This manuscript has been submitted for publication in *Bulletin of the Seismological Society of America (BSSA)*. Please note that, despite having gone through peer-review, the manuscript has yet to be formally accepted for publication. Subsequent versions of this manuscript may have slightly different content. If accepted, the final version of this manuscript will be available via the 'Peer-reviewed Publication DOI' link on the right-hand side of this webpage.

A Stability Analysis of Neural Networks and Its Application to Tsunami Early Warning

Donsub Rim*[1], Sanah Suri[1], Sanghyun Hong[2], Kookjin Lee[3], and Randall J. LeVeque[4]

ABSTRACT

Neural networks (NNs) enable precise modeling of complicated geophysical phenomena but are sensitive to small input changes. In this work, we present a new method for analyzing this instability in NNs. We focus our analysis on adversarial examples, test-time inputs with carefully-crafted human-imperceptible perturbations that expose the worst-case instability in a model's predictions. Our stability analysis is based on a low-rank expansion of NNs on a fixed input, and we apply our analysis to a NN model for tsunami early warning which takes geodetic measurements as the input and forecasts tsunami waveforms. The result is an improved description of local stability that explains adversarial examples generated by a standard gradient-based algorithm, and allows the generation of even worse examples. Our analysis can predict whether noise in the geodetic input will produce an unstable output, and identifies a simple approach to filtering the input that enables more robust forecasting from noisy input.

KEY POINTS

- Neural networks suffer from local instabilities which makes it difficult to deploy in safety-critical applications like natural hazard warning.
- We derive a low-rank expansion for feedforward neural networks, writing it as the sum of a input-*independent* linear map and a input-*dependent* low-rank linear map.
- We are able to reduce the instabilities by performing a component analysis of the low-rank map and using projection-based filtering techniques.

Supplemental Material

INTRODUCTION

In recent years, neural networks (NNs) have achieved a remarkable level of performance in making predictions even when they are trained only on empirical observations (LeCun et al., 2015; Goodfellow et al., 2016). Numerous studies have confirmed that these NNs generalize strikingly well in comparison to previous regression models (LeCun et al., 1989; Hochreiter and Schmidhuber, 1997; Krizhevsky et al., 2017; Vaswani et al., 2017). Following such a strong track record, NNs have been considered as surrogates in modeling various complicated physical phenomena, especially where the existing numerical model depends on highly uncertain input data (e.g. initial conditions) or where the model evaluation cannot be done in real-time (Han et al., 8 21; Lu et al., 2021; Li et al., 2021; Mousavi and Beroza, 2022).

However, computational studies have also observed that NNs are susceptible to instabilities under small input perturbations (Biggio et al., 2013; Szegedy et al., 2014). Such an instability would limit their use in applications where stable and accurate prediction is critical. For example, NN-based tsunami early warning models (Makinoshima et al., 2021; Liu et al., 2021; Mulia et al., 2022; Rim et al., 2022) will need to be stable with respect to uncertainty in its input, as the real input measurement is subject to various measurement noise, sensor malfunctions, or other types of anomalies during a tsunamigenic earthquake event (Titov et al., 2005; Kumar and Ahmed, 2011). As NNs are known to produce erroneous predictions when there is a specific type of noise in the input data, analyzing and controlling these instabilities is of utmost importance in order to use them for safety-critical applications.

Perturbed inputs that exhibit the worst-case instabilities in NNs are called adversarial examples. Many analy-

1. Department of Mathematics and Statistics, Washington University in St. Louis, St. Louis, MO 63135, USA , https://orcid.org/0000-0002-6721-2070 (DR) https://orcid.org/0009-0009-6158-3653 (SS); 2. School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, USA, https://orcid.org/0000-0003-4154-7611 (SH); 3. School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281, USA, https://orcid.org/0000-0003-1557-5862 (KL); 4. Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA, https://orcid.org/0000-0003-1384-4504 (RJL); The authors acknowledge there are no conflicts of interest recorded.

*Corresponding author: rim@wustl.edu

Cite this article as D. Rim, S. Suri, S. Hong, K. Lee, R. J. LeVeque (2022). A Stability Analysis of Neural Networks and Its Application to Tsunami Early Warning, *Bull. Seismol. Soc. Am.* **XX**, 1–11, doi: 00.0000/0000000000.

© Seismological Society of America

ses and mitigation strategies were proposed to understand and address this instability, to varying degrees of success (Szegedy et al., 2014; Tramér et al., 2018; Madry et al., 2018; Cohen et al., 2019; Ilyas et al., 2019; Zhang et al., 2019; Shafahi et al., 2019). But even for relatively simple models such as convolutional neural networks (CNNs), adversarial examples persist and no satisfactory stabilization method has been found. Existing techniques often come at a loss of model utility (Zhang et al., 2019) or at a high computational cost (Cohen et al., 2019). On the other hand, greater data set size and greater test accuracy correlates well with greater accuracy against adversarial examples, according to a broad survey (Miller et al., 2021). This suggests simply scaling up might stabilize the NNs eventually, but the scale needed for desired accuracy is infeasible to achieve at the moment. As a result, stabilizing NNs against adversarial examples still remains an open problem.

In this work, we introduce a stability analysis for feedforward NNs derived directly from the NN architecture. Our analysis is based on a new observation that a feedforward NN f has the expansion

$$f(x) = [F_0 + F_\sigma(x)] \cdot x. \tag{1}$$

In the expansion F_0 is a constant matrix that is independent of the input and F_σ is a low-rank matrix that varies with the input x. We say the latter matrix F_σ is low-rank because its rank is at most the number of nonlinear activations, which is less than the number of layers in the NN; The number of layers is usually much smaller than either input or parameter dimensions in practice. This expansion is new, to the best of our knowledge.

The expansion follows from simple linear algebraic tools like Householder reflection and singular value decomposition (SVD) (Trefethen and Bau, 1997; Groetsch, 2011). Moreover, the key components of the decomposition, i.e., the singular vectors of the low-rank matrix, can be estimated efficiently at the complexity required to evaluate the NN once. This implies that the decomposition can potentially be used during training or evaluation for stabilization. These two qualities, the simplicity of the analysis as well as the computational efficiency, are indications that the analysis can be adapted to more general situations.

Based on an analysis of this expansion, in conjunction with existing empirical evidence, we conjecture that the adversarial examples originate from the input-dependent low-rank map F_σ. To support our thesis in a high-consequence context, we apply the analysis to a NN model for tsunami early warning that was trained on synthetic geodetic signals and tsunami waveform (Rim et al., 2022). The NN model has a CNN architecture, and more specifically it uses 1D convolution and transpose convolution layers. Such a NN is already known to be susceptible to adversarial examples: for example a 1D CNN trained on electrocardiogram data suffers from such examples (Han et al., 2020).

We show that adversarial examples found by the projected gradient descent (PGD) algorithm (Kurakin et al., 2017; Madry et al., 2018) have significant components in the unstable terms in the rank-1 expansion of F_σ and that, conversely, the unstable components of F_σ serve as adversarial input perturbations. Similarly, when the unstable components are filtered, the input ceases to cause large changes in the output. This implies that the expansion in Eq. (1) can be used as a computationally efficient algorithm for detecting and filtering adversarial input perturbations.

LOW-RANK EXPANSION OF NNS AND STABILITY ANALYSIS

This section introduces the low-rank expansion for feedforward NNs, and formulates an explanation of how adversarial examples can appear by analyzing that expansion. We begin by defining feedforward NNs whose activation function is set as the rectified linear unit (ReLU) and deriving an expansion. Then we show that the input-dependent part of the expansion has a low-rank representation. Finally, we describe how certain singular behavior in the expansion can lead to adversarial examples.

We will introduce definitions and notations necessary to describe NNs. Given a sequence $n_\ell \in \mathbb{N}$ for index $\ell = 1, \dots, L$ signifying the layer number, let $x \in \mathbb{R}^{n_0}$ denote the input vector, and $\sigma : \mathbb{R} \to \mathbb{R}$ the nonlinear activation (called ReLU) that sets negative values to zero $\sigma(z) \equiv \max\{z, 0\}$. We will omit the range of ℓ when it is clear from the context. Let $A_\ell : \mathbb{R}^{n_\ell \times n_{\ell-1}}$ be linear maps enumerated by the index ℓ. Linear maps can be identified with a matrix in $\mathbb{R}^{n_\ell \times n_{\ell-1}}$ so we write $A_\ell(z) = A_\ell \cdot z = A_\ell z$, for matrix-vector multiplication, and will often refer to a matrix in place of the corresponding linear map. We denote by $\odot$ the entrywise application of the scalar function σ to any vector: if $z = [z_1, \dots, z_{n_\ell}]^T \in \mathbb{R}^{n_\ell}$ then $\sigma \odot z = [\sigma(z_1), \dots, \sigma(z_{n_\ell})]^T \in \mathbb{R}^{n_\ell}$.

A *feedforward neural network* $f : \mathbb{R}^{n_0} \to \mathbb{R}^{n_L}$ is defined as the alternating composition

$$f(x) = A_L \cdot \sigma \odot A_{L-1} \cdot \dots \odot A_2 \cdot \sigma \odot A_1(x). \tag{2}$$

This is a simplified model widely used in literature analyzing NNs (Szegedy et al., 2014; Goodfellow et al., 2016; DeVore et al., 2021). In practice it is common to have bias terms in the maps A_ℓ or to introduce max-pooling layers, but we will omit these here for the simplicity of exposition. They can also be included in our linearization in a straightforward manner. We will re-introduce these in our computational examples when we experiment with NNs used in practice.

In what follows, we will refer to feedforward NNs simply as NNs. The NN is called *fully-connected* when A_ℓ's are allowed to be dense, and the NN is said to be *convo-*

lutional when they are mostly convolutions or transpose convolutions. When there are skip connections, the NN is said to be *residual*. For detailed explanation of these or other terminology, we refer the reader to Goodfellow et al. (2016).

Motivation

Our principle aim is to determine the local stability properties of the NNs defined by Eq. (2). To motivate our approach, we start by revisiting the well-known analysis presented in Szegedy et al. (2014). Throughout, we will use the ℓ_2-norm denoted by $\|\cdot\|_2$ as the norm of choice. One first observes that σ has a Lipschitz constant $L_\sigma = 1$ since

$$\|\sigma(x) - \sigma(y)\|_2 \le \|\max\{x, 0\} - \max\{y, 0\}\|_2 \le \|x - y\|_2 \quad (3)$$

for all $x, y \in \mathbb{R}^{n_\ell}$. This allows one to derive the bound

$$\begin{aligned}
&\|f(x) - f(y)\|_2 \\
&\le \|A_L \cdot \sigma \odot \ldots \odot A_1(x) \;-\; A_L \cdot \sigma \odot \ldots \odot A_1(y)\|_2 \\
&\le \|A_L\|_2 \|\sigma \odot A_{L-1} \ldots \odot A_1(x) \;-\; \sigma \odot A_{L-1} \ldots \odot A_1(y)\|_2 \\
&\le \|A_L\|_2 \|A_{L-1} \ldots \odot A_1(x) \;-\; A_{L-1} \ldots \odot A_1(y)\|_2 \\
&\quad \vdots \\
&\le \left(\prod_{\ell=1}^{L} \|A_\ell\|_2 \right) \|x - y\|_2.
\end{aligned} \quad (4)$$

Due to this inequality, the Lipschitz constant for f is bounded above by

$$L_f \le \left(\prod_{\ell=1}^{L} \|A_\ell\|_2 \right). \quad (5)$$

The bound implies that one can control the Lipschitz constant L_f of the NN by controlling $\|A_\ell\|_2$. This leads to several strategies that impose certain penalties during training (Bai et al., 2021).

This bound neglects the effect of the nonlinear activation. Consider a matrix F_0 that is derived from the matrices A_ℓ,

$$F_0 \equiv A_L \cdot A_{L-1} \cdot \ldots \cdot A_2 \cdot A_1. \quad (6)$$

Then $F_0 \cdot x$ is equal to $f(x)$ when the output of all the intermediate hidden variables all have non-negative entries. Any instability for the linear map F_0 is straightforward to characterize: Its SVD reveals all the singular vectors with large singular values. So one finds all linear subspaces of the domain that can cause an instability through the SVD (see standard texts, for example Groetsch (2011)).

In contrast, adversarial examples depend nonlinearly on the input x and is not restricted to a fixed linear subspace. This suggests that this instability is inherently nonlinear. Moreover, the upper bound of the Lipschitz constant of f given above in Eq. (5) yields an identical estimate for the Lipschitz constants for both f and F_0. That is to say, the bound cannot distinguish between two types of instabilities that have distinct dependencies on the input. As such, it would be surprising if penalizing the spectral norms of A_ℓ's during training provide a good control of these nonlinear instabilities that are present in f but not in F_0.

Expansion using Householder reflectors

Instead of focusing on F_0 to study the stability of f, we focus on a different representation derived from f. First, we write the ReLU activation in an alternate form. Note that ReLU can be written as

$$\sigma \odot z = \frac{1}{2}(z + |z|), \quad (7)$$

where $|z|$ is the entrywise absolute value of the vector. Since the map $z \longmapsto |z|$ is a reflection, one can write it as a Householder reflection (Trefethen and Bau, 1997) in the form

$$H_z = I - 2 v_z v_z^T. \quad (8)$$

The vector v_z is a scalar multiple of $|z| - z$, given by

$$v_z = \begin{cases} (|z| - z)/\sqrt{2z^T(|z| - z)} & \text{if } |z| \ne z, \\ 0 & \text{if } |z| = z. \end{cases} \quad (9)$$

Then we have that the ReLU function, when interpreted as a matrix, is also a rank-1 perturbation of the identity,

$$\sigma \odot z = \frac{1}{2}(I + H_z) z = (I - v_z v_z^T) z. \quad (10)$$

This is nonlinear since v_z depends on z, but below we will linearize about a fixed z in defining the matrix. Next, denoting the hidden units in the feedforward network by

$$z_{\ell+1} \equiv \sigma \odot A_\ell z_\ell, \quad (11)$$

and enumerating the input dependent vectors v_z appearing in the rank-1 perturbation in Eq. (10) corresponding to the input z_ℓ from the previous layer as in

$$v_\ell \equiv v_{z_{\ell+1}} \quad \text{where } z_{\ell+1} = A_\ell z_\ell, \qquad z_0 = x, \quad (12)$$

we have that

$$\sigma \odot A_\ell z_\ell = (I - v_\ell v_\ell^T)(A_\ell z_\ell) = (A_\ell - v_\ell v_\ell^T A_\ell) z_\ell. \quad (13)$$

Now, we view the map as a rank-1 perturbation of A_ℓ. Denote the perturbation by

$$M_\ell \equiv v_\ell w_\ell^T, \qquad \text{where } w_\ell \equiv -A_\ell^T v_\ell, \quad \ell = 1, \ldots, L-1. \quad (14)$$

So the ReLU σ applied to $A_\ell z_\ell$ is the sum of two linear maps A_ℓ and M_ℓ applied to z_ℓ,

$$\sigma \odot A_\ell z_\ell = (A_\ell + M_\ell) z_\ell, \qquad \ell = 1, \ldots, L-1. \quad (15)$$

Note that each M_ℓ is dependent on the input, that is, $M_\ell = M_\ell(x)$. If we let $M_L = 0$, we have that f can be written as a matrix-vector product

$$f(x) = F(x) \cdot x, \tag{16}$$

where the matrix $F : \mathbb{R}^{n_0} \to \mathbb{R}^{n_L \times n_0}$ is given by

$$F(x) = [A_L + M_L(x)][A_{L-1} + M_{L-1}(x)] \cdots [A_1 + M_1(x)]. \tag{17}$$

We will call F the *matrix representation* of NN f.

One expands the matrix product to see that F is a sum of 2^L linear transformations. Let us denote each of them by F_b where b is an integer represented in L-bits,

$$F_b = B_L B_{L-1} \dots B_1, \quad \text{where } B_\ell \equiv \begin{cases} A_\ell & \text{if } (b)_\ell = 0, \\ M_\ell & \text{if } (b)_\ell = 1. \end{cases} \tag{18}$$

With this notation we may write

$$F = \sum_{b=0}^{2^L-1} F_b. \tag{19}$$

For any $b > 0$, there is a least one ℓ such that $(b)_\ell = 1$ so $B_\ell = M_\ell$ which is rank-1, implying that

$$\text{rank}(F_b) = \text{rank}(B_L \dots B_1) \le \min_\ell(\text{rank}(B_\ell)) \le 1. \tag{20}$$

Writing the terms F_b with $b > 0$ as a separate sum, F is written

$$F(x) = F_0 + F_\sigma(x), \quad \text{where } F_\sigma(x) \equiv \sum_{b=1}^{2^L-1} F_b(x). \tag{21}$$

Here F_σ is dependent on the input x, but F_0 is not. Moreover, F_σ is a sum of rank-1 matrices resulting from the nonlinear activations, and its stability properties cannot be inferred solely from that of F_0. We will show in the next section that it can be viewed as a low-rank perturbation of F_0, and its kernel and range are critical subspaces in relation to adversarial examples.

We briefly remark that Leaky ReLU can be used in place of ReLU above. Leaky ReLU $\hat{\sigma}$ can be written as

$$\hat{\sigma} \odot z = (1-\beta)z + \beta\,|z| \tag{22}$$

for some parameter $\beta \in [\frac{1}{4}, \frac{1}{2}]$. We write

$$\begin{aligned} \hat{\sigma} \odot z &= ((1-\beta)I + \beta H_z)z \\ &= ((1-\beta)I + \beta(I - v_z v_z^T))z = (I - \beta v_z v_z^T)z, \end{aligned} \tag{23}$$

and the expansion in Eq. (21) follows with minor changes.

Low-rank expansion of NNs

Recall that the matrices F_b with $b > 0$ appearing in the input dependent part of F_σ in Eq. (21) and are each rank-1, by Eq. (20). But since the sum in Eq. (21) includes 2^L terms, it seems that F_σ could have very high rank. However, we now show that it is at most rank $L-1$ because of the structure of these matrices, a key fact for our analysis. We show that the domain and range for such F_b belong to fixed linear subspaces independent of b, and that these subspaces have dimensions at most $L-1$. More precisely, we show that there are linear subspaces $\Psi \subset \mathbb{R}^{n_L}$ and $\Phi \subset \mathbb{R}^{n_0}$ satisfying

$$\begin{aligned} &\text{(i) } \text{Range}(F_b) \subset \Phi, \ [\text{Ker}(F_b)]^\perp \subset \Psi \text{ for all } b > 0, \\ &\text{(ii) } \dim(\Phi), \dim(\Psi) \le L-1. \end{aligned} \tag{24}$$

We will proceed by explicitly finding such Φ and Ψ. For any F_b with $b > 0$, let us denote

$$\begin{aligned} \ell_{\min}(b) &\equiv \min\{\ell = 1, \dots, L : (b)_\ell = 1\}, \\ \ell_{\max}(b) &\equiv \max\{\ell = 1, \dots, L : (b)_\ell = 1\}. \end{aligned} \tag{25}$$

For example, given a 9-bit number $b = 001010100$, we would have $\ell_{\min}(b) = 3$ and $\ell_{\max}(b) = 7$. We will sometimes drop the dependence on b and write $\ell_{\min}, \ell_{\max}$ for $\ell_{\min}(b), \ell_{\max}(b)$ for brevity.

Recalling that each term F_b appearing in F_σ was defined as the matrix product

$$F_b = A_L \cdots A_{\ell_{\max}+1} M_{\ell_{\max}} \cdots M_{\ell_{\min}} A_{\ell_{\min}-1} \cdots A_1, \tag{26}$$

and inserting the outer-product form for $M_{\ell_{\min}}$ and $M_{\ell_{\max}}$ as in Eq. (14),

$$F_b = A_L \cdots A_{\ell_{\max}+1}(v_{\ell_{\max}} w_{\ell_{\max}}^T) \cdots (v_{\ell_{\min}} w_{\ell_{\min}}^T) A_{\ell_{\min}-1} \cdots A_1. \tag{27}$$

Grouping the matrix-vector products we may rearrange,

$$\begin{aligned} F_b &= (A_L \cdots A_{\ell_{\max}+1} v_{\ell_{\max}})(w_{\ell_{\max}}^T \cdots v_{\ell_{\min}})(w_{\ell_{\min}}^T A_{\ell_{\min}-1} \cdots A_1) \\ &= (w_{\ell_{\max}}^T \cdots v_{\ell_{\min}})(A_L \cdots A_{\ell_{\max}+1} v_{\ell_{\max}})(w_{\ell_{\min}}^T A_{\ell_{\min}-1} \cdots A_1), \end{aligned} \tag{28}$$

since $w_{\ell_{\max}}^T \cdots v_{\ell_{\min}}$ is scalar. As a result F_b is an outer product of two vectors,

$$F_b = c_b\, \phi_{\ell_{\min}} \psi_{\ell_{\max}}^T \quad \text{for } b > 0, \tag{29}$$

in which the vectors ϕ_ℓ, ψ_ℓ and scalar c_b are defined as follows.

- The vectors ϕ_ℓ and ψ_ℓ are

  $$\phi_\ell \equiv A_L \cdots A_{\ell+1} v_\ell, \quad \psi_\ell \equiv -A_1^T \cdots A_{\ell-1}^T A_\ell^T v_\ell. \tag{30}$$

  Corresponding linear spaces spanned by $\{\phi_\ell\}$ and $\{\psi_\ell\}$ are the linear subspaces from Eq. (24) we are seeking, so we let

  $$\begin{aligned} \Phi &\equiv \text{span}\{\phi_\ell : \ell = 1, \dots, L-1\}, \\ \Psi &\equiv \text{span}\{\psi_\ell : \ell = 1, \dots, L-1\}. \end{aligned} \tag{31}$$

- The scalar coefficients c_b are

  $$c_b \equiv w_{\ell_{\max}}^T B_{\ell_{\max}-1} \cdots B_{\ell_{\min}+1} v_{\ell_{\min}} \tag{32}$$

  where B_ℓ are as in Eq. (18). If this is an empty product, we set $c_b = 1$.

As a result of Eq. (29) the matrices F_b ($b = 1, \dots, 2^L - 1$) have their range in Φ and the orthogonal complement of its kernel in Ψ as required in Eq. (24). We note here that while the

number of terms in the sum for F_σ in Eq. (21) is exponential in L, the domain and range of F_σ belongs to a linear space whose dimension is linear in L.

Put in other terms, the input-dependent part F_σ has the rank-1 expansion,

$$F_\sigma = \sum_{\ell,\ell'=1}^{L-1} C_{\ell\ell'}\, \phi_\ell \psi_{\ell'}^T \quad \text{where } C_{\ell\ell'} = \sum_{b=1}^{2^L-1} \sum_{\substack{\ell_{\min}(b)=\ell \\ \ell_{\max}(b)=\ell'}} c_b. \tag{33}$$

It is natural to take the SVD of the matrix $(C_{\ell\ell'})$ to obtain orthonormal bases $Z = \{\zeta_\ell\}_{\ell=1}^r$ of Φ and $\Xi = \{\xi_\ell\}_{\ell=1}^r$ of Ψ, respectively, that have dimensions $r \le L-1$ and transforms the double sum in Eq. (33) into a single sum. The new sum with $L-1$ terms is expressed as the SVD

$$F_\sigma = \sum_{\ell=1}^{r} d_\ell \zeta_\ell \xi_\ell^T \quad \text{with } r \le L-1, \tag{34}$$

where $d_1 \ge d_2 \ge \cdots \ge d_r$ are singular values, whereas ζ_ℓ and ξ_ℓ are the left and right singular vectors, respectively. We remind the reader that the functions $d_\ell, \zeta_\ell, \xi_\ell$ as well as the rank r depend on x.

To summarize, a NN can be written $f(x) = F(x)\cdot x$ where

$$F(x) = F_0 + F_\sigma(x) \quad \text{with } \operatorname{rank}(F_\sigma(x)) \le L-1. \tag{35}$$

We refer to this expansion as the *low-rank expansion of NN* f. We say the expansion is low-rank in the sense that the rank of F_σ is less than the the number of layers L and in many applications L is significantly smaller than the dimension of the input n_0 or the dimension of the parameters $\sum_{\ell=0}^{L-1} n_\ell n_{\ell+1}$.

Now one could linearize f about a fixed input x_0 by freezing F_σ in Eq. (35). Doing so yields a constant linear map $F_\sigma(x_0)$. In this case Eq. (13) becomes a rank-1 approximation of ReLU which is first order accurate, and one obtains a linearization

$$F(x_0) = F_0 + F_\sigma(x_0). \tag{36}$$

The resulting $F(x_0)\cdot x$ agrees with $f(x)$ up to first order in the neighborhood of x_0.

We make a few remarks regarding the linearization. Observe that the ReLU linearized by freezing the input in its rank-1 perturbation Eq. (10) is equal to the original ReLU output for positive entries, but is not equal otherwise. As a result, this linearization does not preserve the pattern of zeros introduced by the ReLU and it is distinct from other linearizations that do. There are straightforward extensions of Eq. (13) to higher rank approximations that lead to a family of linearizations similar to that in Eq. (35). We mention that there are other linearization of the nonlinear activations (e.g. see Choromanska et al. (2015)).

Adversarial examples from the low-rank expansion

We now discuss the conditions under which an adversarial example can exist although F_0 is stable. We will assume that the induced 2-norm of F_0, and therefore its Lipschitz constant, is not large. Suppose δx is a perturbation about x_0 satisfying

$$\|\delta x\|_2 < \varepsilon \quad \text{and} \quad \frac{\|F_\sigma(x_0)\cdot \delta x\|_2}{\|F_0\cdot \delta x\|_2} \ge 1. \tag{37}$$

Then we have by the triangle inequality

$$\begin{aligned} \|F(x_0)\cdot\delta x\|_2 &= \|[F_0 + F_\sigma(x_0)]\cdot \delta x\|_2 \\ &\ge \|F_\sigma(x_0)\cdot\delta x\|_2 - \|F_0\cdot\delta x\|_2 \\ &\ge \|F_\sigma(x_0)\cdot\delta x\|_2 - \|F_0\|_2\|\delta x\|_2, \end{aligned} \tag{38}$$

then using the rank-1 expansion in Eq. (34),

$$\|F(x_0)\cdot\delta x\|_2 \ge \left\| \sum_{\ell=1}^{r} d_\ell \xi_\ell^T(\delta x)\zeta_\ell \right\|_2 - \|F_0\|_2\|\delta x\|_2. \tag{39}$$

If we choose the perturbation δx satisfying Eq. (37) that is also parallel to the orthonormal basis function $\xi_{\ell*}$ we have

$$\|F(x_0)\|_2 = \frac{\|F(x_0)\cdot\delta x\|_2}{\|\delta x\|_2} \ge d_{\ell*} - \|F_0\|_2. \tag{40}$$

Large variations in the output relative to the input can arise if $d_{\ell*} \sim 1/\varepsilon$ whereas $\|F_0\|_2 \sim 1$. In that case $\|F(x_0)\|_2 \sim 1/\varepsilon$. Since $F(x_0)$ approximates $F(x)$, the NN must also be unstable with respect to the perturbation δx.

NNs satisfying these conditions are readily found in computational experiments, both in toy NNs with randomly assigned weights and in more realistic NNs with trained weights. We will demonstrate this in the following section.

COMPUTATIONAL EXAMPLES

In this section, we will perform computational experiments with two NNs. One is a toy model with randomly given weights without training, and the other is a CNN model trained on simulated geophysical data that was studied in Rim et al. (2022) for the purpose of tsunami early warning.

A Toy Example

We first consider a toy example involving a simple untrained NN with randomly assigned weights. We create a fully-connected NN in the form of Eq. (2) with $L = 7$ layers and dimensions $n_0 = n_1 = \cdots = n_L = 30$. We choose randomly generated weights for all the linear layers A_ℓ, drawn from the standard normal distribution $\mathcal{N}(0,1)$. We also randomly select an input and we expand the NN. The input is a vector whose entries were drawn from the uniform distribution $\mathcal{U}(-1/2,\ 1/2)$.

We compute an orthonormal basis for Φ and Ψ as defined in Eq. (31) by sampling the vectors $\{F_b x_0 : b = 1, \dots, 2^L - 1\}$ and $\{F_b^T f(x_0) : b = 1, \dots, 2^L - 1\}$ then taking their SVD. The

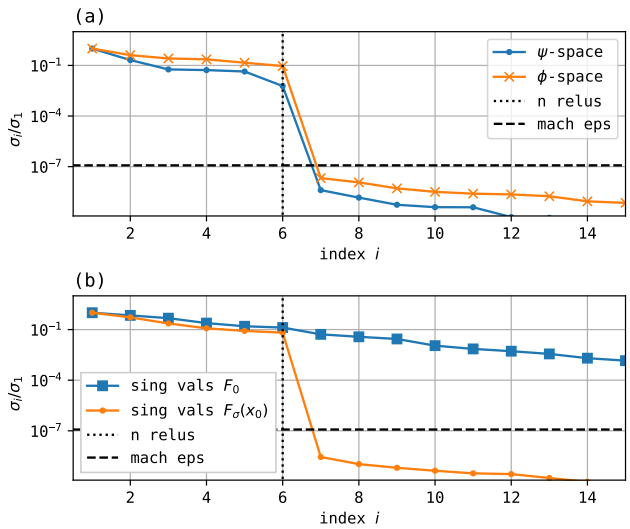

Figure 1. Singular values related to the toy model. (a) Singular values of matrices whose columns are vectors sampled from Φ and Ψ, and (b) singular values of F_0 and $F_\sigma(x_0)$.

singular values are shown in Figure 1 (a). There is a sharp drop in the singular values to single-precision machine epsilon level after the first 6 singular values. This agrees well with our low-rank expansion in Eq. (35), showing that Φ and Ψ are indeed at most $L-1=6$ dimensional.

For this toy model, it is not difficult to compute all the individual terms F_b for $b=1,\dots,2^L-1$ and explicitly form the sum $F_\sigma(x_0)=\sum_{b=1}^{2^L-1}F_b$. Then one obtains the low-rank expansion by taking the SVD of the computed $F_\sigma(x_0)$. The singular values are shown in Figure 1 (b). The index where there is a gap in the singular values agrees precisely with the previous plot. We have computed the correlation between the two sets of singular vectors, and found that the two orthonormal bases span the same space up to the level of numerical precision.

The size of the gap in the singular values suggests that an orthonormal basis for Φ and Ψ can be found simply by sampling $F_b x$ and $F_b^T f(x_0)$ for a number of values b then computing the SVD. Upon computing these bases, one samples the matrix $C_{\ell\ell'}$ and then computes its SVD to estimate the low-rank expansion indirectly.

Next, we attempt to verify whether we can generate adversarial examples from a random toy model. Above we discussed the conditions in which adversarial examples can be precisely identified by the low-rank expansion Eq. (35). We will show that we can meet these conditions with our toy model, if we make use of non-standard noise models for the weights in the linear layers.

We assign randomized weights entrywise by forcing that A_ℓ is lower Hessenberg, that its entries below the first subdiagonal are zero, then drawing the nonzero entries from a normal distribution with a negative bias,

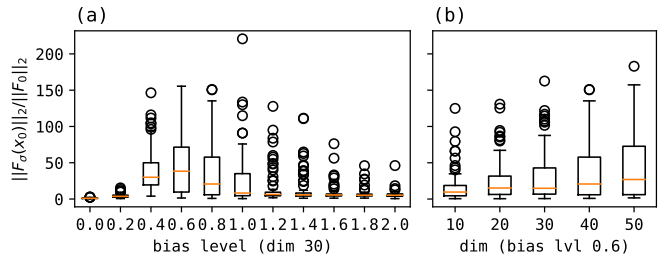

Figure 2. Ratio of spectral norms $R=\|F_\sigma(x_0)\|_2/\|F_0\|_2$ for the toy model with (a) varying dimensions, and (b) varying bias levels B in Eq. (41).

$$(A_\ell)_{ij} \sim \mathcal{N}(-B,1) \quad \text{if } i-j>1 \tag{41}$$

where B is a bias in the range $[0,2]$. We explicitly compute spectral norms of F_0 and $F_\sigma(x_0)$ to compare their ratios

$$R=\frac{\|F_\sigma(x_0)\|_2}{\|F_0\|_2}. \tag{42}$$

Note that if $R\gg 1$ it implies that δx satisfying Eq. (40) with $d_{\ell *}=R\|F_0\|_2$ exists ($\ell_*=1$), resulting in

$$\frac{\|F(x_0)\cdot\delta x\|_2}{\|\delta x\|_2}\geq (R-1)\|F_0\|_2. \tag{43}$$

To observe the dependence on the dimension $m=n_0=\cdots=n_L$ and the bias level B, we calculate the ratio R for 100 randomly drawn NNs for each bias level $B=0.0,0.2,0.4,\dots,2.0$ for fixed dimension $m=30$, then for each dimension $m=10,20,\dots,50$ for fixed bias level $B=0.6$.

The statistics of the ratio is shown in two sets of box plots in Figure 2. The ratio easily reaches 200 and above, meaning that a small perturbation in the input in the direction of the first singular vector (ξ_1 in Eq. (34)) can cause a much larger response in $F_\sigma(x_0)$ compared to that in F_0. One observes that the bias level at around 0.8 maximizes the ratio, and that in the unbiased or highly biased cases the ratio stays modest, although outliers exist in latter case. When the dimension m is increased, the quartiles tend to increase with the dimension. This indicates that among NNs with the same random model for the weights those with larger dimensions are more likely have larger values of R.

Tsunami Early Warning Model

We apply our analysis to a tsunami early warning model (Rim et al., 2022). We compute the low-rank expansion in Eq. (35), estimate the input space Ξ in Eq. (34), then compare the adversarial perturbations found by PGD with perturbations parallel to the singular vectors $\{\xi_\ell\}$ in the expansion. We repeat the experiments for standard noise perturbations.

Model architecture and its expansion

The model is a standard CNN that maps the geodetic time-series measurements coming from 60 Global Navigation

Satellite System (GNSS) stations, each with east (E), north (N), vertical (Z) components. Accordingly, the input dimension is $n_0 = 60 \cdot 3 \cdot 512$ because there are 60 stations and 3 components and the length of the time-series is 512. The output of the neural network is the full tsunami waveform at 3 different gauge locations. The output dimension is $n_L = 3 \cdot 256$, as there are 3 gauges and the length of the time-series is 256. Throughout this section, we will plot inputs and outputs at only two GNSS stations named `bamf` and `lsig` and only show the Z component of the signal, in order to illustrate typical results. We also show the tsunami wave gauge forecasts as only a single gauge location, Gauge 901, which is located in Discovery Bay, WA. The plots for the other components, and at the other GNSS stations or gauge locations, have similar characteristics.

Given an input, the model first applies a sequence of 9 pairs of convolutional and max-pool layers, then applies the 8 transpose convolutional layers. Between each pairs, we apply the Leaky ReLU activation function with negative slope 1/2. The channel output sizes for each convolutional and transpose convolutional layer are 64-64-128-128-128-256-256-512-512 and 512-512-256-256-128-128-64-64, respectively. In the study, an ensemble of 25 NNs were trained. We will analyze the stability of one NN in the ensemble.

For full details regarding the geographical locations of these stations and gauges, as well as the NN model and its training, we refer to the original reference (Rim et al., 2022) (also, see Data and Resources).

This NN architecture includes bias terms and max-pool layers that were absent in our analysis of Eq. (2). Introducing the bias terms result in additional input-dependent terms in our expansion in Eq. (34) that we can include into $F_\sigma(x_0)$ in a straightforward manner, and this does not increase its rank. We also freeze the max-pool layers so that they effectively become permutations, resulting in a linearization. We note here that the technique linearizing the ReLU above can also be applied to max-pool layers, and doing so would result in a different linearization. We will not pursue this here, however.

Basis estimation

We sample the input and output spaces Ψ and Φ by computing the outer product F_b in Eq. (29) for the values

$$b = \begin{bmatrix} 0000...0001 \\ \vdots \\ 0010...0000 \\ 0100...0000 \\ 1000...0000 \end{bmatrix} \quad \begin{vmatrix} 1100...0000 \\ \vdots \\ 1000...0100 \\ 1000...0010 \\ 1000...0001 \end{vmatrix} \quad \begin{bmatrix} 0100...0001 \\ 0010...0001 \\ \vdots \\ 0000...0011 \end{bmatrix}. \quad (44)$$

Then collecting these vectors and taking the SVD, we see a gap in the singular values and obtain an orthonormal bases $\widehat{\Phi} = \{\hat{\phi}_\ell\}$ and $\widehat{\Psi} = \{\hat{\psi}_\ell\}$. See Figure 3 for a plot of singular values of the sampled basis. Next, we use them as surrogates for the sum in Eq. (33)

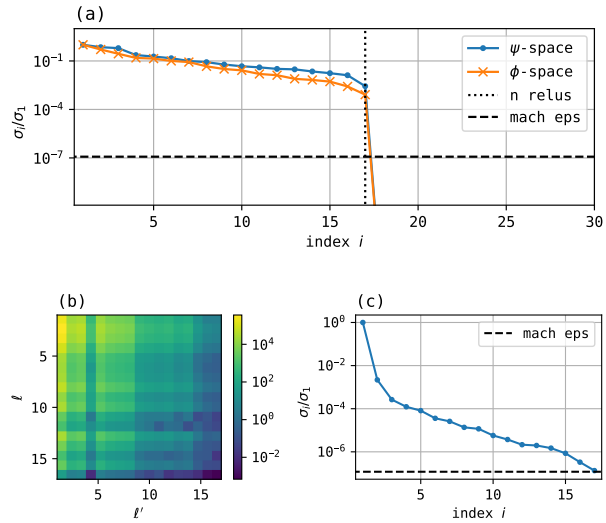

Figure 3. (a) Singular values of matrices whose columns are vectors sampled from Ψ and Φ, (b) logarithmic scale plot of absolute values of $\widehat{C}_{\ell\ell'}$ estimated based on $\widehat{\Psi}$ and $\widehat{\Phi}$, and (c) singular values of the matrix $[\widehat{C}_{\ell\ell'}]$.

$$\widehat{F}_\sigma(x_0) = \sum_{\ell,\ell'} \widehat{C}_{\ell\ell'} \hat{\phi}_\ell \hat{\psi}_{\ell'}. \quad (45)$$

We calculate $\widehat{C}_{\ell\ell'}$ by computing all of F_b for $b = 1, \ldots, 2^L - 1$ and projecting them into the bases $\widehat{\Phi}$ and $\widehat{\Psi}$. Taking the SVD of $\widehat{C}_{\ell\ell'}$, we obtain our approximation $\widehat{Z} = \{\hat{\zeta}_\ell\}$ and $\widehat{\Xi} = \{\hat{\xi}_\ell\}$ of Z and Ξ. A plot of absolute values of $\widehat{C}_{\ell\ell'}$, together with the singular values of $\widehat{C}_{\ell\ell'}$ are shown in Figure 3. First two basis in $\widehat{Z}$ and $\widehat{\Xi}$ are shown in Figure 4.

We remark here that the estimation of $\widehat{C}_{\ell\ell'}$ is a computationally expensive task, whereas the computation of the orthonormal bases $\widehat{\Phi}$ and $\widehat{\Psi}$ only required evaluations of the neural network for a few sample values of b given in Eq. (44). The plot of absolute values of $\widehat{C}_{\ell\ell'}$ shows that the ordering of the bases remains largely unchanged. Therefore the estimated bases $\widehat{\Phi}$ and $\widehat{\Psi}$ can potentially serve as computationally cheap estimates for $\widehat{Z}$ and $\widehat{\Xi}$. As a consequence, the basis in the low-rank expansion can be computed at a small cost at any input x_0, that is, at the cost of one evaluation of the NN.

Adversarial examples from the low-rank expansion

Now that we have an estimate of the low-rank expansion in Eq. (34) for the tsunami model, we perform tests to verify if adversarial examples discovered via optimization lie in the linear subspace Ψ.

We compute an adversarial example $x_0 + \delta x$ about an input x_0 using PGD. In PGD, we attempt to maximize the test loss, a strategy that is sometimes called an untargeted attack. Within the PGD algorithm, we have used both ℓ_∞ and ℓ_2

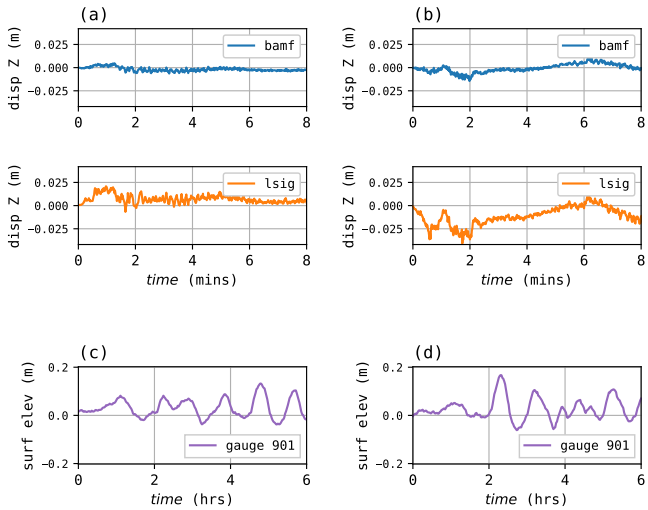

Figure 4. The singular vector pairs (ξ_ℓ, ζ_ℓ) for the domain and the range of F_σ Eq. (34) corresponding to indices $\ell = 1, 2$. (a) The right singular vector ξ_1 plotted as a GNSS input at two selected stations, (c) the left singular vector ζ_1 plotted at a gauge output. Similar plot for ξ_2 (b) and ζ_2 (d).

projections and found that ℓ_∞ projection yields worse adversarial examples. We will use ℓ_∞ throughout but using ℓ_2 projections instead lead to the same conclusions. To measure the size of the perturbations in input and output, we will use the relative ℓ_2 norm.

Next, we compute the projection of δx onto $\widehat{\Psi}$

$$(\delta x)_{\text{proj}} \equiv \widehat{\Psi}\widehat{\Psi}^T(\delta x). \tag{46}$$

then compare the resulting output perturbation. In order to show a fair comparison of the output perturbations from two perturbations δx and $(\delta x)_{\text{proj}}$, we will scale them so that their ℓ_2 norm is 2% of the ℓ_2 norm of the input x_0.

We show in Figures 5 and 6 the adversarial example $x_0 + \delta x$, the input perturbation δx, the resulting output $y_0 + \delta y$ where $y_0 = f(x_0)$. We observe that $(\delta x)_{\text{proj}}$ causes a larger perturbation in the output compared to δx, 57% change versus 27%. Note also that $(\delta x)_{\text{proj}}$ is much more smooth than δx, causing the perturbation to become even less perceptible.

This suggests the possibility that removing from δx its projection to Ψ can lead to a more stable output. So we define the orthogonal complement as

$$(\delta x)_{\text{filter}} \equiv (I - \widehat{\Psi}\widehat{\Psi}^T)(\delta x). \tag{47}$$

We re-scale the filtered input to be 2% of the ℓ_2 norm of the input, as we have for δx and $(\delta x)_{\text{proj}}$, and compute the output of the perturbed input. The results are plot in Figure 7. Now the output perturbation is at a similar level as that of the input perturbation, 3% versus 2%.

Next, following the analysis above, we measure the amount of output perturbation caused by perturbing the

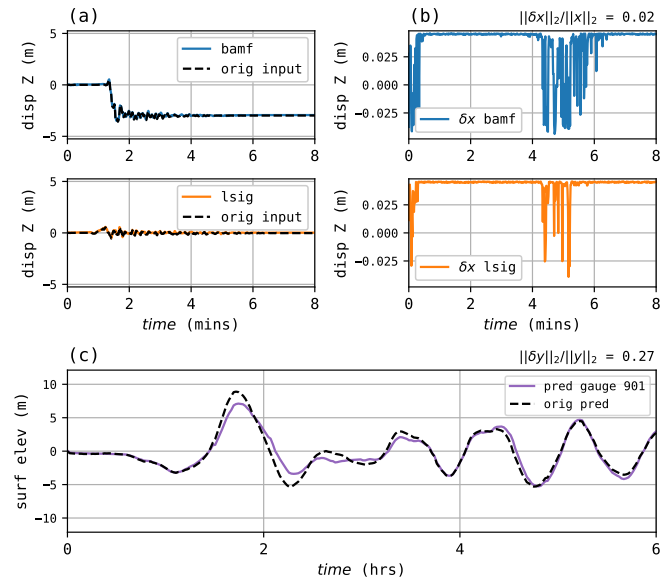

Figure 5. An adversarial example found by PGD. (a) The perturbed input $x_0 + \delta x$ at two selected stations, (b) the perturbation δx, and (c) the resulting perturbed output $f(x_0 + \delta x)$ at gauge 901. An imperceptible 2% change in the input causes a large 27% change in the output.

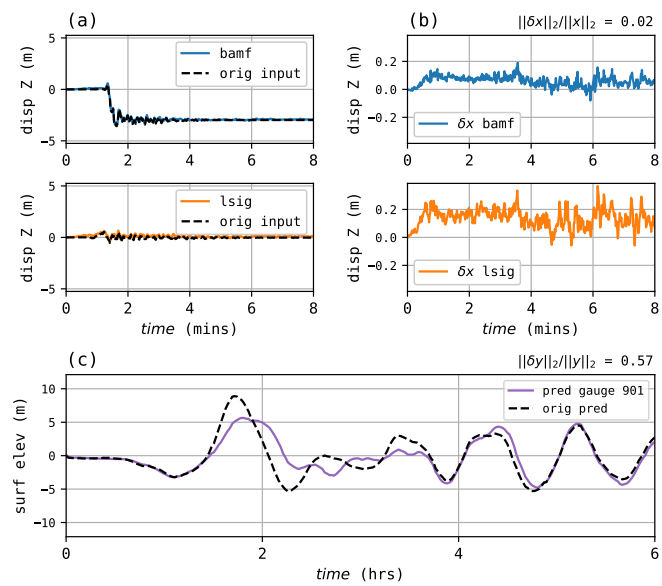

Figure 6. Input and output perturbations for the projected adversarial example. (a) The perturbed input $x_0 + (\delta x)_{\text{proj}}$ at two selected stations, (b) the projected perturbation $(\delta x)_{\text{proj}}$, and (c) the resulting perturbed output $f(x_0 + (\delta x)_{\text{proj}})$. An imperceptible 2% change in input causes a large 57% change in the output.

input using the basis $\widehat{\Xi}$. The results for the first six vectors are shown in Figure 8. Observe that perturbation by ξ_5 causes 83% change in the output, a more severe change even when compared to the adversarial example found by PGD.

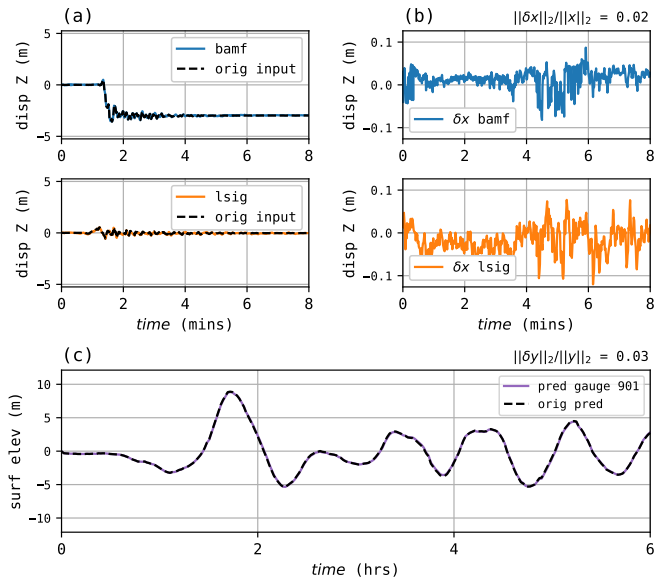

Figure 7. Filtering of an adversarial example found by PGD. (a) The perturbed input $x_0 + (\delta x)_{\text{filter}}$ at two selected stations, (b) the filtered perturbation $(\delta x)_{\text{filter}}$, and (c) the resulting perturbed output $f(x_0 + (\delta x)_{\text{filter}})$. The amount of output perturbation is at 3%, close to that of the input perturbation.

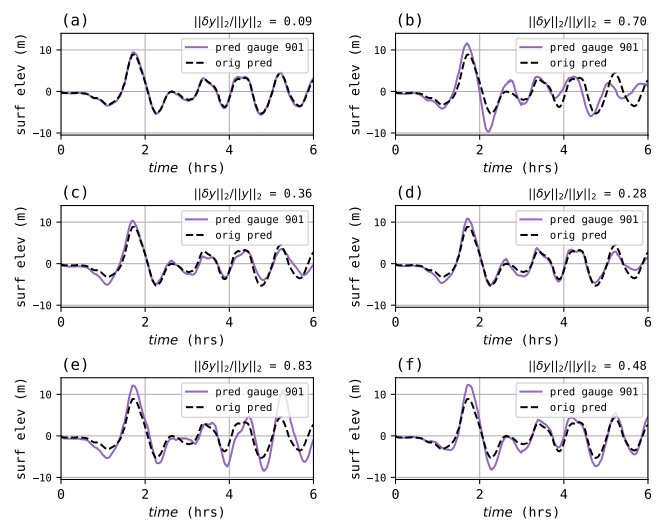

Figure 8. Perturbation of output caused by input perturbation along the basis functions of $\widehat{\Xi}$. (a) to (f) show output perturbation when δx is chosen to be parallel to ξ_ℓ and rescaled to be 2% of the $\|x_0\|_2$, for index ℓ from 1 to 6. These input perturbations cause large changes in the output, even larger than the adversarial perturbations found by PGD in Figure 5.

Noise in GNSS data

We now consider the case that the input perturbation δx is physical noise rather than a specific perturbation in an adversarial example computed by PGD. We generate Brownian noise for each GNSS time-series, then use it as the perturbation δx. The resulting perturbation is shown in

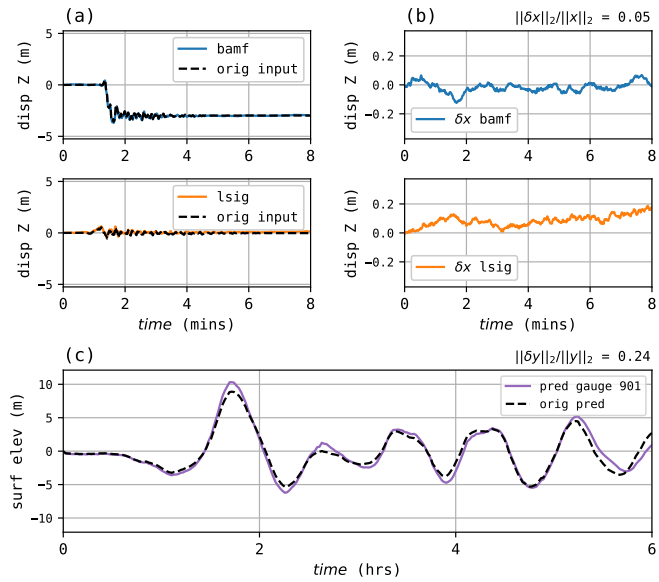

Figure 9. An adversarial example drawn from Brownian noise. (a) The perturbed input $x_0 + \delta x$ at two selected stations, (b) the perturbation δx, and (c) resulting perturbed output $f(x_0 + \delta x)$. Noise level of 5% in the input causes 24% change in the output.

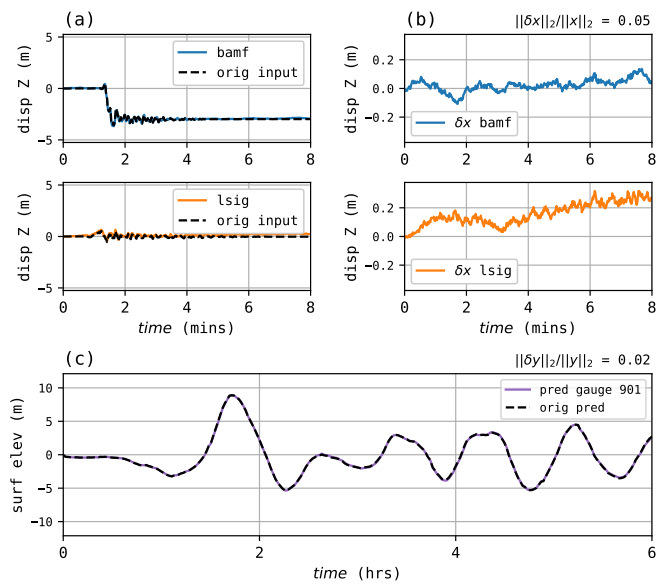

Figure 10. Filtering of Brownian noise. (a) The perturbed input $x_0 + (\delta x)_{\text{filter}}$ at two selected stations, (b) the filtered perturbation $(\delta x)_{\text{filter}}$, and (c) resulting perturbed output $f(x_0 + (\delta x)_{\text{filter}})$. Filtering the input reduces the output perturbation by a factor of 10.

Figure 9. A noise level of 5% causes a 24% change in the output. After filtering, however, the output change is reduced to 3% as shown in Figure 10.

To study the effect of input noise more closely, we experiment with three different types of noise perturbations: white noise, Brownian noise, and power-law noise. For the latter,

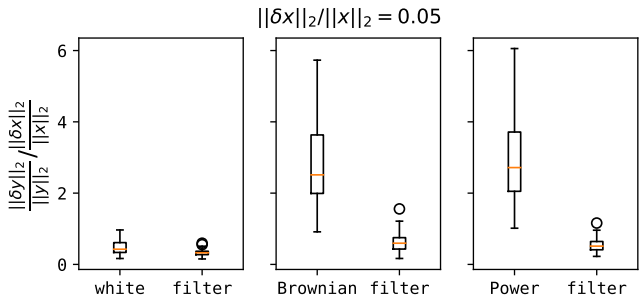

Figure 11. Ratio of input and output changes for different types of noise, white noise, Brownian noise, and power law noise.

the spectral density is chosen as the Gaussian $\exp[-k^2/\sigma^2]$ with $\sigma = \sqrt{10}$. We draw 1000 sample perturbations from each noise and compute the input and output change ratio $||\delta y||_2/||\delta x||_2$. We scale the noise uniformly across samples so that $||\delta x||_2/||x_0||_2$ is around 5%. In addition, we compute the same ratio for the filtered noise $(\delta x)_{\text{filter}}$ as in Eq. (47).

The results are plotted in Figure 11. The NN is not significantly affected by white noise (a similar observation can be made in Figure 7 where filtered perturbation resembles white noise) whereas Brownian noise and power law noise do cause large changes in the output, up to a factor of 5-6. Filtered noise tend to yield significantly smaller output changes.

CONCLUSION

This work proposes a new stability analysis, derived from a new low-rank expansion of feedforward neural networks that uses ReLU activations. Our analysis theoretically explains how adversarial examples can occur. Computational examples demonstrate that the analysis applies to both untrained toy models with random weights and a tsunami warning model trained on data. The analysis suggests new approaches for developing defenses against adversarial examples, which is important for security- or safety-critical applications

DATA AND RESOURCES

We have made use of the data and the NN model from Rim et al. (2022). The geodetic data therein was generated using the MudPy software (Melgar, 2016) that simulates random earthquakes (LeVeque et al., 2016), and the tsunami waveform data was generated using Clawpack (Clawpack Development Team, 2020; Berger et al., 2011). Computational experiments involving NNs were conducted using the PyTorch Library (Paszke et al., 2019).

ACKNOWLEDGEMENTS

The authors thank Robert Baraldi and Christopher M. Liu for carefully reading the manuscript and providing valuable feedback. The authors acknowledge the Research Infrastructure Services (RIS) group at Washington University in St. Louis for providing computational resources and services needed to generate the research results delivered within this paper (URL ris.wustl.edu). The first author thanks Gerrit Welper, Ari Stern, Kui Ren and Weilin Li for helpful discussions.

REFERENCES

Bai, T., J. Luo, J. Zhao, B. Wen, and Q. Wang (2021). Recent Advances in Adversarial Training for Adversarial Robustness. In Z.-H. Zhou (Ed.), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pp. 4312–4321. International Joint Conferences on Artificial Intelligence Organization.

Berger, M. J., D. L. George, R. J. LeVeque, and K. T. Mandli (2011). The geoclaw software for depth-averaged flows with adaptive refinement. *Advances in Water Resources* **34**(9), 1195–1206.

Biggio, B., I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli (2013). Evasion Attacks against Machine Learning at Test Time. In H. Blockeel, K. Kersting, S. Nijssen, and F. Železný (Eds.), *Machine Learning and Knowledge Discovery in Databases*, Berlin, Heidelberg, pp. 387–402. Springer Berlin Heidelberg.

Choromanska, A., M. Henaff, M. Mathieu, G. Ben Arous, and Y. LeCun (2015). The Loss Surfaces of Multilayer Networks. In G. Lebanon and S. V. N. Vishwanathan (Eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, Volume 38 of *Proceedings of Machine Learning Research*, pp. 192–204.

Clawpack Development Team (2020). Clawpack software.

Cohen, J., E. Rosenfeld, and Z. Kolter (2019). Certified Adversarial Robustness via Randomized Smoothing. In K. Chaudhuri and R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, Volume 97 of *Proceedings of Machine Learning Research*, pp. 1310–1320.

DeVore, R., B. Hanin, and G. Petrova (2021). Neural network approximation. *Acta Numerica* **30**, 327–444.

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

Groetsch, C. (2011). *Linear Inverse Problems*, pp. 3–41. New York, NY: Springer New York.

Han, J., A. Jentzen, and W. E (2018-08-21). Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences of the United States of America* **115**(34), 8505 – 8510.

Han, X., Y. Hu, L. Foschini, L. Chinitz, L. Jankelson, and R. Ranganath (2020). Deep learning models for electrocardiograms are susceptible to adversarial attack. *Nature Medicine* **26**(3), 360–363.

Hochreiter, S. and J. Schmidhuber (1997). Long Short-Term Memory. *Neural Computation* **9**(8), 1735–1780.

Ilyas, A., S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry (2019). Adversarial Examples Are Not Bugs, They Are Features. In H. Wallach, H. Larochelle, A. Beygelzimer,

F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Volume 32.

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2017, may). ImageNet Classification with Deep Convolutional Neural Networks. *Commun. ACM* **60**(6), 84–90.

Kumar, D. and I. Ahmed (2011). *Seismic Noise*, pp. 1157–1161. Springer Netherlands.

Kurakin, A., I. J. Goodfellow, and S. Bengio (2017). Adversarial machine learning at scale. In *International Conference on Learning Representations*.

LeCun, Y., Y. Bengio, and G. Hinton (2015, May). Deep learning. *Nature* **521**(7553), 436–444.

LeCun, Y., B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel (1989). Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation* **1**(4), 541–551.

LeVeque, R. J., K. Waagan, F. I. González, D. Rim, and G. Lin (2016, Dec). Generating random earthquake events for probabilistic tsunami hazard assessment. *Pure and Applied Geophysics* **173**(12), 3671–3692.

Li, Z., N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar (2021). Fourier Neural Operator for Parametric Partial Differential Equations. In *International Conference on Learning Representations*.

Liu, C. M., D. Rim, R. Baraldi, and R. J. LeVeque (2021, Dec). Comparison of Machine Learning Approaches for Tsunami Forecasting from Sparse Observations. *Pure and Applied Geophysics* **178**(12), 5129–5153.

Lu, L., P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence* **3**(3), 218–229.

Madry, A., A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu (2018). Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.

Makinoshima, F., Y. Oishi, T. Yamazaki, T. Furumura, and F. Imamura (2021). Early forecasting of tsunami inundation from tsunami and geodetic observation data with convolutional neural networks. *Nature Communications* **12**(1), 2253.

Melgar, D. (2016). Cascadia FakeQuakes waveform data and scenario plots [Data set]. Journal of Geophysical Research. http://doi.org/10.5281/zenodo.59943.

Miller, J. P., R. Taori, A. Raghunathan, S. Sagawa, P. W. Koh, V. Shankar, P. Liang, Y. Carmon, and L. Schmidt (2021). Accuracy on the line: on the strong correlation between out-of-distribution and in-distribution generalization. In *Proceedings of the 38th International Conference on Machine Learning*, Volume 139 of *Proceedings of Machine Learning Research*, pp. 7721–7735. PMLR.

Mousavi, S. M. and G. C. Beroza (2022). Deep-learning seismology. *Science* **377**(6607), eabm4470.

Mulia, I. E., N. Ueda, T. Miyoshi, A. R. Gusman, and K. Satake (2022). Machine learning-based tsunami inundation prediction derived from offshore observations. *Nature Communications* **13**(1), 5489.

Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, Volume 32.

Rim, D., R. Baraldi, C. M. Liu, R. J. LeVeque, and K. Terada (2022). Tsunami Early Warning From Global Navigation Satellite System Data Using Convolutional Neural Networks. *Geophysical Research Letters* **49**(20), e2022GL099511.

Shafahi, A., M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein (2019). Adversarial training for free! In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Volume 32.

Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus (2014). Intriguing properties of neural networks. International Conference on Learning Representations.

Titov, V. V., F. I. Gonzalez, E. N. Bernard, M. C. Eble, H. O. Mofjeld, J. C. Newman, and A. J. Venturato (2005). Real-time tsunami forecasting: Challenges and solutions. *Natural Hazards* **35**(1), 35–41.

Tramér, F., A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel (2018). Ensemble Adversarial Training: Attacks and Defenses. In *International Conference on Learning Representations*.

Trefethen, L. N. and D. Bau (1997). *Numerical Linear Algebra*. Society of Industrial and Applied Mathematics.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin (2017). Attention Is All You Need. In *Advances in Neural Information Processing Systems*, Volume 30.

Zhang, H., Y. Yu, J. Jiao, E. Xing, L. E. Ghaoui, and M. Jordan (2019). Theoretically Principled Trade-off between Robustness and Accuracy. In K. Chaudhuri and R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, Volume 97 of *Proceedings of Machine Learning Research*, pp. 7472–7482.

Manuscript Received 00 Month 0000