

Optimal lower Lipschitz bounds for ReLU layers, saturation, and phase retrieval

Daniel Freeman^a, Daniel Haider^b

^a*Department of Mathematics and Statistics, St. Louis University, Ritter Hall 319, St. Louis, 63103, MO, USA*

^b*Acoustics Research Institute, Austrian Academy of Sciences, Dominikanerbastei 16, Vienna, 1010, AUT*

Abstract

The injectivity of ReLU layers in neural networks, the recovery of vectors from clipped or saturated measurements, and (real) phase retrieval in \mathbb{R}^n allow for a similar problem formulation and characterization using frame theory. In this paper, we revisit all three problems with a unified perspective and derive lower Lipschitz bounds for ReLU layers and clipping which are analogous to the previously known result for phase retrieval and are optimal up to a constant factor.

Keywords: Lower Lipschitz stability, ReLU, saturation, clipping, phase retrieval

1. Introduction

Many important functions in applications arise as a linear operator composed with some simple non-linear function. In engineering, the non-linear component often comes from unwanted technical constraints, such as a limited dynamic range in measurement devices that causes clipping or saturation effects. For example, if the amplitude of an audio signal exceeds the threshold of the recording equipment at any point then the device will often record the resulting waveform as having its top and bottom clipped at the threshold. In contrast to this, there are circumstances where non-linearities are intentionally incorporated into model designs to enhance their expressiveness. This is extensively made use of in neural networks, where linear maps are alternately concatenated with non-linear activation functions to build large and

powerful layer-structured models. In both situations, it is fundamentally important to know whether or not it is possible to recover the original input from the resulting non-linear measurements. In the unintentional case, we want to know how to effectively design measurement devices which are able to compensate for the resulting loss of information. In the intended case, we may want to balance the effect of the non-linearity and the information flow, i.e., get the desired (compression or sparsification) effect *and* invertibility which has been leveraged to extend the usage of a model and enhance its interpretability [2, 18]. In other words, we want to know under which conditions the associated *non-linear measurement operator* is one-to-one on a certain domain of interest.

The maps that we consider here are composed of a linear operator from \mathbb{R}^n to \mathbb{R}^m , followed by a non-linear function that is applied component-wise to the resulting coefficients. Combining engineering and machine learning terminology we will refer to the respective computational steps as *measurement* and *activation*. Let $(\phi_i)_{i=1}^m$ be a collection of measurement vectors in \mathbb{R}^n with $m \geq n$. We say that $(\phi_i)_{i=1}^m$ is a frame for \mathbb{R}^n if the (linear) measurement map $x \mapsto (\langle x, \phi_i \rangle)_{i=1}^m$ is one-to-one from \mathbb{R}^n to \mathbb{R}^m . Let further $\rho : \mathbb{R} \rightarrow \mathbb{R}$ be a non-linear activation function, and $X \subseteq \mathbb{R}^n$ an input domain of choice. The resulting non-linear measurement operator is of the form

$$\begin{aligned} \Phi : X &\rightarrow \mathbb{R}^m \\ x &\mapsto (\rho(\langle x, \phi_i \rangle))_{i=1}^m. \end{aligned} \tag{1}$$

Of course, studying the one-to-one property of Φ is only interesting if the non-linear function ρ is not one-to-one itself. We focus on three examples of activation functions where ρ is not one-to-one, but for which it is possible for the non-linear measurement operator Φ to be one-to-one on X .

A. ReLU layers. Originally introduced to regularize the gradients in deep network architectures during training, these maps have established themselves as a powerful and easy-to-use ingredient for the design of neural networks [14, 19]. A ReLU layer can be written as a non-linear measurement operator according to (1) when setting the activation function ρ to be

$$\text{ReLU}(t) = \max(0, t). \tag{2}$$

Usually, a bias value $\alpha_i \in \mathbb{R}$ is assigned to every measurement vector ϕ_i . The map then becomes $x \mapsto (\text{ReLU}(\langle x, \phi_i \rangle - \alpha_i))_{i=1}^m$. The vector $\alpha \in \mathbb{R}^m$




Figure 1: The plots show the non-linear activation functions used in A.-C.: The ReLU function, the saturation function with $\lambda = 2$, and the magnitude (intensity) function.

comprising the bias values is called the *bias vector*, which acts as a threshold mechanism for the activation. As input domain, we consider $X = \mathbb{R}^n$ [6, 16, 22].

B. λ -Saturation or clipping. In this setting, it is assumed that measurements can only be recorded accurately up to a certain magnitude threshold level $\lambda > 0$ and the maximum $\pm\lambda$ is returned whenever it is exceeded. This is a common occurrence in signal processing and electrical engineering when using a measurement device with a limited dynamic range. The corresponding activation function is the saturation function given by

$$\sigma_\lambda(t) = \text{sign}(t) \cdot \min(|t|, \lambda) \quad (3)$$

Since all the measurements for vectors with very large norm would be saturated, we consider inputs only from a closed ball in \mathbb{R}^n [1, 13, 12, 17]. Using a scaling argument, we may without loss of generality restrict the domain X to be the unit ball $\mathbb{B}_{\mathbb{R}^n} = \{x \in \mathbb{R}^n : \|x\| \leq 1\}$.

C. Phase retrieval. In applications such as electron microscopy and coherent-diffraction imaging, researchers often work with devices which can only record the magnitude of the measurement values, and the phase information is lost. This map is sometimes called the intensity measurement operator [3]. Accordingly, the activation function is given by $t \mapsto |t|$. Since, no matter what, we are only able to reconstruct the input up to a global sign, we consider the quotient space \mathbb{R}^n/\sim as input domain, where the equivalence relation \sim is given by $x \sim y$ if $x = \lambda y$ for some scalar $|\lambda| = 1$.

For all three cases, the injectivity characterizations of the associated non-linear measurement operators have a similar frame theoretic nature [1, 3, 6, 16, 22]. To ensure numerical stability of the recovery maps in applications, the non-linear measurement operators must not only be one-to-one but also be bi-Lipschitz embeddings. Recall that a map $f : X \rightarrow Y$ between metric spaces (X, d_X) and (Y, d_Y) is called *bi-Lipschitz* if there are constants $\kappa_L, \kappa_U > 0$ such that

$$\kappa_L d_X(x, y) \leq d_Y(f(x), f(y)) \leq \kappa_U d_X(x, y) \quad \text{for all } x, y \in X. \quad (4)$$

We call κ_L a lower Lipschitz bound for f and κ_U an upper Lipschitz bound for f . A bound is called *optimal* if the corresponding inequality is sharp. If the measurement vectors $(\phi_i)_{i=1}^m$ form a frame for \mathbb{R}^n then the linear map $x \mapsto (\langle x, \phi_i \rangle)_{i=1}^m$ is a bi-Lipschitz embedding from \mathbb{R}^n into \mathbb{R}^m and the optimal bi-Lipschitz constants correspond exactly to the linear map's smallest and largest singular values. The optimal *frame bounds* of $(\phi_i)_{i=1}^m$ are the greatest positive value A and least value B so that $A\|x\|^2 \leq \sum_{i=1}^m |\langle x, \phi_i \rangle|^2 \leq B\|x\|^2$ holds for all $x \in \mathbb{R}^n$. Thus, A and B are the optimal frame bounds of $(\phi_i)_{i=1}^m$ if and only if \sqrt{A} and \sqrt{B} are the optimal bi-Lipschitz bounds for the linear measurement map $x \mapsto (\langle x, \phi_i \rangle)_{i=1}^m$. We choose to use the terminology of frame bounds, although some papers on the topic express the corresponding inequalities in terms of singular values [4, 22].

All three of the activation functions that we consider are non-expansive mappings, i.e. they have an upper Lipschitz bound of 1. This simple property is very useful when working with non-linear operators, and it is directly exploited in [10] to solve a general class of inverse problems. For us, the fact that the activation functions are 1-Lipschitz give that if B is the upper frame bound of $(\phi_i)_{i=1}^m$, then the optimal upper Lipschitz bound for the non-linear measurement operator Φ always satisfies $\kappa_U \leq \sqrt{B}$. Obtaining lower Lipschitz bounds for Φ is much more difficult, and this will be the main focus of our paper. The following is our main result and provides the first dimension independent lower Lipschitz bounds of κ_L for ReLU layers and λ -saturation.

Theorem 1 (Optimal lower Lipschitz bounds).

Let $(\phi_i)_{i \in I}$ with $|I| = m$ be a frame for \mathbb{R}^n , and let ρ be the activation function for either a ReLU layer with bias α or λ -saturation. Let $X \subseteq \mathbb{R}^n$ and let $D \subseteq \mathbb{R}$ be the closed interval where ρ is one-to-one. For each $x \in X$ let $I_\rho(x) = \{i \in I : \rho(\langle x, \phi_i \rangle) \in D\}$ and let $A(x)$ be the lower frame bound of $(\phi_i)_{i \in I_\rho(x)}$. We set $A_\rho = \min_{x \in X} A(x)$.

(i) If ρ is the activation function for an injective ReLU layer with bias α and κ_L is the optimal lower Lipschitz bound for the map $x \mapsto (\rho(\langle x, \phi_i \rangle))_{i \in I}$ from $X = \mathbb{R}^n$ to \mathbb{R}^m then

$$\frac{1}{2}\sqrt{A_\rho} \leq \kappa_L \leq \sqrt{A_\rho}. \quad (5)$$

(ii) If ρ is the activation function for an injective saturated measurement operator with saturation level $\lambda > 0$ and κ_L is the optimal lower Lipschitz bound for the map $x \mapsto (\rho(\langle x, \phi_i \rangle))_{i \in I}$ from $X = \mathbb{B}_{\mathbb{R}^n}$ to \mathbb{R}^m then

$$\min \left\{ \frac{1}{2}\sqrt{A_\rho}, \lambda \right\} \leq \kappa_L \leq \sqrt{A_\rho}. \quad (6)$$

In the case of ReLU layers with $\alpha = \mathbf{0}$, bounds for κ_L in terms of frame bounds of subcollections of $(\phi_i)_{i \in I}$ was first given in [6]. However, the provided lower bound could be 0 in some circumstances, even when the map was bi-Lipschitz. This was rectified in [22] where it was proven that $\sqrt{\frac{A_\rho}{2m}} \leq \kappa_L$, thus providing the first lower bound on κ_L which is non-zero if and only if κ_L is non-zero. The necessity of m in this inequality was investigated in [23], where it was conjectured that one may replace m with 2. Our bound in (5) verifies this conjecture and provides the first dimension independent lower bound on κ_L for ReLU layers.

For λ -saturation, it was known that if λ is greater than a critical threshold then the non-linear measurement operator is one-to-one if and only if it is bi-Lipschitz [1]. However, lower bounds for κ_L were only known for certain random frames [12, 13, 17], and whether or not the saturated measurement operator is bi-Lipschitz at the critical threshold was listed as an open problem in [1]. Our bound in (6) gives the first explicit lower bound on κ_L , and consequently solves the open problem about the critical threshold.

The paper is structured as following. In Section 2, we review the known injectivity characterizations for ReLU layers, saturation, and phase retrieval, and put them into direct relation as non-linear measurement operators. This sets the stage for studying other settings that allow this formulation in a similar manner, e.g., gating [11], or modulo sampling [5]. In Section 3, we prove Theorem 1 on the optimal lower Lipschitz bounds for ReLU layers and saturation (Corollary 3.2 and 3.6). We show as well how our estimates correspond nicely to the known lower Lipschitz bounds for phase retrieval [4]. Notably, even though we obtain directly analogous lower Lipschitz bounds for all three activation functions, our proof for ReLU and λ -saturation is very different from the previously known proof for phase retrieval.

2. Injectivity of non-linear measurements

For the derivation of the optimal lower Lipschitz bounds for ReLU layers and saturated measurements we make explicit use of the injectivity characterizations of the associated non-linear measurement operators given in [1, 16, 22]. In a nutshell, we obtain injectivity if and only if for any input there are sufficiently many coordinates in the output that are not affected by the activation. In this section, we recall these results as a preparation and find that also the injectivity characterization for phase retrieval can be formulated in a similar manner. This emphasizes the proposed unified perspective as non-linear measurement operators.

2.1. Frames and linear measurements

We will use standard language and tools from frame theory, and we recommend [8, 9] as references. A collection of vectors $(\phi_i)_{i \in I}$ in \mathbb{R}^n with $|I| = m \geq n$ is called a *frame* for \mathbb{R}^n if there exist constants $0 < A \leq B < \infty$ with

$$A\|x\|^2 \leq \sum_{i \in I} |\langle x, \phi_i \rangle|^2 \leq B\|x\|^2 \quad \text{for all } x \in \mathbb{R}^n. \quad (7)$$

The vectors ϕ_i are called *frame elements* and the inner products, or measurements $\langle x, \phi_i \rangle$ are called *frame coefficients* for x . The bounds A and B are called lower and upper frame bounds, respectively. If an inequality is sharp then the corresponding bound is called an *optimal frame bound*. An equivalent formulation of (7) is that the analysis operator associated with

$(\phi_i)_{i \in I}$, given by

$$\begin{aligned} \Theta : \mathbb{R}^n &\rightarrow \mathbb{R}^m \\ x &\mapsto (\langle x, \phi_i \rangle)_{i \in I}, \end{aligned} \tag{8}$$

is one-to-one and the optimal frame bounds coincide with the smallest and largest eigenvalues of the frame operator $\Theta^* \Theta$. Within the paradigm of this paper, the analysis operator is the *linear measurement operator*, i.e., measurement without activation.

2.2. Injectivity of ReLU layers and saturation recovery

For a frame $(\phi_i)_{i \in I}$, a bias vector $\alpha \in \mathbb{R}^m$, and the ReLU function defined in (2), the non-linear measurement operator that corresponds to a ReLU layer is given by

$$\begin{aligned} C_\alpha : \mathbb{R}^n &\rightarrow \mathbb{R}^m \\ x &\mapsto (\text{ReLU}(\langle x, \phi_i \rangle - \alpha_i))_{i \in I}. \end{aligned} \tag{9}$$

To ensure injectivity of C_α we need to check if for any input $x \in \mathbb{R}^n$, the set of frame elements that are not affected by ReLU – i.e., *activated* – forms a frame. Hence, we are interested in those frame elements that satisfy $\langle x, \phi_i \rangle - \alpha_i \geq 0$. Denoting the set of coordinates that get activated for an $x \in \mathbb{R}^n$ as

$$I_\alpha(x) = \{i \in I : \langle x, \phi_i \rangle \geq \alpha_i\}, \tag{10}$$

we have the following characterization of the injectivity of C_α [15, 16, 22].

Theorem 2.1 (Puthawala et al. 2022, Haider et al., 2024). *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\alpha \in \mathbb{R}^m$ a fixed bias. The following are equivalent.*

- (i) *The ReLU layer C_α is one-to-one.*
- (ii) *For all $x \in \mathbb{R}^n$, the activated collection $(\phi_i)_{i \in I_\alpha(x)}$ is a frame for \mathbb{R}^n .*

For λ -saturation, we ask for the injectivity of the saturated measurement operator, which, for a saturation level $\lambda > 0$ and the saturation function σ_λ defined in (3) is given by

$$\begin{aligned} S_\lambda : \mathbb{B}_{\mathbb{R}^n} &\rightarrow \mathbb{R}^m \\ x &\mapsto (\sigma_\lambda(\langle x, \phi_i \rangle))_{i \in I}. \end{aligned} \tag{11}$$

Analogous to the ReLU case, for any $x \in \mathbb{B}_{\mathbb{R}^n}$ we have to check if the set of frame elements that are not affected by the saturation function – i.e., *activated* – forms a frame. For this, the magnitudes of the frame coefficients has to lie below the saturation level λ . Denoting the set of all non-saturated coordinates for x as

$$I_\lambda(x) = \{i \in I : |\langle x, \phi_i \rangle| \leq \lambda\}, \quad (12)$$

we have the following characterization [1].

Theorem 2.2 (Alharbi et al., 2024). *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\lambda > 0$ be a fixed saturation level. The following are equivalent.*

- (i) *The saturated measurement operator S_λ is one-to-one on $\mathbb{B}_{\mathbb{R}^n}$.*
- (ii) *For all $x \in \mathbb{B}_{\mathbb{R}^n}$, the unsaturated collection $(\phi_i)_{i \in I_\lambda(x)}$ is a frame for \mathbb{R}^n .*

2.3. An excursion to other non-linear measurement

In both of the above discussed cases, the characterization of the injectivity reads as that the activated (input-dependent) sub-collections of the measurement vectors $(\phi_i)_{i \in I}$ need to be frames. This argument can be naturally transferred to other suitable activations functions, too. We demonstrate this for the complementary setting to λ -saturation, i.e., where measurements are only recorded accurately if they are larger than a certain threshold $\mu > 0$. The corresponding *gating* activation function is given by $\gamma_\mu(t) = t$ if $|t| \geq \mu$ and $\gamma_\mu(t) = 0$ if $|t| < \mu$. Since very small vectors will never be gated – i.e., *activated* – we consider $X = \mathbb{R}^n \setminus r\mathbb{B}_{\mathbb{R}^n}$ for $r > 0$ as an input domain here. By a scaling argument, we can without loss of generality assume $r = 1$. The gated measurement operator is given by

$$\begin{aligned} G_\mu : \mathbb{R}^n \setminus \mathbb{B}_{\mathbb{R}^n} &\rightarrow \mathbb{R}^m \\ x &\mapsto (\gamma_\mu(\langle x, \phi_i \rangle))_{i \in I}. \end{aligned} \quad (13)$$

Letting $I_\mu(x) = \{i \in I : |\langle x, \phi_i \rangle| \geq \mu\}$, we can use the same proof techniques as in [1] and [16] to show that for a fixed gating level $\mu > 0$, the gated measurement operator G_μ is one-to-one on $\mathbb{R}^n \setminus \mathbb{B}_{\mathbb{R}^n}$ if and only if for any $x \in \mathbb{R}^n \setminus \mathbb{B}_{\mathbb{R}^n}$, the gated collection $(\phi_i)_{i \in I_\mu(x)}$ is a frame for \mathbb{R}^n . The key here is that the domain $\mathbb{R}^n \setminus \mathbb{B}_{\mathbb{R}^n}$ is an open set in \mathbb{R}^n .

For (real) phase retrieval, the characterizing property of the underlying frame to ensure injectivity is the famous complement property [3]. It turns out that we can reformulate it so that the characterization reads analogously to Theorems 2.1, 2.2. We elaborate on this in the following.

2.4. Injectivity of phase retrieval - revisited

In phase retrieval, the non-linear measurement operator is the *intensity measurement operator*, arising as

$$\begin{aligned} \mathcal{M} : \mathbb{R}^n / \sim &\rightarrow \mathbb{R}^m \\ [x]_{\sim} &\mapsto (|\langle x, \phi_i \rangle|)_{i \in I}. \end{aligned} \tag{14}$$

It is well-known that \mathcal{M} is one-to-one if and only if for any subset $J \subseteq I$, either $(\phi_i)_{i \in J}$ or $(\phi_i)_{i \in J^c}$ is a frame for \mathbb{R}^n . This condition is called the *complement property*. While the complement property considers all partitions $\{J, J^c\}$ of I , in the proof of Theorem 2.8 in [3] where the equivalence was first proven, only partitions of a specific form are considered. In particular, one has to ensure that for any pair $x, y \in \mathbb{R}^n$ either the collection of frame elements where the frame coefficients for x and y have the same signs is a frame, or the collection where they have opposite signs is a frame. We shall denote the corresponding sets of coordinates as

$$I^+(x, y) = \{i \in I : \langle x, \phi_i \rangle \langle y, \phi_i \rangle \geq 0\}, \tag{15}$$

$$I^-(x, y) = \{i \in I : \langle x, \phi_i \rangle \langle y, \phi_i \rangle \leq 0\}. \tag{16}$$

Note that if $\langle x, \phi_i \rangle \langle y, \phi_i \rangle = 0$ for some $i \in I$ then $\{I^+(x, y), I^-(x, y)\}$ is not a partition of I . This technicality does not interfere with the result though as it occurs on a nowhere dense subset. Using this perspective, we may reformulate the injectivity characterization of \mathcal{M} stated in [3] in the following sense.

Theorem 2.3 (Reformulation of Balan et al., 2006). *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n . The following are equivalent.*

- (i) *The intensity measurement operator \mathcal{M} is one-to-one.*
- (ii) *For any $J \subseteq I$, either $(\phi_i)_{i \in J}$ or $(\phi_i)_{i \in J^c}$ is a frame for \mathbb{R}^n .*
- (ii*) *For any $x, y \in \mathbb{R}^n$, either $(\phi_i)_{i \in I^+(x, y)}$ or $(\phi_i)_{i \in I^-(x, y)}$ is a frame for \mathbb{R}^n .*

Point (ii*) now says that injectivity of \mathcal{M} can be determined by checking the frame condition for specific sub-collections of $(\phi_i)_{i \in I}$ that depend on x, y - just as for the injectivity of ReLU layers and saturation recovery. While this demonstrates that the conditions for the injectivity of the non-linear measurement operators are of similar nature for all three problems, the one for phase retrieval stands out in a particular way as it depends on one of two sets of vectors being a frame. The reason for this is that the intensity measurement operator \mathcal{M} is defined on the quotient space \mathbb{R}^n / \sim and we need to compensate for the problem of not a priori knowing whether we wish to compare x to y or x to $-y$. Hence, to prove that \mathcal{M} is injective on \mathbb{R}^n / \sim , we need to ensure that if $\mathcal{M}[x]_{\sim} = \mathcal{M}[y]_{\sim}$ for some $x, y \in \mathbb{R}^n$ then $x = y$ or $x = -y$. If $\mathcal{M}[x]_{\sim} = \mathcal{M}[y]_{\sim}$ and $(\phi_i)_{i \in I^+(x,y)}$ is a frame then it follows that $x = y$. Otherwise, if $\mathcal{M}[x]_{\sim} = \mathcal{M}[y]_{\sim}$ and $(\phi_i)_{i \in I^-(x,y)}$ is a frame then $x = -y$. This also contributes to the difficulty in building a reconstruction map for phase retrieval [21, 7, 20]. On the other hand, in the case of ReLU layers and saturation recovery, the activated frame elements can not only be used to determine injectivity but also to build a locally linear reconstruction map [1, 16].

3. Optimal lower Lipschitz bounds

This section represents the main part of the paper. We will derive optimal lower Lipschitz bounds for ReLU layers and saturated measurements, and also present a slightly improved version of the known stability result for phase retrieval.

3.1. Stability of ReLU layers

From the characterization in Theorem 2.1 we can immediately deduce that if a ReLU layer C_α is one-to-one on \mathbb{R}^n then every activated collection $(\phi_i)_{i \in I_\alpha(x)}$ has a positive lower frame bound. We shall refer to the optimal lower frame bound of $(\phi_i)_{i \in I_\alpha(x)}$ as $A_\alpha(x)$. Since there are only finitely many different combinations of activated weight vectors, there is a vector $x^* \in \mathbb{R}^n$ such that $A_\alpha(x^*) \leq A_\alpha(x)$ for all $x \in \mathbb{R}^n$. We denote this value as

$$A_\alpha = \min_{x \in \mathbb{R}^n} A_\alpha(x). \quad (17)$$

Note that for any injective ReLU layer C_α , we can find an open set $U \subseteq \mathbb{R}^n$ so that $I_\alpha(x) = I_\alpha(y)$ for all $x, y \in U$ and that $A_\alpha = A_{\alpha(x)}$ for all $x \in U$.

We have that C_α is affine linear on U , and as U is open, A_α is the largest number so that $A_\alpha \|x - y\|^2 \leq \|C_\alpha(x) - C_\alpha(y)\|^2$ for all $x, y \in U$. Therefore, no lower Lipschitz bound can exceed $\sqrt{A_\alpha}$.

The following theorem provides a lower Lipschitz bound that is specific to the input vectors $x, y \in \mathbb{R}^n$. In [22], the authors consider the line segment from x to y and use that it crosses through at most m polytopes where C_α is linear. Because of this, their lower Lipschitz bound is dependent on m . We use a different argument which only considers the midpoint $\frac{x+y}{2}$. This allows for a lower Lipschitz bound for C_α which is a constant multiple of $\sqrt{A_\alpha}$.

Theorem 3.1. *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\alpha \in \mathbb{R}^m$ a bias vector. For all $x, y \in \mathbb{R}^n$, we have that*

$$\frac{1}{4} A_\alpha \left(\frac{x+y}{2} \right) \|x - y\|^2 \leq \|C_\alpha(x) - C_\alpha(y)\|^2. \quad (18)$$

Proof. We first claim that

$$(I_\alpha(x) \cap I_\alpha(y)) \subseteq I_\alpha \left(\frac{x+y}{2} \right). \quad (19)$$

Indeed, if $i \in (I_\alpha(x) \cap I_\alpha(y))$ then $\langle x, \phi_i \rangle \geq \alpha_i$ and $\langle y, \phi_i \rangle \geq \alpha_i$. Thus, $\langle x+y, \phi_i \rangle \geq 2\alpha_i$, showing that $(I_\alpha(x) \cap I_\alpha(y)) \subseteq I_\alpha \left(\frac{x+y}{2} \right)$.

We now prove that

$$I_\alpha \left(\frac{x+y}{2} \right) \subseteq (I_\alpha(x) \cup I_\alpha(y)). \quad (20)$$

Let $i \in I_\alpha \left(\frac{x+y}{2} \right)$ then $\langle x+y, \phi_i \rangle \geq 2\alpha_i$. Either $\langle x, \phi_i \rangle \geq \alpha_i$ or $\langle y, \phi_i \rangle \geq \alpha_i$. This gives that $i \in (I_\alpha(x) \cup I_\alpha(y))$.

Furthermore, note that if $i \in I_\alpha \left(\frac{x+y}{2} \right)$ then

$$2(\langle x, \phi_i \rangle - \alpha_i) \geq 2\langle x, \phi_i \rangle - \langle x+y, \phi_i \rangle = \langle x-y, \phi_i \rangle.$$

If we have as well that $i \in (I_\alpha(x) \setminus I_\alpha(y)) \cap I_\alpha \left(\frac{x+y}{2} \right)$ then both sides of the above inequality are positive. Hence, we have that

$$2|\langle x, \phi_i \rangle - \alpha_i| \geq |\langle x-y, \phi_i \rangle| \quad \text{for all } i \in (I_\alpha(x) \setminus I_\alpha(y)) \cap I_\alpha \left(\frac{x+y}{2} \right). \quad (21)$$

We now obtain the following estimate for $x, y \in \mathbb{R}^n$.

$$\begin{aligned}
\|C_\alpha(x) - C_\alpha(y)\|^2 &= \sum_{i \in I_\alpha(x) \setminus I_\alpha(y)} |\langle x, \phi_i \rangle - \alpha_i|^2 + \sum_{i \in I_\alpha(y) \setminus I_\alpha(x)} |\langle y, \phi_i \rangle - \alpha_i|^2 \\
&\quad + \sum_{i \in I_\alpha(x) \cap I_\alpha(y)} |\langle x - y, \phi_i \rangle|^2 \\
&\geq \sum_{\substack{i \in (I_\alpha(x) \setminus I_\alpha(y)) \\ \cap I_\alpha(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 + \sum_{\substack{i \in (I_\alpha(y) \setminus I_\alpha(x)) \\ \cap I_\alpha(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 \\
&\quad + \sum_{i \in I_\alpha(x) \cap I_\alpha(y)} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 \quad \text{by (21),} \\
&= \sum_{\substack{i \in (I_\alpha(x) \cup I_\alpha(y)) \\ \cap I_\alpha(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 \quad \text{by (19),} \\
&= \sum_{i \in I_\alpha(\frac{x+y}{2})} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 \quad \text{by (20),} \\
&\geq \frac{1}{4} A_\alpha(\frac{x+y}{2}) \|x - y\|^2.
\end{aligned}$$

□

As $\kappa_L \leq \sqrt{A_\alpha}$, we have the following corollary which proves Theorem 1 for the case of ReLU layers.

Corollary 3.2. *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\alpha \in \mathbb{R}^m$ a bias vector such that the associated ReLU layer C_α is one-to-one. If κ_L is the optimal lower Lipschitz bound of C_α then*

$$\frac{1}{2} \sqrt{A_\alpha} \leq \kappa_L \leq \sqrt{A_\alpha}. \quad (22)$$

When Lipschitz bounds for ReLU layers were first being studied, it was conjectured that it was always the case that $\kappa_L = \sqrt{A_\alpha}$. However, an example is given in [22] where $\kappa_L < \sqrt{A_\alpha}$. We extend this example to give for every $n \in \mathbb{N}$ a simple construction of a frame for \mathbb{R}^n where $\kappa_L = \frac{1}{\sqrt{2}} \sqrt{A_\alpha}$.

Proposition 3.3. *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n with optimal lower frame bound A . The following holds.*

- (i) *The ReLU layer C_α associated with $(\phi_i)_{i \in I} \cup (-\phi_i)_{i \in I}$ and $\alpha = \mathbf{0}$ is one-to-one.*

(ii) $A_\alpha = A$.

(iii) The optimal lower Lipschitz bound for C_α equals $\frac{1}{\sqrt{2}}\sqrt{A}$.

Proof. Point (i) is true since for any $x \in \mathbb{R}^n$ it holds that either $\langle x, \phi_i \rangle \geq 0$ or $-\langle x, \phi_i \rangle \geq 0$. Hence, for all $i \in I$ either ϕ_i or $-\phi_i$ is activated. Since $(\pm\phi_i)_{i \in I}$ is still a frame for any combination of signs, the associated ReLU layer is injective by Theorem 2.1. Moreover, since changing signs of the frame elements does not change the lower frame bound, (ii) immediately follows.

We now prove (iii). Let us denote $I^+(x, y) = \{i \in I : \langle x, \phi_i \rangle \langle y, \phi_i \rangle > 0\}$ and $I^-(x, y) = \{i \in I : \langle x, \phi_i \rangle \langle y, \phi_i \rangle \leq 0\}$. Furthermore, recall the fact that for any $a, b > 0$ it holds that

$$a^2 + b^2 \geq \frac{1}{2}(a + b)^2 \quad (23)$$

with equality if $a = b$. Using this, for any $x, y \in \mathbb{R}^n$ we deduce

$$\begin{aligned} \|C_\alpha(x) - C_\alpha(y)\|^2 &= \sum_{i \in I^+(x, y)} |\langle x, \phi_i \rangle - \langle y, \phi_i \rangle|^2 + \sum_{i \in I^-(x, y)} |\langle x, \phi_i \rangle|^2 + |\langle y, \phi_i \rangle|^2 \\ &\geq \sum_{i \in I^+(x, y)} |\langle x, \phi_i \rangle - \langle y, \phi_i \rangle|^2 + \frac{1}{2} \sum_{i \in I^-(x, y)} |\langle x, \phi_i \rangle - \langle y, \phi_i \rangle|^2 \\ &\geq \frac{1}{2} \sum_{i \in I} |\langle x, \phi_i \rangle - \langle y, \phi_i \rangle|^2 \\ &\geq \frac{1}{2} A \|x - y\|^2. \end{aligned}$$

Thus, we have that $\kappa_L \geq \frac{1}{\sqrt{2}}\sqrt{A}$. Now let $u \in \mathbb{R}^n$ be a unit vector that satisfies $A\|u\|^2 = \sum_{i \in I} |\langle u, \phi_i \rangle|^2$. Setting $x = u$ and $y = -u$, we have that $I^+(x, y) = \emptyset$ and $|\langle x, \phi_i \rangle|^2 = |\langle y, \phi_i \rangle|^2$. With this we obtain

$$\begin{aligned} \|C_\alpha(x) - C_\alpha(y)\|^2 &= \sum_{i \in I^-(x, y)} |\langle x, \phi_i \rangle|^2 + |\langle y, \phi_i \rangle|^2 \\ &= \sum_{i \in I} |\langle x, \phi_i \rangle|^2 + |\langle y, \phi_i \rangle|^2 \quad \text{by invariance of the signs,} \\ &= \frac{1}{2} \sum_{i \in I} |\langle x, \phi_i \rangle - \langle y, \phi_i \rangle|^2 \quad \text{by equality in (23),} \\ &= \frac{1}{2} \sum_{i \in I} |\langle 2u, \phi_i \rangle|^2 \quad \text{by inserting } x = u, y = -u, \\ &= \frac{1}{2} A \|2u\|^2 \quad \text{by design of } u, \\ &= \frac{1}{2} A \|x - y\|^2 \quad \text{as } 2u = x - y. \end{aligned}$$

This gives that $\kappa_L \leq \frac{1}{\sqrt{2}}\sqrt{A}$ and completes the proof. \square

For the special construction of an injective ReLU layer as in the above proposition, it is easy to explicitly compute the optimal lower Lipschitz bound for C_α using a vector where the activated coordinates are disjoint. In a more general case, computing an optimal lower bound is significantly more difficult. Moreover, the bound that we obtain in Proposition 3.3 is still off by a factor of $\frac{1}{\sqrt{2}}$ to the bound that we derived in Corollary 3.2. In the context of neural networks, knowing the best possible estimate for the optimal lower Lipschitz bound for the individual layers is crucial since any errors propagate under the composition of multiple layers. Therefore, our improved lower Lipschitz bound for single ReLU layers becomes even more important when one considers deeper neural networks. Although the numerical experiments in [23] give strong evidence that our factor of $\frac{1}{2}$ is indeed optimal, we are left with the following open problem.

Problem 3.4. *What is the largest $\frac{1}{2} \leq K \leq \frac{1}{\sqrt{2}}$ so that for every frame $(\phi_i)_{i \in I}$ and a bias vector $\alpha \in \mathbb{R}^m$ such that the associated ReLU layer C_α is one-to-one, its optimal lower Lipschitz bound satisfies $\kappa_L \geq K\sqrt{A_\alpha}$?*

3.2. Stability of saturated measurements

Since the condition for recovering a vector from saturated measurements is very closely related to the injectivity of a ReLU layer, it is not surprising that the nature of their lower Lipschitz stability is closely related. As a result, the same technique as in the proof of Theorem 3.1 allows us to derive an explicit lower Lipschitz bound for the saturation recovery problem.

Note that if the saturated measurement operator S_λ is one-to-one on $\mathbb{B}_{\mathbb{R}^n}$ then for every $x \in \mathbb{B}_{\mathbb{R}^n}$ we have that $(\phi_i)_{i \in I_\lambda(x)}$ is a frame of \mathbb{R}^n . We denote $A_\lambda(x)$ to be the optimal lower frame bound for $(\phi_i)_{i \in I_\lambda(x)}$ and denote

$$A_\lambda = \min_{x \in \mathbb{B}_{\mathbb{R}^n}} A_\lambda(x). \quad (24)$$

In [1], it is proven that if λ is greater than some critical threshold then S_λ is injective on $\mathbb{B}_{\mathbb{R}^n}$ if and only if it is bi-Lipschitz. Furthermore, as is the case for ReLU layers, the lower Lipschitz bound satisfies $\kappa_L \leq A_\lambda$. However, the proof does not provide a lower bound on κ_L .

For the presentation of our estimation it will be convenient to introduce some more notation. First, when it comes to saturated coordinates, it is crucial to distinguish between positively and negatively saturated coordinates.

For $x \in \mathbb{B}_{\mathbb{R}^n}$, we set

$$I_\lambda^+(x) = \{i \in I : \langle x, \phi_i \rangle > \lambda\}, \quad (25)$$

$$I_\lambda^-(x) = \{i \in I : \langle x, \phi_i \rangle < -\lambda\}. \quad (26)$$

When considering the situation where for two vectors $x, y \in \mathbb{B}_{\mathbb{R}^n}$ the associated coordinates are both saturated, i.e., $|\langle x, \phi_i \rangle| > \lambda$ and $|\langle y, \phi_i \rangle| > \lambda$, there are two possibilities that can occur to the corresponding frame coefficients:

- If they have the same sign, $i \in (I_\lambda^+(x) \cap I_\lambda^+(y)) \cup (I_\lambda^-(x) \cap I_\lambda^-(y))$, then $|\sigma_\lambda(\langle x, \phi_i \rangle) - \sigma_\lambda(\langle y, \phi_i \rangle)| = 0$.
- If they have opposite signs, $i \in (I_\lambda^+(x) \cap I_\lambda^-(y)) \cup (I_\lambda^-(x) \cap I_\lambda^+(y))$, then $|\sigma_\lambda(\langle x, \phi_i \rangle) - \sigma_\lambda(\langle y, \phi_i \rangle)| = 2\lambda$.

For this reason, the latter case needs to be treated differently as it plays a special role in the estimate. We shall define the corresponding set of coordinates as

$$I_\lambda^\Delta(x, y) = (I_\lambda^+(x) \cap I_\lambda^-(y)) \cup (I_\lambda^-(x) \cap I_\lambda^+(y)). \quad (27)$$

Now we can formulate our result on the lower Lipschitz bound of S_λ that depends on $x, y \in \mathbb{R}^n$, analogous to Theorem 3.1. The global statement follows in Corollary 3.6.

Theorem 3.5. *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\lambda > 0$. For all $x, y \in \mathbb{B}_{\mathbb{R}^n}$, we have that*

$$\left(\frac{1}{4}A_\lambda^\Delta(x, y) + |I_\lambda^\Delta(x, y)| \cdot \lambda^2\right) \|x - y\|^2 \leq \|S_\lambda(x) - S_\lambda(y)\|^2, \quad (28)$$

where $A_\lambda^\Delta(x, y)$ is the optimal lower frame bound of $(\phi_i)_{i \in I_\lambda(\frac{x+y}{2}) \setminus I_\lambda^\Delta(x, y)}$.

It is important to note that if there are no saturated coordinates with opposite signs, i.e., $I_\lambda^\Delta(x, y) = \emptyset$ then $A_\lambda^\Delta(x, y) = A_\lambda(\frac{x+y}{2}) > 0$ is the lower bound. On the other hand, if $I_\lambda^\Delta(x, y) \neq \emptyset$ then $(\phi_i)_{i \in I_\lambda(\frac{x+y}{2}) \setminus I_\lambda^\Delta(x, y)}$ might not be a frame as more frame elements are removed. In this case, it might happen that $A_\lambda^\Delta(x, y) = 0$, and then the lower bound depends on the number of saturated coordinates with opposite signs, given by $|I_\lambda^\Delta(x, y)|$.

Proof. By the same argument as for the ReLU case, we have that

$$(I_\lambda(x) \cap I_\lambda(y)) \subseteq I_\lambda\left(\frac{x+y}{2}\right). \quad (29)$$

Similarly, it holds that

$$(I_\lambda(\frac{x+y}{2}) \setminus I_\lambda^\Delta(x, y)) \subseteq (I_\lambda(x) \cup I_\lambda(y)). \quad (30)$$

Now assume that $i \in (I_\lambda(x) \setminus I_\lambda(y)) \cap I_\lambda(\frac{x+y}{2})$ and observe the following for the two possible cases of saturation for $\langle y, \phi_i \rangle$. If $i \in I_\lambda^+(y)$, then $2|\langle x, \phi_i \rangle - \lambda| \geq |\langle x - y, \phi_i \rangle|$, and if $i \in I_\lambda^-(y)$, then $2|\langle x, \phi_i \rangle + \lambda| \geq |\langle x - y, \phi_i \rangle|$. So in general, we have for all $i \in (I_\lambda(x) \setminus I_\lambda(y)) \cap I_\lambda(\frac{x+y}{2})$ that

$$2|\langle x, \phi_i \rangle - \text{sign}(\langle y, \phi_i \rangle)\lambda| \geq |\langle x - y, \phi_i \rangle|. \quad (31)$$

Using this, we now can deduce the following estimate for $x, y \in \mathbb{B}_{\mathbb{R}^n}$.

$$\begin{aligned} & \|S_\lambda(x) - S_\lambda(y)\|^2 \\ &= \sum_{i \in I_\lambda(x) \setminus I_\lambda(y)} |\langle x, \phi_i \rangle - \text{sign}(\langle y, \phi_i \rangle)\lambda|^2 + \sum_{i \in I_\lambda(y) \setminus I_\lambda(x)} |\langle y, \phi_i \rangle - \text{sign}(\langle x, \phi_i \rangle)\lambda|^2 \\ &\quad + \sum_{i \in I_\lambda(x) \cap I_\lambda(y)} |\langle x - y, \phi_i \rangle|^2 + \sum_{i \in I_\lambda^\Delta(x, y)} 4\lambda^2 \\ &\geq \sum_{\substack{i \in (I_\lambda(x) \setminus I_\lambda(y)) \\ \cap I_\lambda(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 + \sum_{\substack{i \in (I_\lambda(y) \setminus I_\lambda(x)) \\ \cap I_\lambda(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 \\ &\quad + \sum_{i \in I_\lambda(x) \cap I_\lambda(y)} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 + |I_\lambda^\Delta(x, y)| \cdot 4\lambda^2 \quad \text{by (31),} \\ &= \sum_{\substack{i \in (I_\lambda(x) \cup I_\lambda(y)) \\ \cap I_\lambda(\frac{x+y}{2})}} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 + |I_\lambda^\Delta(x, y)| \cdot 4\lambda^2 \quad \text{by (29),} \\ &= \sum_{i \in I_\lambda(\frac{x+y}{2}) \setminus I_\lambda^\Delta(x, y)} \frac{1}{4} |\langle x - y, \phi_i \rangle|^2 + |I_\lambda^\Delta(x, y)| \cdot 4\lambda^2 \quad \text{by (30),} \\ &\geq \frac{1}{4} A_\lambda^\Delta(x, y) \|x - y\|^2 + |I_\lambda^\Delta(x, y)| \cdot 4\lambda^2. \\ &\geq (\frac{1}{4} A_\lambda^\Delta(x, y) + |I_\lambda^\Delta(x, y)| \cdot \lambda^2) \|x - y\|^2 \quad \text{as } 2 \geq \|x - y\|. \end{aligned}$$

□

The following corollary gives Theorem 1 for the case of λ -saturation.

Corollary 3.6. *Let $(\phi_i)_{i \in I}$ be a frame for \mathbb{R}^n and $\lambda > 0$ such that S_λ is one-to-one on $\mathbb{B}_{\mathbb{R}^n}$. If κ_L is the optimal lower Lipschitz bound of S_λ then*

$$\min \left\{ \frac{1}{2} \sqrt{A_\lambda}, \lambda \right\} \leq \kappa_L \leq \sqrt{A_\lambda}. \quad (32)$$

Proof. Recall that the upper bound $\kappa_L \leq \sqrt{A_\lambda}$ was proven in [1]. The lower bound follows from Theorem 3.5 by the following argument. Let $x, y \in \mathbb{B}_{\mathbb{R}^n}$. We first assume that $|I_\lambda^\Delta(x, y)| = 0$ and hence $A_\lambda^\Delta(x, y) = A_\lambda(\frac{x+y}{2}) \geq A_\lambda$. Theorem 3.5 gives that

$$\frac{1}{2}\sqrt{A_\lambda}\|x - y\| \leq \|S_\lambda(x) - S_\lambda(y)\|.$$

We now assume that $|I_\lambda^\Delta(x, y)| \geq 1$. In this case $(\phi_i)_{i \in I_\lambda(\frac{x+y}{2}) \setminus I_\lambda^\Delta(x, y)}$ might not form a frame. The worst case is that $A_\lambda^\Delta(x, y) = 0$ and $|I_\lambda^\Delta(x, y)| = 1$. Theorem 3.5 then gives that

$$\lambda\|x - y\| \leq \|S_\lambda(x) - S_\lambda(y)\|.$$

Thus, we have that $\min\{\frac{1}{2}\sqrt{A_\lambda}, \lambda\} \leq \kappa_L$. \square

In the case of using the ReLU activation function, we obtained an optimal lower bound (up to a factor of at most $\frac{1}{\sqrt{2}}$) for the lower Lipschitz bound of ReLU layer. For the case of λ -saturation, we have a lower bound for the lower Lipschitz bound of the λ -saturated measurement operator of $\min\{\frac{1}{2}\sqrt{A_\lambda}, \lambda\}$. The lower bound necessarily depends on both A_λ and λ . Note that $\frac{1}{2}\sqrt{A_\lambda}$ is optimal up to a constant factor, but it is not clear what the dependence on λ should be. We are left with the following open problem.

Problem 3.7. *What is the function $f(A_\lambda, \lambda)$ so that for every frame $(\phi_i)_{i \in I}$ and saturation level $\lambda > 0$ such that the associated saturated measurement operator S_λ is one-to-one on $\mathbb{B}_{\mathbb{R}^n}$, its optimal lower Lipschitz bound satisfies $\kappa_L \geq f(A_\lambda, \lambda)$?*

Notably, for frames in \mathbb{R}^n with $n + 1$ elements, we can remove λ from the minimum such that the lower Lipschitz bound becomes simply $\frac{1}{2}\sqrt{A_\lambda}$.

Proposition 3.8. *Let $(\phi_i)_{i=1}^{n+1}$ be a frame for \mathbb{R}^n and $\lambda > 0$ such that S_λ is one-to-one on $\mathbb{B}_{\mathbb{R}^n}$. The optimal lower Lipschitz bound for S_λ satisfies*

$$\frac{1}{2}\sqrt{A_\lambda} \leq \kappa_L \leq \sqrt{A_\lambda}. \quad (33)$$

If additionally, $(\phi_i)_{i=1}^{n+1}$ is a finite unit norm tight frame for \mathbb{R}^n and $1 > \lambda > 0$ then the optimal lower Lipschitz bound for S_λ satisfies

$$\frac{1}{2\sqrt{n}} \leq \kappa_L \leq \frac{1}{\sqrt{n}}. \quad (34)$$

Proof. First, recall that by Theorem 3.5, injectivity of S_λ implies that for any $x \in \mathbb{B}_{\mathbb{R}^n}$ the non-saturated coordinates are a frame and therefore $|I_\lambda(x)| \geq n$. Since there are only $n + 1$ frame elements, only up to one can be saturated. Let $x, y \in \mathbb{R}^n$, we distinguish two cases. If $I_\alpha^\Delta(x, y) = \emptyset$ then by Theorem 3.5, we have that

$$\frac{1}{4}A_\lambda\|x - y\|^2 \leq \|S_\lambda(x) - S_\lambda(y)\|^2.$$

We now assume that $I_\alpha^\Delta(x, y) \neq \emptyset$. As S_λ is injective and the frame has $n + 1$ elements, the same single coordinate is saturated for both x and y . That is, $|I_\alpha^\Delta(x, y)| = 1$ and $I_\lambda(x) = I_\lambda(y) = I \setminus I_\alpha^\Delta(x, y)$. We deduce

$$\|S_\lambda(x) - S_\lambda(y)\|^2 = \sum_{i \in I_\alpha(x)} |\langle x - y, \phi_i \rangle|^2 + 4\lambda^2 \geq A_\lambda\|x - y\|^2.$$

Thus, in both cases we have the lower bound in (33).

To prove (34), note that in the case that $(\phi_i)_{i=1}^{n+1}$ is a frame of unit vectors and $1 > \lambda > 0$, we have that A_λ coincides with the smallest optimal lower frame bound of an n -element subset of $(\phi_i)_{i=1}^{n+1}$. Now recall that the frame bound of a unit norm tight frame in \mathbb{R}^n with $n + 1$ elements is given by $\frac{n+1}{n}$. Since any vector of a unit norm tight frame is a tight frame for its own span with bound 1 it follows that the optimal lower bound after removing an element becomes $\frac{n+1}{n} - 1 = \frac{1}{n}$. We deduce that $A_\lambda = \frac{1}{n}$, which together with the first statement finishes the proof. \square

Let us now turn to the open problem about the critical saturation level. First, it is known that for any finite frame $(\phi_i)_{i \in I}$ for \mathbb{R}^n there is a critical value $\lambda_c > 0$ so that S_λ is one-to-one on the unit ball of \mathbb{R}^n if and only if $\lambda \geq \lambda_c$ [1]. Furthermore, if $\lambda > \lambda_c$ then S_λ is bi-Lipschitz stable. Whether saturation recovery is stable also at the critical saturation level was posed as an open problem in [1]. A positive answer immediately follows from Theorem 3.5 as the derived lower Lipschitz bound depends only on the fact that the associated saturated measurement operator S_{λ_c} is one-to-one on the unit ball.

Corollary 3.9. *Saturation recovery is stable at the critical saturation level. In particular, if $(\phi_i)_{i \in I}$ is a finite frame of \mathbb{R}^n and $\lambda > 0$ then S_λ is injective on $\mathbb{B}_{\mathbb{R}^n}$ if and only if S_λ is bi-Lipschitz.*

We note that for the gated measurement operator, the proof technique to derive a lower Lipschitz bound is not directly applicable since $\mathbb{R}^n \setminus \mathbb{B}_{\mathbb{R}^n}$ as the domain is not convex. By a modification, however, one might find a similar result. We leave this as an open problem for the future.

3.3. Stability of phase retrieval - revisited

Phase retrieval was the first of the three non-linear activation function problems to be considered in the context of frame theory [3]. Results on phase retrieval therefore have served as an important inspiration for approaching the injectivity and stability of ReLU layers and saturation recovery. Now, by considering Theorem 2.3 (*ii**) we can turn around the direction of inspiration and revisit the stability of phase retrieval to slightly improve the known results from [4].

To formulate their result on the optimal lower Lipschitz bound in phase retrieval, the authors of [4] make use of the concept of the σ -strong complement property. A frame $(\phi_i)_{i \in I}$ fulfills the σ -strong complement property if for any $J \subseteq I$ the larger of the lower frame bounds of $(\phi_i)_{i \in J}$ and $(\phi_i)_{i \in J^c}$ is at least σ^2 . As with characterizing injectivity, it turns out that it is sufficient to only consider the collections $(\phi_i)_{i \in I^+(x,y)}$ and $(\phi_i)_{i \in I^-(x,y)}$ for any $x, y \in \mathbb{R}^n$ to obtain the result. Denoting by $A_{|\cdot|}^+(x, y)$ the optimal lower frame bound of $(\phi_i)_{i \in I^+(x,y)}$ and by $A_{|\cdot|}^-(x, y)$ the optimal lower frame bound of $(\phi_i)_{i \in I^-(x,y)}$ and defining

$$A_{|\cdot|} = \min_{x,y \in \mathbb{R}^n} \max\{A_{|\cdot|}^+(x, y), A_{|\cdot|}^-(x, y)\}, \quad (35)$$

the following was proven.

Theorem 3.10 (Bandeira et al., 2014). *Let κ_L be the optimal lower Lipschitz bound of \mathcal{M} , and let σ be the largest value such that the σ -strong complement property holds then*

$$\sigma \leq \kappa_L \leq \sqrt{2}\sigma. \quad (36)$$

Moreover,

$$\sqrt{A_{|\cdot|}} \leq \kappa_L \leq \sqrt{2A_{|\cdot|}}. \quad (37)$$

Although not explicitly stated in the paper, the moreover part of Theorem 3.10 follows directly from the proof of Theorem 18 in [4]. Now, since σ^2 is obtained by taking a minimum over all partitions of $(\phi_i)_{i \in I}$ into two sets and $A_{|\cdot|}$ is obtained by taking a minimum only over partitions of a particular form, it must hold that $\sigma^2 \leq A_{|\cdot|}$. Furthermore, it follows from Theorem 3.10 that $\sigma^2 \leq A_{|\cdot|} \leq 2\sigma^2$. Indeed, there are cases where $\sigma^2 = A_{|\cdot|}$ and cases where $A_{|\cdot|} = 2\sigma^2$. One can check that if $(\phi_i)_{i \in I}$ is the Mercedes-Benz frame in \mathbb{R}^2 then $A_{|\cdot|} = \sigma^2$ and if $(\phi_i)_{i \in I}$ consists of two copies of a single non-zero vector in \mathbb{R} then $A_{|\cdot|} = 2\sigma^2$. Thus, the following gives a slight improvement of the original statement in [4].

Corollary 3.11. *The optimal lower Lipschitz bound of \mathcal{M} satisfies*

$$\sqrt{A_{|\cdot|}} \leq \kappa_L \leq \sqrt{2}\sigma. \quad (38)$$

Funding

The research of D. Freeman was supported by NSF grant 2154931. The research of D. Haider was supported by the Marietta-Blau Fellowship of the OeAD (MMC-2023-06983).

References

- [1] W. Alharbi, D. Freeman, D. Ghoreishi, B. Johnson, and N. L. Randraharivony. Declipping and the recovery of vectors from saturated measurements. *ArXiv*, abs/2402.03237, 2024.
- [2] L. Ardizzone, J. Kruse, C. Rother, and U. Köthe. Analyzing inverse problems with invertible neural networks. In *7th International Conference on Learning Representations (ICLR)*, 2019.
- [3] R. Balan, P. Casazza, and D. Edidin. On signal reconstruction without phase. *Applied and Computational Harmonic Analysis*, 20(3):345–356, 2006.
- [4] A. S. Bandeira, J. Cahill, D. G. Mixon, and A. A. Nelson. Saving phase: Injectivity and stability for phase retrieval. *Applied and Computational Harmonic Analysis*, 37(1):106–125, 2014.
- [5] A. Bhandari, F. Krahmer, and R. Raskar. On unlimited sampling and reconstruction. *IEEE Transactions on Signal Processing*, PP(99), 2020.
- [6] J. Bruna, A. Szlam, and Y. Lecun. Signal recovery from pooling representations. In *International Conference on Machine Learning (ICML)*, pages 1585–1598, 2014.
- [7] E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.

- [8] P. G. Casazza and G. Kutyniok. *Finite frames: Theory and applications*. Springer, 2012.
- [9] O. Christensen. *An Introduction to Frames and Riezs Bases*. Birkhäuser, 2003.
- [10] P. L. Combettes and Z. C. Woodstock. Signal recovery from inconsistent nonlinear observations. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 5872–5876, 2022.
- [11] M. Ehler. Shrinkage rules for variational minimization problems and applications to analytical ultracentrifugation. *Journal of Inverse and Ill-posed Problems*, 19(4-5):593–614, 2011.
- [12] S. Foucart and J. Li. Sparse recovery from inaccurate saturated measurements. *Acta Appl. Math.*, 158:49–66, 2018.
- [13] S. Foucart and T. Needham. Sparse recovery from saturated measurements. *Inf. Inference*, 6(2):196–212, 2017.
- [14] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In G. Gordon, D. Dunson, and M. Dudík, editors, *International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 315–323, 2011.
- [15] D. Haider, P. Balazs, and M. Ehler. Convex geometry of ReLU-layers, injectivity on the ball and local reconstruction. In *International Conference on Machine Learning (ICML)*, 2023.
- [16] D. Haider, M. Ehler, and P. Balazs. Injectivity of ReLU layers: Tools from frame theory. *ArXiv*, abs/2406.15856, 2024.
- [17] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk. Democracy in action: quantization, saturation, and compressive sensing. *Appl. Comput. Harmon. Anal.*, 31(3):429–443, 2011.
- [18] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5188–5196, 2015.

- [19] V. Nair and G. E. Hinton. Rectified linear units improve restricted Boltzmann machines. In *International Conference on International Conference on Machine Learning (ICML)*, page 807–814, 2010.
- [20] R. Nenov, D.-K. Nguyen, P. Balazs, and R. Ioan Boț. Accelerated Griffin-Lim algorithm: A fast and provably converging numerical method for phase retrieval. *IEEE Transactions on Signal Processing*, 72:190–202, 2024.
- [21] Z. Průša, P. Balazs, and P. L. Søndergaard. A non-iterative method for (re)construction of phase from stft magnitude. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(5):1154–1164, 2017.
- [22] M. Puthawala, K. Kothari, M. Lassas, I. Dokmanić, and M. de Hoop. Globally injective ReLU networks. *Journal of Machine Learning Research*, 23(105):1–55, 2022.
- [23] C. Rausch. Refining the inverse Lipschitz constant for injective ReLU networks. Master’s thesis, South Dakota State University, 2024.