C\|_{(\infty), (\infty, [\eta]^{-1})}^2}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+} + \|D\|_{\infty, \infty} \right) \quad (18)$$

$$X = \Phi(AX + BU), \quad \mu_{\infty, [\eta]^{-1}}(A) \leq \gamma.$$

Certified adversarial robustness via Lipschitz bounds Given a nominal input $u \in \mathbb{R}^r$, we consider any perturbed input v within an ℓ_{∞} -ball of radius ε around u . In this case, we have

$$\|y_u - y_v\|_{\infty} \leq \text{Lip}_{u \rightarrow y} \|u - v\|_{\infty} \leq \text{Lip}_{u \rightarrow y} \varepsilon. \quad (19)$$

Then we define $\text{margin}(u) = (y_u)_i - \max_{j \neq i} (y_u)_j$, where $(y_u)_i$ is the logit corresponding to the (correct) label i for the input u . Then provided $L\varepsilon \leq \frac{1}{2} \text{margin}(u)$, NEMON is certifiably robust to any perturbed input v within an ℓ_{∞} -ball of radius ε centered at u .

Backpropagation of gradients via average iteration From [El Ghaoui et al., 2021] we now show how the average iteration can be used to perform backpropagation via the implicit function theorem. For simplicity, we assume that each activation function ϕ_i is differentiable and consider mini-batches of size 1, i.e., we have $X = x \in \mathbb{R}^n$, $U = u \in \mathbb{R}^r$ and $\hat{Y} = \hat{y} \in \mathbb{R}^q$. Let x^* be the unique solution of the fixed-point equation (11). Then the chain rule implies

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial A} &= (\nabla_{x^*} \mathcal{L}) x^\top, & \frac{\partial \mathcal{L}}{\partial B} &= (\nabla_{x^*} \mathcal{L}) u^\top, \\ \frac{\partial \mathcal{L}}{\partial C} &= (\nabla_{\hat{y}} \mathcal{L}) x^\top, & \frac{\partial \mathcal{L}}{\partial D} &= (\nabla_{\hat{y}} \mathcal{L}) u^\top.\end{aligned}$$

Since \mathcal{L} depends explicitly on \hat{y} , computing $\nabla_{\hat{y}} \mathcal{L}$ is straightforward. Computing $\nabla_{x^*} \mathcal{L}$ is more complicated since x^* is defined only implicitly. However, it be shown that

$$\nabla_{x^*} \mathcal{L} = (C(I - D\Phi A)^{-1} D\Phi)^\top \nabla_{\hat{y}} \mathcal{L}.$$

Since $\mu_{\infty, [\eta]^{-1}}(A) < 1$, by Lemma 8 we get that $\mu_{\infty, [\eta]^{-1}}(D\Phi A) < 1$. This implies that the matrix $G := (I_n - D\Phi A)^{-1} D\Phi \in \mathbb{R}^{n \times n}$ exists and is the solution to the following fixed-point equation [El Ghaoui et al., 2021, Section 6.2]

$$G = D\Phi(AG + I_n). \quad (20)$$

Moreover, $\mu_{\infty, [\eta]^{-1}}(D\Phi A) < 1$ and Theorem 2 together imply that the fixed-point equation (20) has a unique solution G^* and, for every $0 < \alpha \leq \alpha^* := (1 - \min_i (A_{ii})_-)^{-1}$, the average iterations

$$G_{k+1} = (1 - \alpha)G_k + \alpha D\Phi(AG_k + I_n), \quad \text{for all } k \in \mathbb{Z}_{\geq 0}$$

are contracting with the minimum contraction factor $1 - \alpha^*(1 - \mu_{\infty, [\eta]^{-1}}(A)_+)$ at step size α^* .

6 Theoretical and numerical comparisons

In this section, we provide a comprehensive comparison of our framework with the state-of-the-art implicit neural networks².

6.1 Implicit network models

We start by reviewing the existing models for implicit networks in the literature.

Implicit deep learning model. [El Ghaoui et al., 2021] proposes a class of implicit neural networks with input-output behavior described by (11). It is shown that a sufficient condition for existence and uniqueness of a solution and convergence of the Picard iterations for the fixed point equation $x = \Phi(Ax + Bu)$ is $\lambda_{\text{pf}}(|A|) < 1$, where $|A|$ denotes the entrywise absolute value of the matrix A and λ_{pf} denotes the Perron-Frobenius eigenvalue. For training, the optimization problem (16) is used where the constraint $\mu_{\infty, [\eta]^{-1}}(A) \leq \gamma$ is replaced by $\|A\|_\infty \leq \gamma$ [El Ghaoui et al., 2021, Equation 6.3].³ It is easy to see that our well-posedness condition in Corollary 5(i) is less conservative than $\lambda_{\text{pf}}(|A|) < 1$ and its convex relaxation $\|A\|_\infty < 1$.

Monotone operator deep equilibrium network (MON). [Winston and Kolter, 2020] proposes to use monotone operator theory to guarantee well-posedness of the fixed-point equation as well as its convergence to the solutions. The input-output behavior of the network is described by (11). For training, the optimization problem (16) is used where the constraint $\mu_{\infty, [\eta]^{-1}}(A) \leq \gamma$ is replaced by $I_n - \frac{1}{2}(A + A^\top) \succeq (1 - \gamma)I_n$. In order to ensure that this constraint is always satisfied in the training procedure, the weight matrix A is parametrized as $A = \gamma I_n - W^\top W - Z + Z^\top$, for arbitrary $W, Z \in \mathbb{R}^{n \times n}$ [Winston and Kolter, 2020, Appendix D].⁴ In the context of contraction theory,

$$I_n - \frac{1}{2}(A + A^\top) \succeq (1 - \gamma)I_n \iff \mu_2(A) \leq \gamma,$$

which is shown in Appendix A. Thus, the parametrization $A = \gamma I_n - W^\top W - Z + Z^\top$ can be considered as the ℓ_2 -norm version of the parametrization described by equation (17). In other words, the monotone operator network formulation is a Euclidean transcription of the framework we propose in this paper.

²All models were trained using Google Colab with a Tesla P100-PCIE-16GB GPU.

³The implicit deep learning implementation is available at <https://github.com/beeperman/idl>.

⁴The MON implementation is available at https://github.com/locuslab/monotone_op_net.

6.2 MNIST experiments




Figure 1: Performance comparison between the NEMON model with $\mu_\infty(A) \leq 0.95$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$ on the MNIST dataset. The curves are generated by mean accuracy and mean loss over 5 different runs while light envelopes around the curves correspond to the standard deviation over the runs. Average best accuracy for the NEMON model is 0.9772, while it is 0.9721 for implicit deep learning model and 0.9762 for the MON model.

In the digit classification dataset MNIST, input data are 28×28 pixel images of handwritten digits between 0-9. There are 60000 training images and 10000 test images. For training, images are reshaped into 784 dimensional column vectors and entries are scaled into the range $[0, 1]$. As a loss function, we use the cross-entropy. All models are of order $n = 100$, used the ReLU activation function $\phi_i(x) = (x)_+$, and are trained with a batch size of 300 over 10 epochs with a learning rate of 1.5×10^{-2} . Curves for accuracy and loss versus epochs for the three models are shown in Figure 1. Regarding training times, using the average iteration, NEMON took, on average, 12 forward iterations, 13 backward iterations, and 9.8 seconds to train per epoch. Using the Peaceman-Rachford iteration, MON took, on average, 17 forward iterations, 16 backward iterations, and 9.5 seconds to train per epoch. Using the Picard iteration, the implicit deep learning model took, on average, 10 forward iterations, 5 backward iterations, and 5.8 seconds to train per epoch. We observe that the NEMON model performs better than the implicit deep learning model and has a comparable performance to MON.




Figure 2: On the left is a plot of test error versus Lipschitz constant for the implicit deep learning model with $\|A\|_\infty \leq 0.95$ and for NEMON with $\mu_\infty(A) \leq 0.95$ and parametrized by the regularization hyperparameter λ . We define the test error as 1 minus the accuracy. On the right is a plot of accuracy versus ℓ_∞ perturbation of a deterministic adversarial image inversion attack where we additionally include the MON model with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$.

We also study the robustness of the NEMON model compared to the implicit deep learning model and the MON model on the MNIST dataset. We train various models regularized by the input-output Lipschitz constant as in (18). Additionally, to verify robustness of the different models, we consider several adversarial attacks and plot the accuracy versus perturbation of such an attack. In Figure 2, we consider a continuous image inversion attack [Hosseini et al., 2017], where each pixel is perturbed

in the direction of pixel value inversion with amplitude given by the ℓ_∞ perturbation. For more details on this and other types of adversarial perturbations, we refer to Appendix D. We observe that for $\lambda = 10^{-5}$, the regularized NEMON model achieves a two order of magnitude decrease in its input-output Lipschitz constant compared to the un-regularized NEMON models. In addition, we see that the implicit deep learning model and the MON model are more sensitive to the continuous image inversion attack than NEMON. Moreover, as the regularization parameter λ increases, the NEMON model becomes increasingly robust to this attack.

6.3 CIFAR-10 experiments

In the image classification dataset CIFAR-10, input data are 32×32 color images in 10 classes. There are 50000 training images and 10000 test images. We compare our proposed NEMON model with a convolutional structure to a single convolutional layer MON model. Each model used 81 channels. We train both models with a batch size of 256 and a learning rate of 10^{-3} for 40 epochs. For training, using the average iteration, NEMON took, on average, 10 forward iterations, 10 backward iterations, and 75.0 seconds per epoch to train. Using the Peaceman-Rachford iteration, MON took, on average, 5 forward iterations, 5 backward iterations, and 101.8 seconds per epoch to train.

We focus primarily on the robustness of NEMON and MON with respect to ℓ_∞ -norm bounded perturbations on CIFAR-10. To this end, we additionally trained two NEMON models with regularization parameters $\lambda \in \{10^{-4}, 10^{-5}\}$. In Figure 3, on the left is a plot of the certified robustness of each of the models via their ℓ_∞ -Lipschitz constants. For MON, we got the ℓ_∞ -Lipschitz bound using the method in [Pabbaraju et al., 2021] for the ℓ_2 -Lipschitz bound and using the upper bound $\|u\|_2 \leq \sqrt{rs^2}\|u\|_\infty$. On the right is a plot of the accuracy of different models with respect to the projected gradient descent attack. We observe that the un-regularized and regularized NEMON models are more robust to ℓ_∞ -norm bounded perturbations than is MON.




Figure 3: On the left is a plot of certified robustness via the Lipschitz constants of MON with the constraint $I_n - \frac{1}{2}(A + A^T) \succeq I_n$ and NEMON with the constraint $\mu_\infty(A) \leq 0$. On the right is a plot of accuracy versus ℓ_∞ perturbation of the projected gradient descent attack.

7 Conclusion

Using non-Euclidean contraction theory, we propose a framework to study stability of fixed-point equations. We apply this framework to analyze well-posedness and convergence of implicit neural networks and to design an efficient training algorithm to incorporate robustness guarantees. For future research, we envision that our framework is applicable to study stability and robustness of implicit learning models with additional structure such as graph neural networks.

8 Acknowledgments

The authors thank Ian Manchester for stimulating discussions about contraction theory and the anonymous reviewers for their insightful feedback. This work was supported in part by DTRA under grant HDTRA1-19-1-0017, AFOSR under grant FA9550-22-1-0059, and NSF Graduate Research Fellowship under grant 2139319.

References

- A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, 2019. URL <https://arxiv.org/abs/1910.12430>.
- B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter. Differentiable MPC for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, 2018. URL <https://arxiv.org/abs/1810.13400>.
- S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, 2019. URL <https://arxiv.org/abs/1909.01377>.
- H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2 edition, 2017. ISBN 978-3-319-48310-8.
- N. Carlini and D. Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *ACM Workshop on Artificial Intelligence and Security*, pages 3–14, 2017. doi:[10.1145/3128572.3140444](https://doi.org/10.1145/3128572.3140444).
- R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, 2018. URL <https://arxiv.org/abs/1806.07366>.
- P. L. Combettes and J.-C. Pesquet. Lipschitz certificates for layered network structures driven by averaged activation operators. *SIAM Journal on Mathematics of Data Science*, 2(2):529–557, 2020. doi:[10.1137/19M1272780](https://doi.org/10.1137/19M1272780).
- P. L. Combettes and J.-C. Pesquet. Fixed point strategies in data science. *IEEE Transactions on Signal Processing*, 2021. doi:[10.1109/TSP.2021.3069677](https://doi.org/10.1109/TSP.2021.3069677).
- A. Davydov, S. Jafarpour, and F. Bullo. Non-Euclidean contraction theory for robust nonlinear stability. 2021. URL <https://arxiv.org/abs/2103.12263>.
- C. A. Desoer and H. Haneda. The measure of a matrix as a tool to analyze computer algorithms for circuit analysis. *IEEE Transactions on Circuit Theory*, 19(5):480–486, 1972. doi:[10.1109/TCT.1972.1083507](https://doi.org/10.1109/TCT.1972.1083507).
- L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Tsai. Implicit deep learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021. doi:[10.1137/20M1358517](https://doi.org/10.1137/20M1358517).
- Y. Fang and T. G. Kincaid. Stability analysis of dynamical neural networks. *IEEE Transactions on Neural Networks*, 7(4):996–1006, 1996. doi:[10.1109/72.508941](https://doi.org/10.1109/72.508941).
- M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. J. Pappas. Efficient and accurate estimation of Lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, 2019. URL <https://arxiv.org/abs/1906.04893>.
- M. Fazlyab, M. Morari, and G. J. Pappas. Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming. *IEEE Transactions on Automatic Control*, 2020. doi:[10.1109/TAC.2020.3046193](https://doi.org/10.1109/TAC.2020.3046193).
- I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015. URL <https://arxiv.org/abs/1412.6572>.
- W. He and J. Cao. Exponential synchronization of chaotic neural networks: a matrix measure approach. *Nonlinear Dynamics*, 55:55–65, 2009. doi:[10.1007/s11071-008-9344-4](https://doi.org/10.1007/s11071-008-9344-4).
- H. Hosseini, B. Xiao, M. Jaiswal, and R. Poovendran. On the limitation of convolutional neural networks in recognizing negative images. In *IEEE International Conference on Machine Learning and Applications*, pages 352–358, 2017. doi:[10.1109/ICMLA.2017.0-136](https://doi.org/10.1109/ICMLA.2017.0-136).

- A. Kag, Z. Zhang, and V. Saligrama. RNNs incrementally evolving on an equilibrium manifold: A panacea for vanishing and exploding gradients? In International Conference on Learning Representations, 2020. URL <https://openreview.net/forum?id=HylpqA4FwS>.
- K. Kawaguchi. On the theory of implicit deep learning: Global convergence with implicit layers. In International Conference on Learning Representations, 2021. URL <https://openreview.net/forum?id=p-NZIuwqhI4>.
- V. Krishnan, A. A. A. Makdah, and F. Pasqualetti. Lipschitz bounds and provably robust training by Laplacian smoothing. In Advances in Neural Information Processing Systems, 2020. URL <https://arxiv.org/abs/2006.03712>.
- T. C. Lim. On fixed point stability for set-valued contractive mappings with applications to generalized differential equations. Journal of Mathematical Analysis and Applications, 110(2):436–441, 1985. doi:[10.1016/0022-247X\(85\)90306-3](https://doi.org/10.1016/0022-247X(85)90306-3).
- W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. Automatica, 34(6): 683–696, 1998. doi:[10.1016/S0005-1098\(98\)00019-3](https://doi.org/10.1016/S0005-1098(98)00019-3).
- A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In International Conference on Machine Learning, 2018. URL <https://arxiv.org/abs/1706.06083>.
- C. Pabbaraju, E. Winston, and J. Z. Kolter. Estimating Lipschitz constants of monotone deep equilibrium models. In International Conference on Learning Representations, 2021. URL <https://openreview.net/forum?id=VcB4QkSfy0>.
- N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In IEEE Symposium on Security and Privacy, pages 582–597, 2016. doi:[10.1109/SP.2016.41](https://doi.org/10.1109/SP.2016.41).
- P. Pauli, A. Koch, J. Berberich, P. Kohler, and F. Allgower. Training robust neural networks using Lipschitz bounds. IEEE Control Systems Letters, 2021. doi:[10.1109/LCSYS.2021.3050444](https://doi.org/10.1109/LCSYS.2021.3050444).
- A. Pavlov, A. Pogromsky, N. Van de Wouw, and H. Nijmeijer. Convergent dynamics, a tribute to Boris Pavlovich Demidovich. Systems & Control Letters, 52(3-4):257–261, 2004. doi:[10.1016/j.sysconle.2004.02.003](https://doi.org/10.1016/j.sysconle.2004.02.003).
- H. Qiao, J. Peng, and Z.-B. Xu. Nonlinear measures: A new approach to exponential stability analysis for Hopfield-type neural networks. IEEE Transactions on Neural Networks, 12(2):360–370, 2001. doi:[10.1109/72.914530](https://doi.org/10.1109/72.914530).
- M. Revay, R. Wang, and I. R. Manchester. Lipschitz bounded equilibrium networks. 2020. URL <https://arxiv.org/abs/2010.01732>.
- M. Revay, R. Wang, and I. 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](https://doi.org/10.1109/LCSYS.2020.3038221).
- E. K. Ryu and S. Boyd. Primer on monotone operator methods. Applied Computational Mathematics, 15(1):3–43, 2016.
- C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In International Conference on Learning Representations, 2014. URL <https://arxiv.org/abs/1312.6199>.
- E. Winston and J. Z. Kolter. Monotone operator equilibrium networks. In Advances in Neural Information Processing Systems, 2020. URL <https://arxiv.org/abs/2006.08591>.
- E. Wong and J. Z. Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In International Conference on Machine Learning, pages 5286–5295, 2018. URL <http://proceedings.mlr.press/v80/wong18a.html>.

A A comprehensive review of non-Euclidean contraction theory

Matrix measures Let $\|\cdot\|$ be a norm on \mathbb{R}^n and its induced norm on $\mathbb{R}^{n \times n}$. The matrix measure of $A \in \mathbb{R}^{n \times n}$ with respect to $\|\cdot\|$ is

$$\mu(A) := \lim_{h \rightarrow 0^+} \frac{\|I_n + hA\| - 1}{h}. \quad (21)$$

It is well known that this limit is well posed because the right-hand side is non-increasing in h , due to the convexity of the norm. For arbitrary $n \times n$ matrices A and B , the following properties hold:

$$\text{sub-additivity:} \quad \mu(A + B) \leq \mu(A) + \mu(B), \quad (22a)$$

$$\text{weak homogeneity:} \quad \mu(\alpha A) = \alpha \mu(A), \quad \forall \alpha \geq 0, \quad (22b)$$

$$\text{convexity:} \quad \mu(\theta A + (1 - \theta)B) \leq \theta \mu(A) + (1 - \theta) \mu(B), \quad \forall \theta \in [0, 1], \quad (22c)$$

$$\text{norm/spectrum:} \quad -\|A\| \leq -\mu(-A) \leq \Re(\lambda) \leq \mu(A) \leq \|A\|, \quad \forall \lambda \in \text{spec}(A), \quad (22d)$$

$$\text{translation:} \quad \mu(A + cI_n) = \mu(A) + c, \quad \forall c \in \mathbb{R}, \quad (22e)$$

$$\text{product:} \quad \max\{-\mu(A), -\mu(-A)\} \|x\| \leq \|Ax\|, \quad \forall x \in \mathbb{R}^n, \quad (22f)$$

$$\text{norm of inverse:} \quad \mu(A) < 0 \implies \|A^{-1}\| \leq -1/\mu(A). \quad (22g)$$

Note that convexity is an immediate consequence of sub-additivity and weak homogeneity. Additionally, by property (22d), the matrix measure is upper bounded by the matrix norm and may be negative. We refer to [Desoer and Haneda, 1972], and references therein, for the proof of these and additional properties enjoyed by matrix measures.

We will be specifically interested in diagonally weighted ℓ_1 and ℓ_∞ norms defined by

$$\|x\|_{1, [\eta]} = \sum_i \eta_i |x_i| \quad \text{and} \quad \|x\|_{\infty, [\eta]^{-1}} = \max_i \frac{1}{\eta_i} |x_i|, \quad (23)$$

where, given a positive vector $\eta \in \mathbb{R}_{>0}^n$, we use $[\eta]$ to denote the diagonal matrix with diagonal entries η . The corresponding matrix norms and measures are

$$\|A\|_{1, [\eta]} = \max_{j \in \{1, \dots, n\}} \sum_{i=1}^n \frac{\eta_i}{\eta_j} |a_{ij}|, \quad \mu_{1, [\eta]}(A) = \max_{j \in \{1, \dots, n\}} \left(a_{jj} + \sum_{i=1, i \neq j}^n |a_{ij}| \frac{\eta_i}{\eta_j} \right), \quad (24)$$

$$\|A\|_{\infty, [\eta]^{-1}} = \max_{i \in \{1, \dots, n\}} \sum_{j=1}^n \frac{\eta_j}{\eta_i} |a_{ij}|, \quad \mu_{\infty, [\eta]^{-1}}(A) = \max_{i \in \{1, \dots, n\}} \left(a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| \frac{\eta_j}{\eta_i} \right). \quad (25)$$

Finally, we include the Euclidean norm ℓ_2 . Given a positive definite P , we define the weighted ℓ_2 norm by

$$\|x\|_{2, P^{1/2}} = \sqrt{x^\top P x}.$$

Then the following equalities are well known, e.g., see [Desoer and Haneda, 1972, Davydov et al., 2021],

$$\mu_{2, P^{1/2}}(A) = \lambda_{\max} \left(\frac{PAP^{-1} + A^\top}{2} \right) = \min\{b \in \mathbb{R} \mid A^\top P + PA \preceq 2bP\} \quad (26)$$

$$= \max\{x^\top P A x \mid x^\top P x = 1\}. \quad (27)$$

Weak pairings We briefly review the notion of a weak pairing (WP) on \mathbb{R}^n from [Davydov et al., 2021]. A WP on \mathbb{R}^n is a map $\llbracket \cdot, \cdot \rrbracket : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying:

- (i) (sub-additivity and continuity of first argument) $\llbracket x_1 + x_2, y \rrbracket \leq \llbracket x_1, y \rrbracket + \llbracket x_2, y \rrbracket$, for all $x_1, x_2, y \in \mathbb{R}^n$ and $\llbracket \cdot, \cdot \rrbracket$ is continuous in its first argument,
- (ii) (weak homogeneity) $\llbracket \alpha x, y \rrbracket = \llbracket x, \alpha y \rrbracket = \alpha \llbracket x, y \rrbracket$ and $\llbracket -x, -y \rrbracket = \llbracket x, y \rrbracket$, for all $x, y \in \mathbb{R}^n, \alpha \geq 0$,
- (iii) (positive definiteness) $\llbracket x, x \rrbracket > 0$, for all $x \neq 0_n$,
- (iv) (Cauchy-Schwarz inequality) $|\llbracket x, y \rrbracket| \leq \llbracket x, x \rrbracket^{1/2} \llbracket y, y \rrbracket^{1/2}$, for all $x, y \in \mathbb{R}^n$.

For every norm $\|\cdot\|$ on \mathbb{R}^n , there exists a (possibly not unique) compatible WP $\llbracket \cdot, \cdot \rrbracket$ such that $\|x\|^2 = \llbracket x, x \rrbracket$, for every $x \in \mathbb{R}^n$. If the norm is induced by an inner product, the WP coincides with the inner product.

Specifically, from [Davydov et al., 2021, Table III], we introduce the WPs $\llbracket \cdot, \cdot \rrbracket_{1, [\eta]} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ and $\llbracket \cdot, \cdot \rrbracket_{\infty, [\eta]^{-1}} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, defined by

$$\llbracket x, y \rrbracket_{1, [\eta]} = \|y\|_{1, [\eta]} \text{sign}(y)^\top [\eta] x \quad \text{and} \quad \llbracket x, y \rrbracket_{\infty, [\eta]^{-1}} = \max_{i \in I_\infty([\eta]^{-1}y)} \eta_i^{-2} y_i x_i. \quad (28)$$

where $I_\infty(x) = \{i \in \{1, \dots, n\} \mid |x_i| = \|x\|_\infty\}$. One can show the so-called Lumer equalities (generalizing equation (27)):

$$\mu_{1, [\eta]}(A) = \max_{\|x\|_{1, [\eta]}=1} \text{sign}(x)^\top [\eta] A x, \quad (29)$$

$$\mu_{\infty, [\eta]^{-1}}(A) = \max_{\|x\|_{\infty, [\eta]^{-1}}=1} \max_{i \in I_\infty([\eta]^{-1}x)} ([\eta]^{-1}x)_i ([\eta]^{-1}A x)_i. \quad (30)$$

Lipschitz maps Given a norm $\|\cdot\|$ with induced matrix norm $\|\cdot\|$ and induced matrix measure $\mu(\cdot)$, a map $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is Lipschitz continuous with constant $\text{Lip}(F) \in \mathbb{R}_{\geq 0}$ if

$$\|F(x_1) - F(x_2)\| \leq \text{Lip}(F) \|x_1 - x_2\| \quad \text{for all } x_1, x_2 \in \mathbb{R}^n. \quad (31)$$

If the map F is differentiable, then F is Lipschitz continuous with constant $\text{Lip}(F)$ if and only if

$$\|DF(x)\| \leq \text{Lip}(F) \quad \text{for all } x \in \mathbb{R}^n. \quad (32)$$

One-sided Lipschitz maps Given a norm $\|\cdot\|$ with compatible WP $\llbracket \cdot, \cdot \rrbracket$ and associated matrix measure $\mu(\cdot)$, a continuous map $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is one-sided Lipschitz continuous with constant $\text{osL}(F) \in \mathbb{R}$ if

$$\llbracket F(x_1) - F(x_2), x_1 - x_2 \rrbracket \leq \text{osL}(F) \|x_1 - x_2\|^2 \quad \text{for all } x_1, x_2 \in \mathbb{R}^n. \quad (33)$$

If the map F is differentiable, then F is one-sided Lipschitz continuous with constant $\text{osL}(F) \in \mathbb{R}$ if and only if

$$\mu(DF(x)) \leq \text{osL}(F) \quad \text{for all } x \in \mathbb{R}^n. \quad (34)$$

In other words, when the map F is differentiable, the two definitions (33) and (34) are equivalent. Note that (i) the one-sided Lipschitz constant is upper bounded by the Lipschitz constant, (ii) a Lipschitz map is always one-sided Lipschitz, but the converse is not necessarily true, and (iii) the one-sided Lipschitz constant may be negative. For instance, consider the scalar function $f(x) = -x - x^3$. It is easy to check that this function is not globally Lipschitz and $\text{Lip}(f) = \infty$. However, f is one-sided Lipschitz with $\text{osL}(f) = -1$.

In the following example, we compare the regions $\text{Lip}(A) < 1$ and $\text{osL}(A) < 1$ for a matrix $A \in \mathbb{R}^{2 \times 2}$ with respect to the ℓ_∞ -norm.

Example 6. Let $A = \begin{bmatrix} a & b \\ b & a \end{bmatrix}$, it is easy to see that condition $\text{Lip}(A) < 1$ for ℓ_∞ -norm can be written as $\|A\|_\infty = |a| + |b| < 1$. One can also define the average operator A_α using parameter $\alpha \in (0, 1]$ as follows:

$$A_\alpha = (1 - \alpha)I_2 + \alpha A.$$

Figure 4 compares the regions $\text{Lip}(A) < 1$, $\text{Lip}(A_\alpha) < 1$, and $\text{osL}(A) < 1$ based on the parameters a and b . It can be shown that as $\alpha \rightarrow 0^+$, the condition $\text{Lip}(A_\alpha) < 1$ converges to $\text{osL}(A) < 1$.

B Novel results about non-Euclidean matrix measures

In this appendix we provide some results regarding the matrix measure and matrix norm for weighted ℓ_1 and ℓ_∞ -norms.

Lemma 7 (Non-Euclidean contraction estimates). Let $A = [a_{ij}] \in \mathbb{R}^{n \times n}$ and $\eta \in \mathbb{R}_{>0}^n$,




Figure 4: The left figure shows the region $\text{Lip}(A) \leq 1$, the middle figure shows the region $\text{Lip}(A_\alpha) \leq 1$ for $\alpha = \frac{1}{2}$, and the right figure shows $\text{osL}(A) \leq 1$. Both Lip and osL are with respect to the ℓ_∞ -norm

(i) For every $\alpha \in \mathbb{R}$ such that $|\alpha| \leq (\max_i |a_{ii}|)^{-1}$,

$$\begin{aligned} \|I_n + \alpha A\|_{1, [\eta]} &= 1 + \alpha \mu_{1, [\eta]}(A), \\ \|I_n + \alpha A\|_{\infty, [\eta]}^{-1} &= 1 + \alpha \mu_{\infty, [\eta]}^{-1}(A). \end{aligned}$$

(ii) the minimizer and minimum value of $\min_{\alpha \geq 0} \|I_n + \alpha A\|_{\infty, [\eta]}^{-1}$ can be computed via the linear program:

$$\begin{aligned} \min_{\alpha, t} \quad & t \\ & 1 + \alpha(a_{ii} + r_i) \leq t, \quad i \in \{1, \dots, n\}, \\ & -1 + \alpha(-a_{ii} + r_i) \leq t, \quad i \in \{1, \dots, n\}, \\ & \alpha \geq 0. \end{aligned}$$

$$\text{where } r_i = \sum_{j \neq i} \frac{\eta_j}{\eta_i} |a_{ij}|.$$

Proof. Regarding part (i), we compute

$$\|I_n + \alpha A\|_{\infty, [\eta]}^{-1} = \max_{i \in \{1, \dots, n\}} \left\{ |1 + \alpha a_{ii}| + \alpha \sum_{j=1, j \neq i}^n \frac{\eta_j}{\eta_i} |a_{ij}| \right\}. \quad (35)$$

Since $|\alpha| \leq (\max_i |a_{ii}|)^{-1}$, we know $|\alpha| |a_{ii}| \leq 1$ for all $i \in \{1, \dots, n\}$. Therefore $1 + \alpha a_{ii} \geq 0$ and $|1 + \alpha a_{ii}| = 1 + \alpha a_{ii}$, for every $i \in \{1, \dots, n\}$. In summary, replacing in (35),

$$\|I_n + \alpha A\|_{\infty, [\eta]}^{-1} = \max_{i \in \{1, \dots, n\}} \left\{ 1 + \alpha a_{ii} + \alpha \sum_{j=1, j \neq i}^n \frac{\eta_j}{\eta_i} |a_{ij}| \right\} = 1 + \alpha \mu_{\infty, [\eta]}^{-1}(A).$$

The proof of the formula relating the weighted ℓ_1 -norm and the weighted ℓ_1 -matrix measure will follow mutatis mutandis to the above proof for ℓ_∞ -norm and we omit it in the interest of brevity.

Regarding part (ii), using formula (35), we get

$$\begin{aligned} \|I_n + \alpha A\|_{\infty, [\eta]}^{-1} &= \max_{i \in \{1, \dots, n\}} \{ |1 + \alpha a_{ii}| + \alpha r_i \} \\ &= \max_{i \in \{1, \dots, n\}} \{ 1 + \alpha a_{ii} + \alpha r_i, -1 - \alpha a_{ii} + \alpha r_i \}. \end{aligned}$$

The result then follows. \square

The following results are related to [Fang and Kincaid, 1996, Theorem 3.8] and [He and Cao, 2009, Lemma 3] and, indirectly, to [Qiao et al., 2001]. In comparison with [Fang and Kincaid, 1996, He and Cao, 2009], we prove sharper bounds for a more general setting.

Lemma 8 (Matrix measure inequalities under multiplicative scalings). *For each $A \in \mathbb{R}^{n \times n}$, $C \in \mathbb{R}^{n \times n}$ diagonal positive, and $\eta \in \mathbb{R}_{>0}^n$,*

- (i) $\max_{d \in [0,1]^n} \mu_{\infty, [\eta]}(-C + [d]A) = \max \{ \mu_{\infty, [\eta]}(-C), \mu_{\infty, [\eta]}(-C + A) \}$, and
- (ii) $\max_{d \in [0,1]^n} \mu_{1, [\eta]}(-C + A[d]) = \max \{ \mu_{1, [\eta]}(-C), \mu_{1, [\eta]}(-C + A) \}$.

Proof. Define the short-hand $r_i = a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| \eta_i / \eta_j$ and note

$$\begin{aligned} \mu_{\infty, [\eta]}(-C) &= \max_{i \in \{1, \dots, n\}} \{-c_i\}, \quad \mu_{\infty, [\eta]}(-C + A) = \max_{i \in \{1, \dots, n\}} \{-c_i + r_i\}, \quad \text{and} \\ \mu_{\infty, [\eta]}(-C + [d]A) &= \max_{i \in \{1, \dots, n\}} \{-c_i + d_i r_i\}. \end{aligned}$$

Since $0 \leq d_i \leq 1$, we note

$$\begin{aligned} r_i \leq 0 &\implies d_i r_i \leq 0 \implies -c_i + d_i r_i \leq -c_i, \\ r_i > 0 &\implies d_i r_i \geq 0 \implies -c_i + d_i r_i \leq -c_i + r_i. \end{aligned}$$

Therefore

$$\begin{aligned} \max_{d \in [0,1]^n} \max_{i: r_i \leq 0} \{-c_i + d_i r_i\} &= \max_{i: r_i \leq 0} \max_{d_i \in [0,1]} \{-c_i + d_i r_i\} = \max_{i: r_i \leq 0} \{-c_i\} \leq \mu_{\infty, [\eta]}(-C), \\ \max_{d \in [0,1]^n} \max_{i: r_i > 0} \{-c_i + d_i r_i\} &= \max_{i: r_i > 0} \max_{d_i \in [0,1]} \{-c_i + d_i r_i\} = \max_{i: r_i > 0} \{-c_i + r_i\} \leq \mu_{\infty, [\eta]}(-C + A). \end{aligned}$$

In summary

$$\begin{aligned} \max_{d \in [0,1]^n} \mu_{\infty, [\eta]}(-C + [d]A) &= \max_{d \in [0,1]^n} \max_{i \in \{1, \dots, n\}} \{-c_i + d_i r_i\} \\ &= \max_{d \in [0,1]^n} \max \left\{ \max_{i: r_i \leq 0} \{-c_i + d_i r_i\}, \max_{i: r_i > 0} \{-c_i + d_i r_i\} \right\} \\ &\leq \max \{ \mu_{\infty, [\eta]}(-C), \mu_{\infty, [\eta]}(-C + A) \}. \end{aligned}$$

On the other hand, we note that

$$\begin{aligned} \max_{d \in [0,1]^n} \mu_{\infty, [\eta]}([d]A - C) &\geq \max \{ \mu_{\infty, [\eta]}([0_n]A - C), \mu_{\infty, [\eta]}([1_n]A - C) \} \\ &= \max \{ \mu_{\infty, [\eta]}(-C), \mu_{\infty, [\eta]}(-C + A) \}, \end{aligned}$$

thereby proving the equality in statement (i). Next, recall $\mu_{1, [\eta]}(B) = \mu_{\infty, [\eta]}(B^\top)$ for all B and compute

$$\begin{aligned} \max_{d \in [0,1]^n} \mu_{1, [\eta]}(-C + A[d]) &= \max_{d \in [0,1]^n} \mu_{\infty, [\eta]}(-C + [d]A^\top) \\ &= \max \{ \mu_{\infty, [\eta]}(-C), \mu_{\infty, [\eta]}(-C + A^\top) \} \\ &= \max \{ \mu_{1, [\eta]}(-C), \mu_{\infty, [\eta]}(-C + A) \}. \end{aligned}$$

This concludes the proof of statement (ii). \square

In the same style as [Winston and Kolter, 2020, Proposition 1] and [Revay et al., 2020, Theorems 1 and 2], the next lemma provides a parametrization of all matrices satisfying a μ_∞ constraint.

Lemma 9 (Parametrization of matrices with bounded ℓ_∞ measure). *For any $\gamma \in \mathbb{R}$,*

- (i) *given any $A \in \mathbb{R}^{n \times n}$ with $\mu_\infty(A) \leq \gamma$, there exists a $T \in \mathbb{R}^{n \times n}$ such that $A = T - \text{diag}(|T| \mathbb{1}_n) + \gamma I_n$,*
- (ii) *given any $T \in \mathbb{R}^{n \times n}$, the matrix $A = T - \text{diag}(|T| \mathbb{1}_n) + \gamma I_n \in \mathbb{R}^{n \times n}$ satisfies $\mu_\infty(A) \leq \gamma$,*

where we let $|T|$ denote the entry-wise absolute value of T .

Proof. Regarding statement (i), define

$$\begin{aligned} t_{ij} &= a_{ij} && \text{for all } i \neq j \in \{1, \dots, n\}, \\ t_{ii} &= \frac{1}{2} \left(a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| - \gamma \right), && \text{for } i \in \{1, \dots, n\}. \end{aligned}$$

Because $\mu_\infty(A) \leq \gamma$, we know $a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| \leq \gamma$ for each i . This implies that $t_{ii} \leq 0$ and therefore $t_{ii} - |t_{ii}| = a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| - \gamma$. It is an easy transcription now to show that this equality and the off-diagonal equality $t_{ij} = a_{ij}$ together imply $A = T - \text{diag}(|T|1_n) + \gamma I_n$.

Regarding statement (ii), note that $a_{ij} = t_{ij}$ for all $j \neq i$, and $a_{ii} = t_{ii} - \sum_{j=1}^n |t_{ij}| + \gamma$. Then, for all i ,

$$\begin{aligned} a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| &= (t_{ii} - \sum_{j=1}^n |t_{ij}| + \gamma) + \sum_{j=1, j \neq i}^n |t_{ij}| \\ &= t_{ii} - |t_{ii}| + \gamma = \begin{cases} \gamma, & \text{if } t_{ii} \geq 0, \\ -2|t_{ii}| + \gamma, & \text{if } t_{ii} < 0. \end{cases} \end{aligned}$$

Therefore, $a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| \leq \gamma$ for all i and, in turn, $\mu_\infty(A) \leq \gamma$. \square

We conclude with a simple graph-theoretical interpretation of the main well-posedness condition $\mu_\infty(A) < 1$. Loosely speaking, we call $-a_{ii}$ the self-attenuation of neuron i and $\sum_{j=1, j \neq i}^n |a_{ij}|$ the strength of its outgoing synapses. Then

$$\begin{aligned} \mu_\infty(A) < 1 &\iff a_{ii} + \sum_{j=1, j \neq i}^n |a_{ij}| < 1 \quad \text{for all } i \\ &\iff \text{for each neuron, strength of outgoing synapses} < 1 + \text{self-attenuation}. \end{aligned} \quad (36)$$

C Proofs and additional results on non-differentiable activation functions

C.1 Proofs of Theorems 1 and 2

Proof of Theorem 1. Regarding (ii) \implies (i), note that, for every $x \in \mathbb{R}^n$ and every $0 < \alpha \leq \alpha^*$,

$$\mu(DF_\alpha(x)) \leq \|DF_\alpha(x)\| \leq \gamma_{\ell, c}(\alpha).$$

As a result, $\alpha\mu(DF(x)) = \mu(DF_\alpha(x)) - 1 + \alpha \leq -1 + \alpha + \gamma_{\ell, c}(\alpha)$. Thus,

$$\mu(DF(x)) \leq 1 - \frac{1 - \gamma_{\ell, c}(\alpha)}{\alpha}, \quad \text{for all } x \in \mathbb{R}^n.$$

By choosing $\alpha = \hat{\alpha} = \frac{2c}{(2c+\ell+1)(\ell+1)} < \frac{c}{(c+\ell+1)(\ell+1)}$, we get

$$\mu(DF(x)) \leq 1 - \frac{1 - \gamma_{\ell, c}(\hat{\alpha})}{\hat{\alpha}} = 1 - \frac{1 - (1 - \hat{\alpha}c)}{\hat{\alpha}} = 1 - c, \quad \text{for all } x \in \mathbb{R}^n.$$

Thus, $\sup_{x \in \mathbb{R}^n} \mu(DF(x)) \leq 1 - c$. This implies that $\text{osL}(F) \leq 1 - c$.

Regarding (i) \implies (ii), using the mean value theorem for vector valued functions, we compute

$$\|F_\alpha(x) - F_\alpha(y)\| = \left\| \int_0^1 DF_\alpha(tx + (1-t)y) dt (x - y) \right\| \leq \|\overline{DF}_\alpha(x, y)\| \|x - y\|,$$

where $\overline{DF}_\alpha(x, y) = \int_0^1 DF_\alpha(tx + (1-t)y) dt$, for every $x, y \in \mathbb{R}^n$.

Next, to obtain an upper bound on $\|\overline{DF}_\alpha(x, y)\|$, we first derive a lower bound on $\|\overline{DF}_\alpha^{-1}(x, y)\|$. We start by noting that, the product property (22f) implies $\|Av\| \geq -\mu(-A)\|v\|$, for every $v \in \mathbb{R}^n$ and every $A \in \mathbb{R}^{n \times n}$. Therefore, for every $v \in \mathbb{R}^n$,

$$\|\overline{DF}_\alpha^{-1}(x, y)v\| \geq -\mu(-\overline{DF}_\alpha(x, y))\|v\|. \quad (37)$$

Since $\overline{DF}_\alpha(x, y) = I_n + \alpha(-I_n + \overline{DF}(x, y))$ and $\alpha < \frac{c}{(c+\ell+1)(\ell+1)} \leq \frac{1}{\ell+1}$, we can use the Neumann series to get

$$\overline{DF}_\alpha^{-1}(x, y) = \sum_{i=0}^{\infty} (-1)^i \alpha^i (-I_n + \overline{DF}(x, y))^i. \quad (38)$$

We first compute an upper bound for $\mu(\overline{DF}(x))$. Since $\text{osL}(\mathbf{F}) \leq 1 - c$, by the subadditive property (22a) of the matrix measures, we get

$$\begin{aligned}\mu(-I_n + \overline{DF}(x, y)) &= \mu\left(\int_0^1 (-I_n + DF(tx + (1-t)y))dt\right) \\ &\leq \int_0^1 \mu(-I_n + DF(tx + (1-t)y))dt \leq -c.\end{aligned}\quad (39)$$

Now, we use equation (38) to obtain

$$\begin{aligned}\|\overline{DF}_\alpha^{-1}(x, y)v\| &\geq -\mu\left(\sum_{i=0}^{\infty} (-1)^{i+1} \alpha^i (-I_n + \overline{DF}(x, y))^i\right) \|v\| \\ &\geq -\left(\mu(-I_n) + \alpha\mu(-I_n + \overline{DF}(x, y))\right. \\ &\quad \left.+ \sum_{i=2}^{\infty} \alpha^i \mu((-1)^{i+1} (-I_n + \overline{DF}(x, y))^i)\right) \|v\| \\ &\geq (1 + \alpha c - \sum_{i=2}^{\infty} (\alpha(\ell+1))^i) \|v\| = \left(1 + \alpha c - \frac{\alpha^2(\ell+1)^2}{1 - \alpha(\ell+1)}\right) \|v\|,\end{aligned}\quad (40)$$

where the first inequality holds by (37), the second inequality holds by subadditive property of the matrix measures (22a), and the third inequality holds because, using (39) and (22d), we obtain the upper bound:

$$\mu((-1)^{i+1} (-I_n + \overline{DF}(x, y))^i) \leq \|(-I_n + \overline{DF}(x, y))^i\| \leq (1 + \ell)^i, \quad \text{for all } i \in \mathbb{Z}_{\geq 0}.$$

Note that $\alpha \in]0, \frac{c}{(c+\ell+1)(\ell+1)}[$. Equation (40) implies that, for each $w \in \mathbb{R}^n$ and $v = \overline{DF}_\alpha(x, y)w$,

$$\frac{\|\overline{DF}_\alpha(x, y)w\|}{\|w\|} = \frac{\|v\|}{\|\overline{DF}_\alpha^{-1}(x, y)v\|} \leq \gamma_{\ell, c}(\alpha).$$

As a result, $\|\overline{DF}_\alpha(x, y)\| \leq \gamma_{\ell, c}(\alpha)$ and

$$\|\mathbf{F}_\alpha(x) - \mathbf{F}_\alpha(y)\| \leq \gamma_{\ell, c}(\alpha) \|x - y\|, \quad \text{for all } x, y \in \mathbb{R}^n.$$

Regarding parts (iii) and (iv), a straightforward calculation shows that, if $0 < \alpha < \frac{c}{(c+\ell+1)(\ell+1)}$, then $1/\left(1 + \alpha c - \frac{\alpha^2(\ell+1)^2}{1 - \alpha(\ell+1)}\right) < 1$. The result then follows from the Banach fixed-point theorem.

Regarding part (v), we define the function $\xi :]0, \frac{c}{(c+\ell+1)(\ell+1)}[\rightarrow \mathbb{R}_{>0}$ by $\xi(\alpha) = 1 + \alpha c - \frac{\alpha^2(\ell+1)^2}{1 - \alpha(\ell+1)}$. Then it is clear that $\xi(\alpha) = 1/\gamma_{\ell, c}(\alpha)$. Note that

$$\begin{aligned}\frac{d\xi}{d\alpha} &= (c + \ell + 1) - \frac{\ell + 1}{(1 - \alpha(\ell + 1))^2}, \\ \frac{d^2\xi}{d\alpha^2} &= -\frac{2(\ell + 1)^2}{(1 - \alpha(\ell + 1))^3}.\end{aligned}$$

Since $\frac{d^2\xi}{d\alpha^2} \leq 0$, we conclude that ξ is a concave function on $]0, \frac{c}{(c+\ell+1)(\ell+1)}[$ and its maximum is achieved at α^* for which $\frac{d\xi}{d\alpha}(\alpha^*) = 0$. By a straightforward calculation, we get

$$\alpha^* = \frac{\kappa}{c} \left(1 - \frac{1}{\sqrt{1 + 1/\kappa}}\right)$$

and it is easy to see that the optimal value is as claimed in the theorem statement. \square

Proof of Theorem 2. We restrict ourselves to the norm $\|\cdot\|_{\infty, [\eta]^{-1}}$; the proof for $\|\cdot\|_{1, [\eta]}$ is similar and omitted in the interest of brevity.

Regarding part (i), first we note that $\text{diagL}(F) \leq \text{osL}(F) < 1$, since for every $i \in \{1, \dots, n\}$ and every $x \in \mathbb{R}^n$

$$DF_{ii}(x) \leq DF_{ii}(x) + \sum_{j \neq i} |DF_{ij}(x)| \frac{\eta_i}{\eta_j} = \mu_{\infty, [\eta]^{-1}}(DF(x)) \leq \text{osL}(F) < 1. \quad (41)$$

This implies that $\frac{1}{1 - \text{diagL}(F)} > 0$ and $(1 - \text{osL}(F))/(1 - \text{diagL}(F)) \leq 1$. Moreover, for every $x \in \mathbb{R}^n$,

$$\|(1 - \alpha)I_n + \alpha DF(x)\|_{\infty, [\eta]^{-1}} = \|I_n + \alpha(-I_n + DF(x))\|_{\infty, [\eta]^{-1}}.$$

Next, we study the diagonal entries of $-I_n + DF(x)$. By the definition of $\text{diagL}(F)$ and by equation (41),

$$\begin{aligned} -1 + \text{diagL}(F) &\leq -1 + DF_{ii}(x) < 0 && (\text{for every } i \in \{1, \dots, n\} \text{ and } x) \\ \implies |1 - \text{diagL}(F)| &\geq |-1 + DF_{ii}(x)| \\ \implies 1 - \text{diagL}(F) &\geq \max_i |-1 + DF_{ii}(x)| \\ \implies \frac{1}{1 - \text{diagL}(F)} &\leq \frac{1}{\max_i |-1 + DF_{ii}(x)|}. \end{aligned}$$

Therefore, $\alpha \leq \frac{1}{\max_i |-1 + DF_{ii}(x)|}$ and we can use Lemma 7(i) to deduce that

$$\begin{aligned} \|(1 - \alpha)I_n + \alpha DF(x)\|_{\infty, [\eta]^{-1}} &= 1 + \alpha \mu_{\infty, [\eta]^{-1}}(-I_n + DF(x)) \\ &= 1 + \alpha(-1 + \mu_{\infty, [\eta]^{-1}}(DF(x))) && \text{for all } x \in \mathbb{R}^n \\ &\leq 1 + \alpha(-1 + \text{osL}(F)) = 1 - \alpha(1 - \text{osL}(F)) < 1. \end{aligned}$$

where the second equality follows from the translation property (22e) of matrix measures, and the inequality holds because $\mu_{\infty, [\eta]^{-1}}(DF(x)) \leq \text{osL}(F)$ for all x , and the last inequality holds because $\text{osL}(F) < 1$. This means that $\text{Lip}(F_\alpha) < 1$, for every $0 < \alpha \leq \frac{1}{1 - \text{diagL}(F)}$ and the result follows from the Banach fixed-point theorem.

Regarding part (ii), we note the contraction factor is a strictly decreasing function of α . At $\alpha = 0$ the factor is 1 and at the maximum of value of α that is, at $\alpha^* = (1 - \text{diagL}(F))^{-1}$ the contraction factor is still positive since $(1 - \text{osL}(F))/(1 - \text{diagL}(F)) \leq 1$. Hence the minimum contraction factor is achieved at α^* . \square

C.2 Proof of Theorem 3 and comparison with the literature

Before we prove Theorem 3, it is useful to compare it with similar results in the literature. The result in [Lim, 1985, Lemma 1] is more general than Theorem 3 by allowing F to be a multi-valued map defined on a metric space. However, Theorem 3(ii) uses the one-side Lipschitz constant and provides a tighter upper bound on the distance between fixed-points of F compared to its counterpart in [Lim, 1985, Lemma 1].

Proof of Theorem 3. Let $\llbracket \cdot, \cdot \rrbracket$ be a WP for the norm $\|\cdot\|_{\mathcal{X}}$ on \mathbb{R}^n .

Regarding part (i), for every $u \in \mathbb{R}^m$, we note that by definition of $\text{osL}_x(F)$, for every $u \in \mathbb{R}^r$,

$$\llbracket F(x, u) - F(y, u), x - y \rrbracket \leq \text{osL}_x(F) \|x - y\|_{\mathcal{X}}^2,$$

This implies that $\text{osL}(F_u) \leq \text{osL}_x(F) < 1$, for every $u \in \mathbb{R}^r$. Thus, by Theorem 1(iii), F_u has a unique fixed-point x_u^* .

Regarding part (ii), let $\llbracket \cdot, \cdot \rrbracket$ be a WP for the norm $\|\cdot\|_{\mathcal{X}}$ on \mathbb{R}^n and compute

$$\begin{aligned} \|x_u^* - x_v^*\|_{\mathcal{X}}^2 &= \llbracket x_u^* - x_v^*, x_u^* - x_v^* \rrbracket && (\text{by compatibility}) \\ &= \llbracket F_u(x_u^*) - F_v(x_v^*), x_u^* - x_v^* \rrbracket \\ &\leq \llbracket F_u(x_u^*) - F_u(x_v^*), x_u^* - x_v^* \rrbracket + \llbracket F_u(x_v^*) - F_v(x_v^*), x_u^* - x_v^* \rrbracket && (\text{by sub-additivity}) \\ &\leq \text{osL}_x(F) \|x_u^* - x_v^*\|_{\mathcal{X}}^2 + \|F_u(x_v^*) - F_v(x_v^*)\|_{\mathcal{X}} \|x_u^* - x_v^*\|_{\mathcal{X}} && (\text{by Cauchy-Schwarz}) \\ &\leq \text{osL}_x(F) \|x_u^* - x_v^*\|_{\mathcal{X}}^2 + \text{Lip}_u(F) \|u - v\|_{\mathcal{U}} \|x_u^* - x_v^*\|_{\mathcal{X}}. \end{aligned}$$

This implies that $(1 - \text{osL}_x(F)) \|x_u^* - x_v^*\|_{\mathcal{X}} \leq \text{Lip}_u(F) \|u - v\|_{\mathcal{U}}$ and the result of part (ii) follows. \square

C.3 Non-differentiable fixed-point problems

In many machine learning applications, the activation functions are continuous but non-differentiable and thus our results in Sections 3 do not directly apply to these problems. In this subsection, we focus on a specific form of the fixed-point equation (5), where $F = \Phi \circ H$ and $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a diagonal activation function with absolutely continuous components and $H : \mathbb{R}^n \times \mathbb{R}^r \rightarrow \mathbb{R}^n$ is a differentiable function. It can be shown that, for this class of systems, conclusions of Theorems 1, 2, and 3 still hold with respect to weighted ℓ_∞ -norms. Here, we present a result which extends Theorems 2 and 3 for $H(x, u) = G(x) + Bu$ given some $B \in \mathbb{R}^{n \times r}$ and with respect to the norm $\|\cdot\|_{\infty, [\eta]}^{-1}$.

Theorem 10 (Fixed points for non-differentiable activation functions). *Consider the norm $\|\cdot\|_{\infty, [\eta]}^{-1}$ on \mathbb{R}^n for some $\eta > 0$ and the norm $\|\cdot\|_{\mathcal{U}}$ on \mathbb{R}^r . Additionally, consider the following perturbed fixed point problem:*

$$x = \Phi(G(x) + Bu) := \Phi^G(x, u),$$

where $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a diagonal function given by $(\phi_1(x_1), \dots, \phi_n(x_n))$ with non-expansive and weakly increasing ϕ_i , $G : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a continuously differentiable function, and $B \in \mathbb{R}^{n \times r}$. Define the average map $\Phi_\alpha^G(x, u) := (1 - \alpha)x + \Phi^G(x, u)$ and pick $\text{diagL}(G)_- \in [-\text{Lip}(G), \text{osL}(G)]$ such that

$$\text{diagL}(G)_- \leq \min_i \inf_{x \in \mathbb{R}^n} DG_{ii}(x)_-.$$

Assume that $\text{osL}(G) < 1$. Then,

- (i) for every $u \in \mathbb{R}^r$, the map $\Phi^G(\cdot, u)$ has a unique fixed-point x_u^* ;
- (ii) for every $0 < \alpha \leq \frac{1}{1 - \text{diagL}(G)_-}$ and every $u \in \mathbb{R}^r$, $\Phi_\alpha^G(\cdot, u)$ is a contraction map with contraction factor $1 - \alpha(1 - \text{osL}(G)_+)$;
- (iii) for every $u, v \in \mathbb{R}^r$, we have $\|x_u^* - x_v^*\|_{\infty, [\eta]}^{-1} \leq \frac{\text{Lip}_u \Phi^G}{1 - \text{osL}(G)_+} \|u - v\|_{\mathcal{U}}$.

Proof of Theorem 10. Regarding part (i), the assumptions on each scalar activation function imply that (i) $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is non-expansive with respect to $\|\cdot\|_{\infty, [\eta]}^{-1}$ and (ii) for every $p, q \in \mathbb{R}$, there exists $\theta_i \in [0, 1]$ such that $\phi_i(p) - \phi_i(q) = \theta_i(p - q)$ or in the matrix form $\Phi(\mathbf{p}) - \Phi(\mathbf{q}) = \Theta(\mathbf{p} - \mathbf{q})$ where Θ is a diagonal matrix with diagonal elements $\theta_i \in [0, 1]$ and $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$. As a result, we have

$$\begin{aligned} \|\Phi_\alpha^G(x_1, u) - \Phi_\alpha^G(x_2, u)\|_{\infty, [\eta]}^{-1} &= \|(1 - \alpha)(x_1 - x_2) + \alpha\Theta(G(x_1) - G(x_2))\|_{\infty, [\eta]}^{-1} \\ &\leq \sup_{y \in \mathbb{R}^n} \|I_n + \alpha(-I_n + \Theta DG(y))\|_{\infty, [\eta]}^{-1} \|x_1 - x_2\|_{\infty, [\eta]}^{-1}. \end{aligned}$$

where the inequality holds by the mean value theorem. Then, for every $\alpha \in]0, \frac{1}{1 - \text{diagL}(\Theta DG)}]$,

$$\begin{aligned} \|I_n + \alpha(-I_n + \Theta DG(y))\|_{\infty, [\eta]}^{-1} &= 1 + \alpha \mu_{\infty, [\eta]}^{-1}(-I_n + \Theta DG(y)) \\ &\leq 1 + \alpha(-1 + \mu_{\infty, [\eta]}^{-1}(\Theta DG(y))) \\ &\leq 1 + \alpha(-1 + \mu_{\infty, [\eta]}^{-1}(DG(y))_+) \\ &\leq 1 - \alpha(1 - \text{osL}(G)_+) < 1, \end{aligned}$$

where the first equality holds by Lemma 7(i), the second inequality holds by subadditive property of matrix measures (22a), and the third inequality holds by Lemma 8(i). Moreover, since $\theta_i \in [0, 1]$, we have $\theta_i DG_{ii} \geq (DG_{ii})_-$, for every $i \in \{1, \dots, n\}$. This means that

$$\text{diagL}(\Theta DG) = \min_i \inf_{y \in \mathbb{R}^n} (\Theta DG(y))_{ii} \geq \min_i \inf_{y \in \mathbb{R}^n} (DG_{ii}(y))_- = \text{diagL}(G)_-.$$

This implies that, for every $\alpha \in]0, \frac{1}{1 - \text{diagL}(G)_-}]$,

$$\|\Phi_\alpha^G(x_1, u) - \Phi_\alpha^G(x_2, u)\|_{\infty, [\eta]}^{-1} \leq (1 - \alpha(1 - \text{osL}(G)_+)) \|x_1 - x_2\|_{\infty, [\eta]}^{-1}.$$

Since $1 - \alpha(1 - \text{osL}(G)_+) < 1$, the map $\Phi_\alpha^G(\cdot, u)$ is a contraction for every $\alpha \in]0, \frac{1}{1 - \text{diagL}(G)_-}]$. This concludes the proof of parts (i) and (ii),

Regarding part (iii), from formula (33) for the one-sided Lipschitz constant and formula (28) for the relevant WP, we obtain that, for all $x_1, x_2 \in \mathbb{R}^n$,

$$\begin{aligned}
& \llbracket \Phi(G(x_1) + Bu) - \Phi(G(x_2) + Bu), x_1 - x_2 \rrbracket_{\infty, [\eta]^{-1}} \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i (\phi_i((G(x_1) + Bu)_i) - \phi_i((G(x_2) + Bu)_i)) \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i ((G(x_1) + Bu)_i - (G(x_2) + Bu)_i) \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i (G(x_1) - G(x_2))_i,
\end{aligned}$$

Next, we recall Lumer's equality (30) and write it as

$$\text{osL}(G) = \sup_{x_1 \neq x_2} \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i (G(x_1) - G(x_2))_i.$$

Next, we consider two cases. Suppose that $\text{osL}(G) \leq 0$. Since $\theta_i \in [0, 1]$ for all i , we obtain

$$\llbracket \Phi(G(x_1) + Bu) - \Phi(G(x_2) + Bu), x_1 - x_2 \rrbracket_{\infty, [\eta]^{-1}} \leq 0,$$

since the maximum value is achieved at $\theta_i = 0$ for all i . Alternatively, suppose that $\text{osL}(G) > 0$. Then

$$\begin{aligned}
& \llbracket \Phi(G(x_1) + Bu) - \Phi(G(x_2) + Bu), x_1 - x_2 \rrbracket_{\infty, [\eta]^{-1}} \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i (G(x_1) - G(x_2))_i \\
&\leq \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i (G(x_1) - G(x_2))_i \leq \text{osL}(G) \|x_1 - x_2\|_{\infty, [\eta]^{-1}}^2,
\end{aligned}$$

since the maximum value is achieved at $\theta_i = 1$ for all i . This means that $\text{osL}(\Phi^G) = \text{osL}(G)_+$. Now we compute

$$\begin{aligned}
\|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}}^2 &= \llbracket x_u^* - x_v^*, x_u^* - x_v^* \rrbracket_{\infty, [\eta]^{-1}} \\
&= \llbracket \Phi_u^G(x_u^*) - \Phi_v^G(x_v^*), x_u^* - x_v^* \rrbracket_{\infty, [\eta]^{-1}} \\
&\leq \llbracket \Phi_u^G(x_u^*) - \Phi_u^G(x_v^*), x_u^* - x_v^* \rrbracket_{\infty, [\eta]^{-1}} + \llbracket \Phi_u^G(x_v^*) - \Phi_v^G(x_v^*), x_u^* - x_v^* \rrbracket_{\infty, [\eta]^{-1}} \\
&\leq \text{osL}(G)_+ \|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}}^2 + \|\Phi_u^G(x_v^*) - \Phi_v^G(x_v^*)\|_{\infty, [\eta]^{-1}} \|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}} \\
&\leq \text{osL}(G)_+ \|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}}^2 + \text{Lip}_u(\Phi^G) \|u - v\|_{\mathcal{U}} \|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}}.
\end{aligned}$$

This implies that $(1 - \text{osL}(G)_+) \|x_u^* - x_v^*\|_{\infty, [\eta]^{-1}} \leq \text{Lip}_u(\Phi^G) \|u - v\|_{\mathcal{U}}$ and the result follows. \square

C.4 Proofs of results in Section 4

Proof of Theorem 4. The assumptions on each scalar activation function imply that (i) $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is non-expansive with respect to $\|\cdot\|_{\infty, [\eta]^{-1}}$, and (ii) for every $p, q \in \mathbb{R}$, there exists $\theta_i \in [0, 1]$ such that $\phi_i(p) - \phi_i(q) = \theta_i(p - q)$. Regarding the equality $\text{osL}_x(\mathbf{N}) = \mu_{\infty, [\eta]^{-1}}(A)_+$, from formula (33) for the one-sided Lipschitz constant and formula (28) for the relevant WP, we obtain that, for all $x_1, x_2 \in \mathbb{R}^n$,

$$\begin{aligned}
& \llbracket \Phi(Ax_1 + Bu) - \Phi(Ax_2 + Bu), x_1 - x_2 \rrbracket_{\infty, [\eta]^{-1}} \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i (\phi_i((Ax_1 + Bu)_i) - \phi_i((Ax_2 + Bu)_i)) \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i ((Ax_1 + Bu)_i - (Ax_2 + Bu)_i) \\
&= \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i (Ax_1 - Ax_2)_i,
\end{aligned}$$

Next, we recall Lumer's equality (30) and write it as

$$\mu_{\infty, [\eta]^{-1}}(A) = \sup_{x_1 \neq x_2} \max_{i \in I_\infty([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i ((Ax_1)_i - (Ax_2)_i).$$

Next, we consider two cases. Suppose that $\mu_{\infty, [\eta]^{-1}}(A) \leq 0$. Since $\theta_i \in [0, 1]$ for all i , we obtain

$$\|\Phi(Ax_1 + Bu) - \Phi(Ax_2 + Bu), x_1 - x_2\|_{\infty, [\eta]^{-1}} \leq 0,$$

since the maximum value is achieved at $\theta_i = 0$ for all i . Alternatively, suppose that $\mu_{\infty, [\eta]^{-1}}(A) > 0$. Then

$$\begin{aligned} & \|\Phi(Ax_1 + Bu) - \Phi(Ax_2 + Bu), x_1 - x_2\|_{\infty, [\eta]^{-1}} \\ &= \max_{i \in I_{\infty}([\eta]^{-1}(x_1 - x_2))} \theta_i \eta_i^{-2} (x_1 - x_2)_i (Ax_1 - Ax_2)_i \\ &\leq \max_{i \in I_{\infty}([\eta]^{-1}(x_1 - x_2))} \eta_i^{-2} (x_1 - x_2)_i (Ax_1 - Ax_2)_i \leq \mu_{\infty, [\eta]^{-1}}(A) \|x_1 - x_2\|_{\infty, [\eta]^{-1}}^2, \end{aligned}$$

since the maximum value is achieved at $\theta_i = 1$ for all i . This concludes the proof of formula $\text{osL}_x(\mathbf{N}) = \mu_{\infty, [\eta]^{-1}}(A)_+$. Next, since Φ is non-expansive, we compute

$$\begin{aligned} \|\mathbf{N}(x_1, u) - \mathbf{N}(x_2, u)\|_{\infty, [\eta]^{-1}} &= \|\Phi(Ax_1 + Bu) - \Phi(Ax_2 + Bu)\|_{\infty, [\eta]^{-1}} \\ &\leq \|(Ax_1 + Bu) - (Ax_2 + Bu)\|_{\infty, [\eta]^{-1}} \\ &\leq \|A(x_1 - x_2)\|_{\infty, [\eta]^{-1}} \leq \|A\|_{\infty, [\eta]^{-1}} \|x_1 - x_2\|_{\infty, [\eta]^{-1}}, \end{aligned}$$

proving the formula $\text{Lip}_x(\mathbf{N}) = \|A\|_{\infty, [\eta]^{-1}}$. The proof of the formula $\text{Lip}_u(\mathbf{N}) = \|B\|_{(\infty, [\eta]^{-1}), \mathcal{U}}$ is essentially identical. Finally, if each ϕ_i is differentiable then we compute

$$\begin{aligned} \text{diagL}(\mathbf{N}) &= \min_{i \in \{1, \dots, n\}} \inf_{x \in \mathbb{R}^n, u \in \mathbb{R}^r} D\mathbf{N}_{ii}(x, u) = \min_{i \in \{1, \dots, n\}} \inf_{x \in \mathbb{R}^n, u \in \mathbb{R}^r} \phi'_i((Ax + Bu)_i) A_{ii} \\ &\leq \min_{i \in \{1, \dots, n\}} \begin{cases} 0, & \text{if } A_{ii} > 0 \\ A_{ii}, & \text{if } A_{ii} \leq 0 \end{cases} = \min_{i \in \{1, \dots, n\}} (A_{ii})_-, \end{aligned} \quad (42)$$

because of the properties of the activation functions. Now suppose that there exists $i \in \{1, \dots, n\}$ such that ϕ_i is not differentiable. Using Theorem 10(ii) with $\mathbf{G} = A$, $\text{diagL}(\mathbf{N})$ is chosen to be equal to be $\text{diagL}(A)_-$ which in turn is equal to $\min_{i \in \{1, \dots, n\}} (A_{ii})_-$. \square

Proof of Corollary 5. The results are immediate consequences of Theorem 2 (or more generally Theorem 10 for non-differentiable activation functions) and of the Lipschitz estimates in Theorem 4. \square

D Adversarial attacks on implicit neural networks

In this appendix, we study the effect of different adversarial attacks on the existing implicit network models as well as to the NEMON model.

D.1 Attack models

First, we review several attack models that are used in the literature to study the input-output resilience of neural networks. Each attack consists of a model for generating suitable perturbations of the test input data. Perturbations with respect to these attacks were generated using the Foolbox software package⁵.

Continuous image inversion. The continuous image inversion attack is defined by:

$$U_{\text{adversarial}} = U + \varepsilon \text{sign} \left(\frac{1}{2} \mathbf{1}_r \mathbf{1}_m^\top - U \right). \quad (43)$$

It is clear that this attack is independent of the neural network model. Plots of perturbed MNIST images under the continuous image inversion attack are shown in Figure 6. In Figure 2, the right plot compares the accuracy of the NEMON model, the implicit deep learning model [El Ghaoui et al., 2021], and the MON model [Winston and Kolter, 2020] for $\varepsilon \in [0.0, 5]$.

Uniform additive ℓ_∞ -noise. For this attack, the test images are perturbed by an additive noise with ℓ_∞ magnitude sampled uniformly from the interval $[0, 1]$. Plots of perturbed MNIST images under uniform additive ℓ_∞ -noise are shown in Figure 6. Figure 7 shows scatter plots of the accuracy of the NEMON model, the implicit deep learning model, and the MON model over 1000 sample attacks.

⁵The Foolbox implementation is licensed under the MIT License and is available at <https://github.com/bethgelab/foolbox>.




Figure 5: Images of MNIST handwritten digits perturbed by the continuous image inversion attack. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude $\varepsilon = 0.1$, etc.

Fast gradient sign method. Given input data $U \in \mathbb{R}^{r \times m}$ and output labels $Y \in \mathbb{R}^{q \times m}$, the fast gradient sign method (FGSM) generates adversarial inputs via the formula

$$U_{\text{adversarial}} = U + \varepsilon \text{sign} \left(\frac{\partial \mathcal{L}}{\partial U}(Y, CX + DU) \right), \quad (44)$$

where \mathcal{L} is the loss function used to train the network and ε provides the ℓ_∞ amplitude of the perturbation. Plots of perturbed MNIST images under the FGSM are shown in Figure 8. Plots of accuracy versus ℓ_∞ perturbation under the FGSM are shown in Figure 9.

Projected gradient descent method. The projected gradient descent method (PGDM) can be thought of as perturbing the input with several steps of the FGSM. The PGDM attack can be defined for any norm, but for consistency, we reproduce it only for the ℓ_∞ -norm. For the input data $U \in \mathbb{R}^{r \times m}$ and outputs $Y \in \mathbb{R}^{q \times m}$, PGDM defines the finite sequence of perturbations $\{\delta_k\}_{k=1}^M$ by

$$\delta_{k+1} = \text{Proj}_{\overline{B(\varepsilon)}} \left(\delta_k + \alpha \text{sign} \left(\frac{\partial \mathcal{L}}{\partial U}(Y, CX + D(U + \delta_k)) \right) \right), \quad (45)$$

where M is some prescribed maximum number of steps, α is a stepsize, and $\text{Proj}_{\overline{B(\varepsilon)}}$ is the ℓ_2 orthogonal projection operator onto the entrywise ℓ_∞ closed ball with radius ε . This projection operator corresponds to clipping each entry of the matrix so that it is in the range $[-\varepsilon, \varepsilon]$. Then, the perturbed input is simply

$$U_{\text{adversarial}} = U + \delta_M.$$

Plots of perturbed MNIST images under the PGDM are shown in Figure 10. Plots of accuracy versus ℓ_∞ perturbation under the PGDM are shown in Figure 11.

D.2 Other methods to decrease the ℓ_∞ Lipschitz constant

Recall that the input-output Lipschitz constant of the model (11) with both $\|\cdot\|_{\mathcal{U}}$ and $\|\cdot\|_{\mathcal{Y}}$ equal to the ℓ_∞ -norm is given by

$$\text{Lip}_{u \rightarrow y} = \frac{\|B\|_{(\infty, [\eta]^{-1}), (\infty)} \|C\|_{(\infty), (\infty, [\eta]^{-1})}}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+} + \|D\|_{\infty, \infty}.$$




Figure 6: Images of MNIST handwritten digits as perturbed by uniform additive ℓ_∞ noise. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.2 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude $\varepsilon = 0.2$, etc.

When input data, U , is perturbed, the perturbation is directly fed into the output Y via the output equation $Y = CX + DU$. For this reason, a simple change to attempt to minimize the effect of input perturbations on the output is to replace the DU term in the output equation by a static bias, i.e.,

$$Y = CX + b\mathbb{1}_m^\top,$$

where $b \in \mathbb{R}^q$. This simple modification to the model changes the input-output Lipschitz constant to

$$\text{Lip}_{u \rightarrow y} = \frac{\|B\|_{(\infty, [\eta]^{-1}), (\infty)} \|C\|_{(\infty), (\infty, [\eta]^{-1})}}{1 - \mu_{\infty, [\eta]^{-1}}(A)_+}.$$

Finally, another degree of freedom is the parameter $\gamma < 1$ in the constraint $\mu_{\infty, [\eta]^{-1}}(A) \leq \gamma$. In all previously shown experiments on MNIST, we selected $\gamma = 0.95$. From the expression for the input-output Lipschitz constant of the network (15), $\mu_{\infty, [\eta]^{-1}}(A) = 0.95$ leads to a small denominator, resulting in a relatively large input-output Lipschitz constant. A simple modification to moderate the Lipschitz constant is to impose $\mu_{\infty, [\eta]^{-1}}(A) \leq \epsilon$ for some small $\epsilon \geq 0$. This attempts to maximize the denominator in the expression for the Lipschitz constant.

For these modifications to the models, plots of accuracy versus ℓ_∞ -perturbation generated by the FGSM are shown in Figure 12. In this figure, we set $\epsilon = 0.05$ for the NEMON models. For comparison, the well-posedness condition for MON is set to be $\mu_2(A) \leq \epsilon$. We do not modify the condition $\|A\|_\infty \leq 0.95$ as imposing the constraint $\|A\|_\infty \leq \epsilon$ is overly restrictive and would result in a significant drop in accuracy.

D.3 Robustness of implicit neural networks on the MNIST dataset

In this section, we compare the performance of the NEMON model with $\mu_\infty(A) \leq 0.95$ to the implicit deep learning model with $\|A\|_\infty \leq 0.95$ and to the monotone operator equilibrium network (MON) with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$ with respect to the attacks described in the previous subsection on the MNIST dataset.

For the continuous image inversion attack, Figure 2 shows the curves for accuracy versus ℓ_∞ -amplitude of the perturbation. We observe that, compared to the NEMON model, the implicit deep




Figure 7: Scatter plots of accuracy versus ℓ_∞ perturbation as generated by uniform additive ℓ_∞ noise over 1000 trials. Plots are shown for the NEMON model $\mu_\infty(A) \leq 0.95$ with $\lambda \in \{10^{-1}, 10^{-2.5}, 10^{-5}, 0\}$, the implicit deep learning model $\|A\|_\infty \leq 0.95$, and the monotone operator equilibrium network (MON) with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$.

learning model and the monotone operator equilibrium network (MON) have larger drops in accuracy for small perturbations. For the NEMON model, as λ increases, the accuracy at zero perturbation decreases. However, as λ increases, the overall robustness of NEMON improves as its accuracy does not decrease substantially even for large amplitudes of perturbation.

For uniform additive ℓ_∞ -noise, scatter plots with accuracy versus ℓ_∞ amplitude of the perturbation are shown in Figure 7. We see that the NEMON model with $\lambda = 0$, the implicit deep learning model, and the MON model all perform comparably. The NEMON models with $\lambda = 10^{-1}$ and $\lambda = 10^{-2.5}$ both see improved robustness as their accuracy does not drop as noticeably with ℓ_∞ amplitude of the perturbation. Surprisingly, the NEMON model with $\lambda = 10^{-5}$ seems to be less robust than the NEMON model with $\lambda = 0$.

For the FGSM, Figure 9 shows the curves for accuracy versus ℓ_∞ amplitude of the perturbation. We see that the NEMON models with $\lambda = 10^{-5}$ and $\lambda = 10^{-4}$ are the least robust, followed by the NEMON model with $\lambda = 0$ and the MON. Only for $\lambda \in \{10^{-2.5}, 10^{-2}, 10^{-1}\}$ do we see an




Figure 8: Images of MNIST handwritten digits as perturbed by the FGSM. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude $\varepsilon = 0.1$, etc.

improvement in robustness for the NEMON model at the price of a decrease in nominal accuracy. Note that for the FGSM, each model experiences different perturbations.

For the PGDM, Figure 11 shows the curves for accuracy versus ℓ_∞ amplitude of the perturbation. We see that the results are comparable with the perturbation generated by the FGSM, with the exception that the implicit deep learning model now performs comparably with the monotone operator equilibrium model. Note that for the PGDM, each model experiences different perturbations.

Finally, we compare the performance of the models with the modification that the output equation is $Y = CX + b\mathbb{1}_m^\top$. Figure 12 shows the curves for accuracy versus ℓ_∞ amplitude of the FGSM perturbation for the NEMON model with $\mu_\infty(A) \leq 0.05$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and the monotone operator equilibrium model with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$. For these modifications in the models, we see improvement in overall accuracy compared to original models of implicit networks (11) shown in Figure 9. Additionally, we observe comparable performance in the NEMON model with $\lambda = 0$ and the implicit deep learning model, with the MON performing slightly better than both. For the NEMON model with $\lambda = 10^{-4}$, the accuracy at zero perturbation is comparable to the NEMON model with $\lambda = 0$ and the overall robustness of the NEMON model to the FGSM attack is significantly improved. However, as λ increases, we see that the nominal accuracy and overall robustness of the NEMON models deteriorate.

D.4 Robustness of implicit neural networks on the CIFAR-10 dataset

In this section, we compare the performance of the NEMON model with $\mu_\infty(A) \leq 0$ to MON with $I_n - \frac{1}{2}(A + A^\top) \succeq I_n$ with respect to the FGSM attack described in the previous subsection on the CIFAR-10 dataset.

For the FGSM attack on the CIFAR-10 dataset, Figure 13 shows the accuracy versus the ℓ_∞ amplitude of the perturbation for the regularized and un-regularized NEMON model and the MON model. We observe that un-regularized NEMON model is more accurate than MON for all amplitudes of perturbation. For example, at ℓ_∞ -perturbation equal to 0.1, the accuracy of un-regularized NEMON is 39% whereas the accuracy of MON at this attack amplitude is 35%. Moreover, the regularized NEMON with the regularization parameter $\lambda = 10^{-4}$ has a clean performance accuracy




Figure 9: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0.95$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$.

of 66% which is lower than the clean accuracy of both MON and the un-regularized NEMON. However, the regularized NEMON demonstrates a consistent improvement in accuracy for sizeable ℓ_∞ -perturbations. For example, at an ℓ_∞ -perturbation equal to 0.15, the accuracy of the regularized NEMON model is 29% whereas the accuracy of MON at this attack amplitude is 24%.




Figure 10: Images of MNIST handwritten digits as perturbed by the PGDM. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude $\varepsilon = 0.1$, etc.




Figure 11: Plot of accuracy versus ℓ_∞ perturbation as generated by the PGDM for the NEMON model with $\mu_\infty(A) \leq 0.95$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^T) \succeq 0.05I_n$.




Figure 12: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0.05$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$. The output equation is $Y = CX + b\mathbb{1}_m^\top$.




Figure 13: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0$ and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq I_n$.