

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 $[[\cdot, \cdot]] : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying:

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

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

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

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 $\eta \in \mathbb{R}_{>0}^n$,

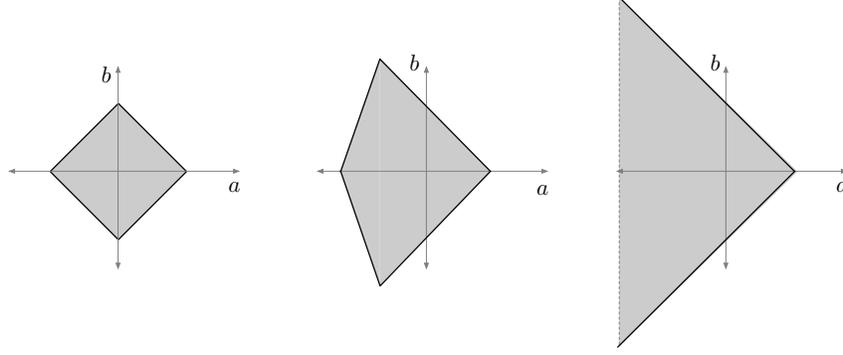


Figure 4: The left figure shows the region $\text{Lip}(A) \leq 1$, the middle figure shows the region $\text{Lip}(A_\alpha) \leq 1$ for $\alpha = \frac{1}{2}$, and the right figure shows $\text{osL}(A) \leq 1$. Both Lip and osL are with respect to the ℓ_∞ -norm

(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

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

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

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

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

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 $\text{osL}(F) \leq 1 - c$, 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

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

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

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

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 $x_1, x_2 \in \mathbb{R}^n$,

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

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

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

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

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

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

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

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

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

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

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

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 $G = A$, $\text{diagL}(\mathbf{N})$ is chosen to be equal to $\text{diagL}(A)_-$ 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 $\varepsilon \in [0.0, 0.5]$.

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>.

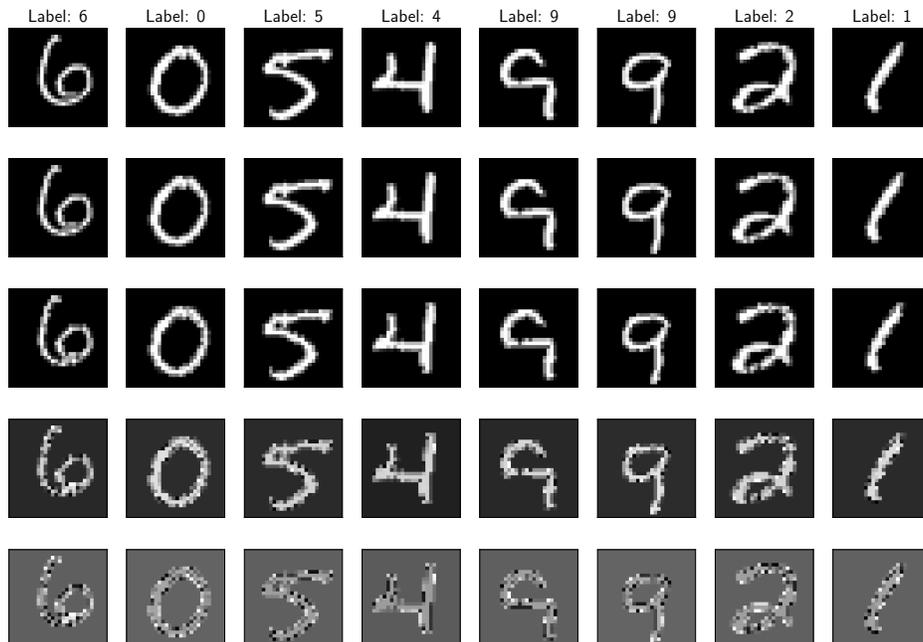


Figure 5: Images of MNIST handwritten digits perturbed by the continuous image inversion attack. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed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

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

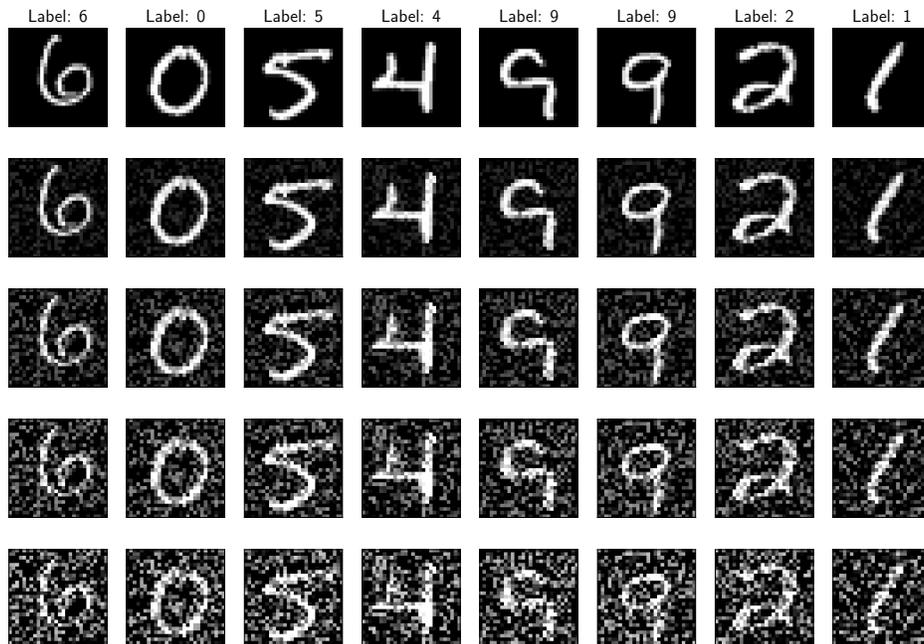


Figure 6: Images of MNIST handwritten digits as perturbed by uniform additive ℓ_∞ noise. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.2 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed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

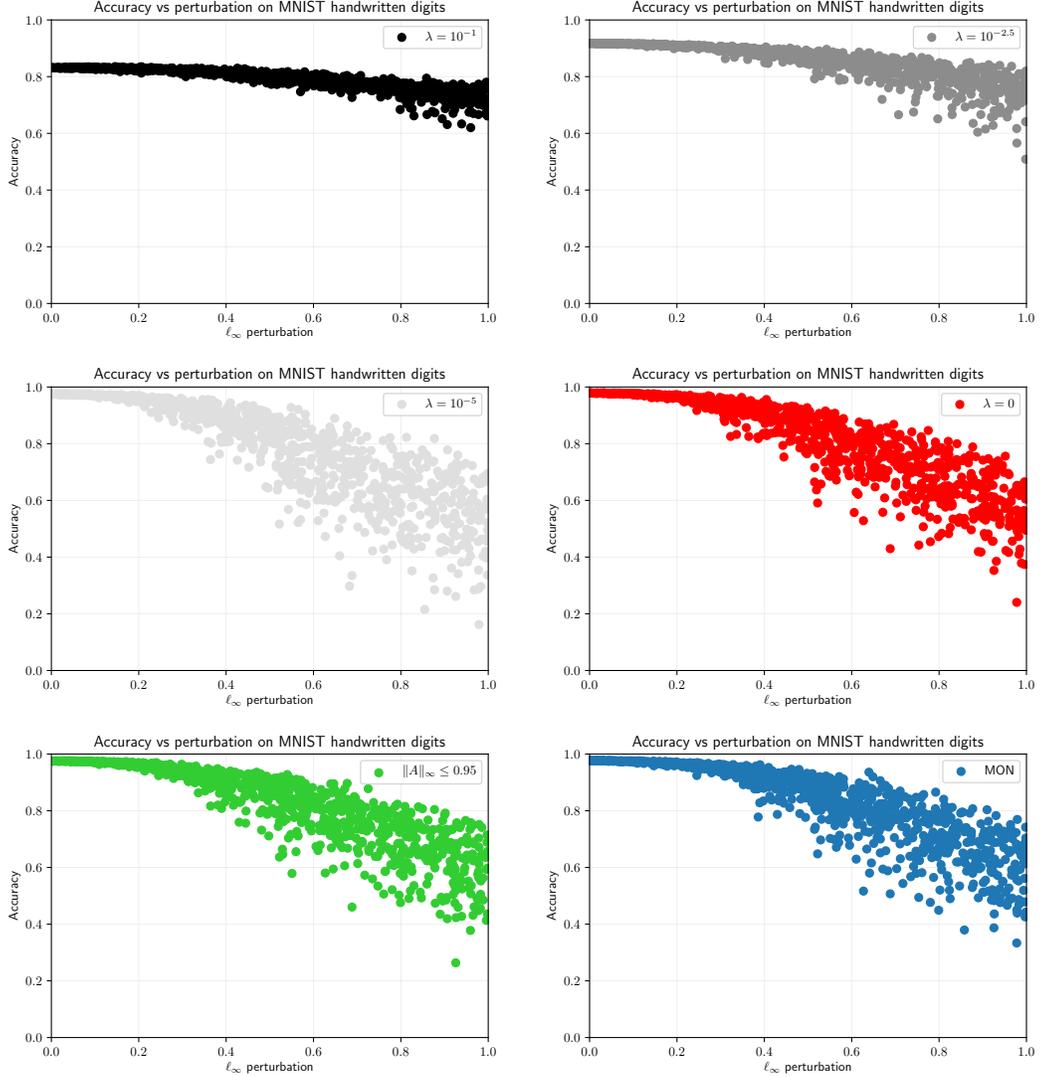


Figure 7: Scatter plots of accuracy versus ℓ_∞ perturbation as generated by uniform additive ℓ_∞ noise over 1000 trials. Plots are shown for the NEMON model $\mu_\infty(A) \leq 0.95$ with $\lambda \in \{10^{-1}, 10^{-2.5}, 10^{-5}, 0\}$, the implicit deep learning model $\|A\|_\infty \leq 0.95$, and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

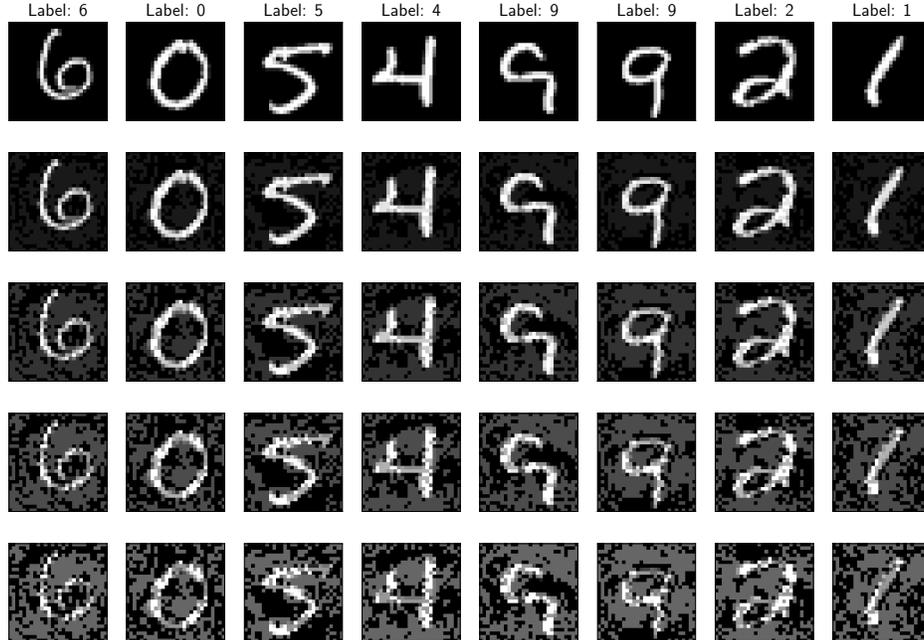


Figure 8: Images of MNIST handwritten digits as perturbed by the FGSM. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude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

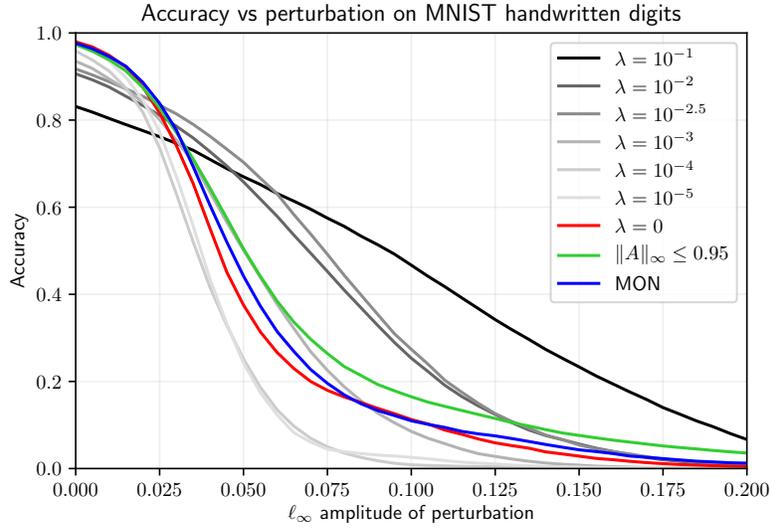


Figure 9: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0.95$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^T) \succeq 0.05I_n$.

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%.

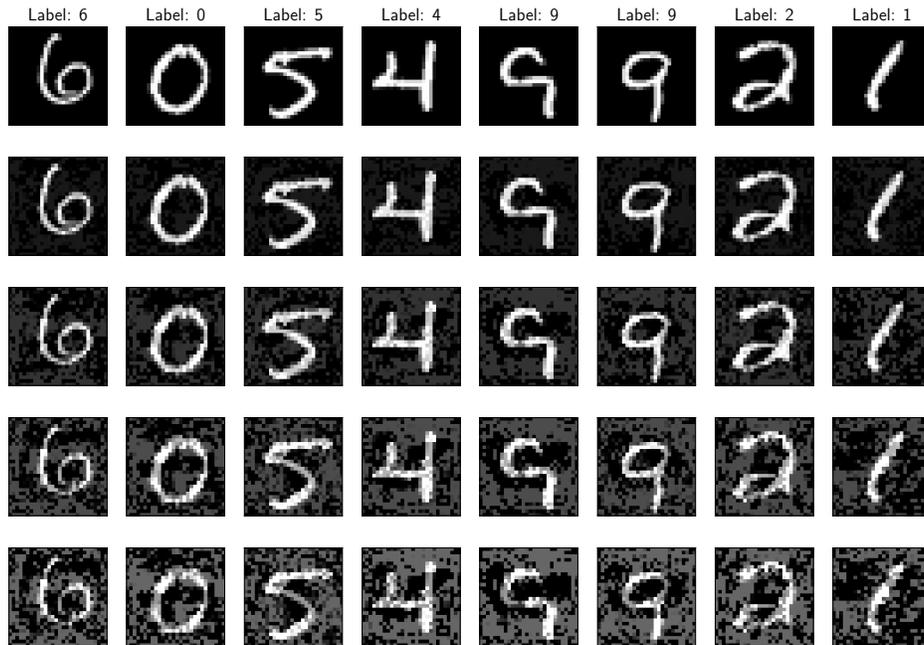


Figure 10: Images of MNIST handwritten digits as perturbed by the PGDM. For $i \in \{1, \dots, 5\}$, row i corresponds to an ℓ_∞ perturbation amplitude $\varepsilon = 0.1 \times (i - 1)$. In other words, the top row has unperturbed images, the second row has images that is perturbed by an ℓ_∞ amplitude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

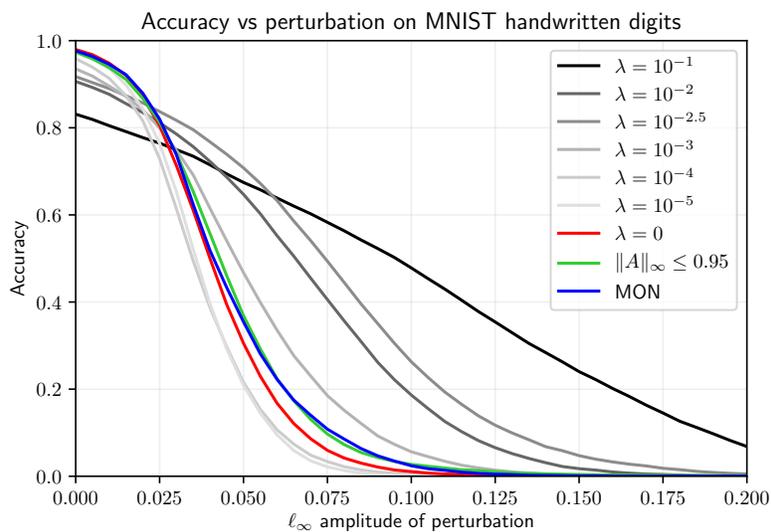


Figure 11: Plot of accuracy versus ℓ_∞ perturbation as generated by the PGDM for the NEMON model with $\mu_\infty(A) \leq 0.95$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^T) \succeq 0.05I_n$.

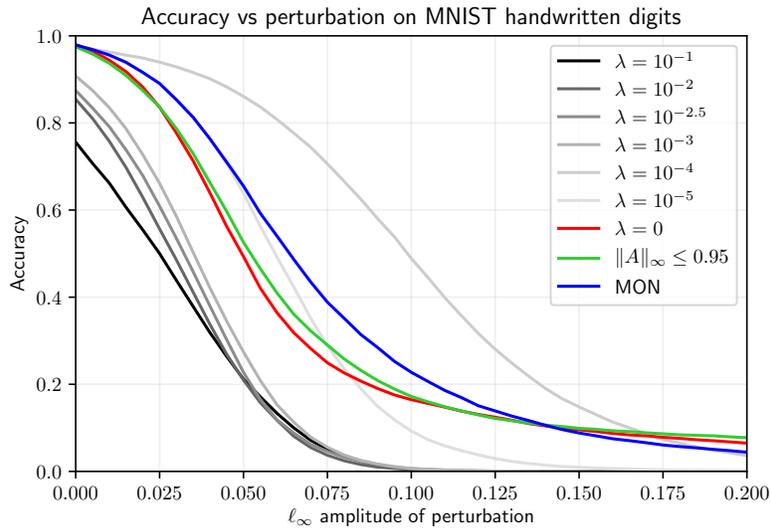


Figure 12: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0.05$, the implicit deep learning model with $\|A\|_\infty \leq 0.95$, and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq 0.05I_n$. The output equation is $Y = CX + b\mathbf{1}_m^\top$.

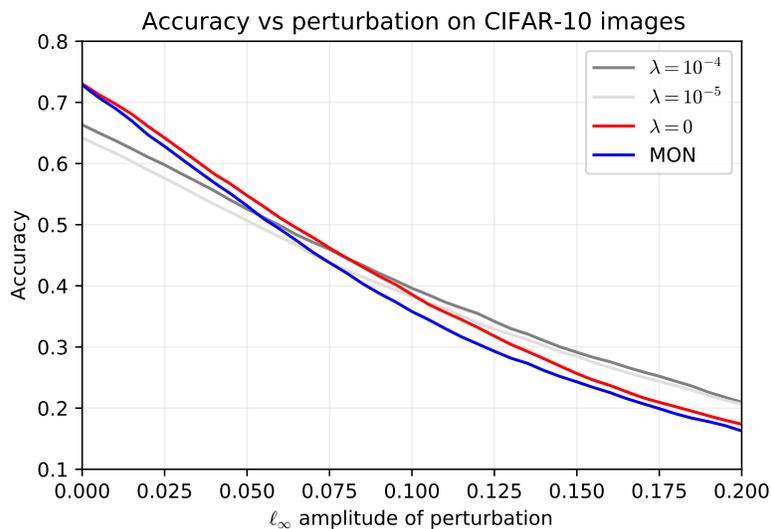


Figure 13: Plot of accuracy versus ℓ_∞ perturbation as generated by the FGSM for the NEMON model with $\mu_\infty(A) \leq 0$ and MON with $I_n - \frac{1}{2}(A + A^\top) \succeq I_n$.