



Contents lists available at [ScienceDirect](#)

## Neural Networks

journal homepage: [www.elsevier.com/locate/neunet](http://www.elsevier.com/locate/neunet)



### Highlights

#### **Bidirectionally self-normalizing neural networks**

*Neural Networks xxx (xxxx) xxx*

Yao Lu<sup>\*</sup>, Stephen Gould, Thalaiyasingam Ajanthan

- The vanishing/exploding gradient problem in training deep neural networks is addressed.
- The problem is provably solved under mild conditions using high-dimensional probability theory.
- Experiments show neural network of 200 layers can be trained without linearizing the networks.

**Graphical abstract and Research highlights will be displayed in online search result lists, the online contents list and the online article, but **will not appear in the article PDF file or print** unless it is mentioned in the journal specific style requirement. They are displayed in the proof pdf for review purpose only.**



Contents lists available at ScienceDirect

## Neural Networks

journal homepage: [www.elsevier.com/locate/neunet](http://www.elsevier.com/locate/neunet)

## Bidirectionally self-normalizing neural networks

Yao Lu<sup>a,b,\*</sup>, Stephen Gould<sup>a</sup>, Thalaiyasingam Ajanthan<sup>a,c</sup><sup>a</sup> Australian National University, Australia<sup>b</sup> Peking University, China<sup>c</sup> Amazon

## ARTICLE INFO

## Article history:

Received 11 October 2022

Received in revised form 9 August 2023

Accepted 11 August 2023

Available online xxxxx

## Keywords:

Neural networks

Vanishing/exploding gradient problem

Training

Optimization

## ABSTRACT

The problem of vanishing and exploding gradients has been a long-standing obstacle that hinders the effective training of neural networks. Despite various tricks and techniques that have been employed to alleviate the problem in practice, there still lacks satisfactory theories or provable solutions. In this paper, we address the problem from the perspective of high-dimensional probability theory. We provide a rigorous result that shows, under mild conditions, how the vanishing/exploding gradients problem disappears with high probability if the neural networks have sufficient width. Our main idea is to constrain both forward and backward signal propagation in a nonlinear neural network through a new class of activation functions, namely Gaussian–Poincaré normalized functions, and orthogonal weight matrices. Experiments on both synthetic and real-world data validate our theory and confirm its effectiveness on very deep neural networks when applied in practice.

© 2023 Elsevier Ltd. All rights reserved.

## 1. Introduction

Neural networks have brought unprecedented performance in various artificial intelligence tasks (Ciregan, Meier, & Schmidhuber, 2012; Graves, Mohamed, & Hinton, 2013; Krizhevsky, Sutskever, & Hinton, 2012; Silver et al., 2017). However, despite decades of research, training neural networks is still mostly guided by empirical observations and successful training often requires various heuristics and extensive hyperparameter tuning. It is therefore desirable to understand the cause of the difficulty in neural network training and to propose theoretically sound solutions.

A major difficulty is the vanishing/exploding gradients problem (Bengio, Simard, & Frasconi, 1994; Glorot & Bengio, 2010; Hochreiter, 1991; Pascanu, Mikolov, & Bengio, 2013; Philipp, Song, & Carbonell, 2018). That is, the norm of the gradient in each layer is either growing or shrinking at an exponential rate as the gradient signal is propagated from the top layer to bottom layer. For deep neural networks, this problem might cause numerical overflow and make the optimization problem intrinsically difficult, as the gradient in each layer has vastly different magnitude and therefore the optimization landscape becomes pathological. One might attempt to solve the problem by simply normalizing the gradient in each layer. Indeed, the adaptive gradient optimization methods (Duchi, Hazan, & Singer, 2011;

Kingma & Ba, 2015; Tieleman & Hinton, 2012) implement this idea and have been widely used in practice. However, one might also wonder if there is a solution more intrinsic to deep neural networks, whose internal structure if well-exploited would lead to further advances. For example, making the activation function more linear (Glorot, Bordes, & Bengio, 2011) and adding residual connections (He, Zhang, Ren, & Sun, 2016) might alleviate the problem, though they do not solve it in principle.

To enable the trainability of deep neural networks, batch normalization (Ioffe & Szegedy, 2015) was proposed in recent years and achieved widespread empirical success. Batch normalization is a differentiable operation which normalizes its inputs based on mini-batch statistics and is inserted between the linear and nonlinear layers. It is reported that batch normalization can accelerate neural network training significantly (Ioffe & Szegedy, 2015). However, batch normalization does not solve the vanishing/exploding gradients problem (Philipp et al., 2018). Indeed it is proved that batch normalization can actually worsen the problem (Yang, Pennington, Rao, Sohl-Dickstein, & Schoenholz, 2019). Besides, batch normalization requires separate training and testing phases and can be ineffective when the mini-batch size is small (Ioffe, 2017).

Alternatively, self-normalizing neural networks (Klambauer, Unterthiner, Mayr, & Hochreiter, 2017) and dynamical isometry theory (Pennington, Schoenholz, & Ganguli, 2017) were proposed to combat the vanishing/exploding gradients problem. In self-normalizing neural networks, a new activation function, scaled exponential linear unit (SELU), was devised to ensure the output of each unit to have zero mean and unit variance. In dynamical

\* Corresponding author at: Peking University, China.

E-mail addresses: [yaolubrain@gmail.com](mailto:yaolubrain@gmail.com) (Y. Lu), [stephen.gould@anu.edu.au](mailto:stephen.gould@anu.edu.au) (S. Gould), [thalaiyasingam.ajanthan@anu.edu.au](mailto:thalaiyasingam.ajanthan@anu.edu.au) (T. Ajanthan).

isometry theory, all singular values of the input–output Jacobian matrix are constrained to be close to one at initialization. This amounts to initializing the functionality of a neural network to be close to an orthogonal matrix. While the two theories dispense batch normalization, it is shown that neural networks with SELU still suffer from the vanishing/exploding gradients problem and dynamical isometry restricts the functionality of neural networks to be close to linear (pseudo-linearity) (Philipp et al., 2018).

In this paper, we follow the above line of research to investigate neural network trainability. Our contributions are three-fold: First, we propose a new type of neural networks that consist of orthogonal weight matrices and a new class of activation functions which we call Gaussian–Poincaré normalized (GPN) functions. We show many common activation functions can be easily transformed into their respective GPN versions. Second, we rigorously prove that the vanishing/exploding gradients problem disappears with high probability in the neural networks if the width of each layer is sufficiently large. Third, with experiments on synthetic and real-world data, we confirm that the vanishing/exploding gradients problem is solved to large extent in the neural networks while nonlinear functionality is maintained.

## 2. Theory

In this section, we introduce bidirectionally self-normalizing neural networks (BSNNs) formally and analyze its properties. All the proofs of our results are left to Appendix. To simplify the analysis, we define neural network in a restricted sense as the following.

**Definition 1 (Neural Network).** A neural network is a function from  $\mathbb{R}^d$  to  $\mathbb{R}^d$  that for  $l = 1, \dots, L$

$$\mathbf{h}^{(l)} = \mathbf{W}^{(l)}\mathbf{x}^{(l)}, \quad \mathbf{x}^{(l+1)} = \phi(\mathbf{h}^{(l)}), \quad (1)$$

where  $\mathbf{W}^{(l)} \in \mathbb{R}^{d \times d}$ ,  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  is a differentiable function applied element-wise to a vector,  $\mathbf{x}^{(1)} \in \mathbb{R}^d$  is the input and  $\mathbf{x}^{(L+1)} \in \mathbb{R}^d$  is the output.

Under this definition,  $\phi$  is called the activation function,  $\{\mathbf{W}^{(l)}\}_{l=1}^L$  are called the parameters,  $d$  is called the width and  $L$  is called the depth and superscript  $(l)$  denotes the  $l$ th layer of a neural network. The above formulation is similar to Pennington et al. (2017) but we omit the bias term in (1) for simplicity as it plays no role in our analysis.

Let  $E$  be the objective function of  $\{\mathbf{W}^{(l)}\}_{l=1}^L$  and  $\mathbf{D}^{(l)} = \text{diag}(\phi'(h_1^{(l)}), \dots, \phi'(h_d^{(l)}))$ , where  $\phi'$  denotes the derivative of  $\phi$  and  $h_i^{(l)}$  denotes the  $i$ th element of  $\mathbf{h}^{(l)}$ . Now, the error signal is back propagated via

$$\mathbf{y}^{(l)} = \mathbf{D}^{(l)} \frac{\partial E}{\partial \mathbf{x}^{(l+1)}}, \quad (2)$$

$$\mathbf{y}^{(l)} = \mathbf{D}^{(l)} (\mathbf{W}^{(l+1)})^T \mathbf{y}^{(l+1)}, \quad (3)$$

and the gradient of the weight matrix for layer  $l$  can be computed as

$$\frac{\partial E}{\partial \mathbf{W}^{(l)}} = \mathbf{y}^{(l)} (\mathbf{x}^{(l)})^T. \quad (4)$$

To solve the vanishing/exploding gradients problem, we constrain the forward signal  $\mathbf{x}^{(l)}$  and the backward signal  $\mathbf{y}^{(l)}$  in order to constrain the norm of the gradient. This leads to the following.

**Definition 2 (Bidirectional Self-Normalization).** A neural network is bidirectionally self-normalizing if

$$\|\mathbf{x}^{(1)}\|_2 = \dots = \|\mathbf{x}^{(L)}\|_2 = \sqrt{d}, \quad (5)$$

$$\|\mathbf{y}^{(1)}\|_2 = \dots = \|\mathbf{y}^{(L)}\|_2 = \left\| \frac{\partial E}{\partial \mathbf{x}^{(L+1)}} \right\|_2. \quad (6)$$

**Proposition 1.** If a neural network is bidirectionally self-normalizing, then

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(1)}} \right\|_F = \dots = \left\| \frac{\partial E}{\partial \mathbf{W}^{(L)}} \right\|_F. \quad (7)$$

In the rest of this section, we derive the conditions under which bidirectional self-normalization is achievable for a neural network.

### 2.1. Constraints on weight matrices

We constrain the weight matrices to be orthogonal since multiplication by an orthogonal matrix preserves the norm of a vector. For linear neural networks, this guarantees bidirectional self-normalization and its further benefits are discussed in Saxe, McClelland, and Ganguli (2014). Even for nonlinear neural networks, orthogonal constraints are shown to improve the trainability with proper scaling (Mishkin & Matas, 2016; Pennington et al., 2017).

### 2.2. Constraints on activation functions

To achieve bidirectional self-normalization for a nonlinear network, it is not enough only to constrain the weight matrices. We also need to constrain the activation function in such a way that both forward and backward signals are normalized. To this end, we propose the following constraint.

**Definition 3 (Gaussian–Poincaré Normalization).** Function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  is Gaussian–Poincaré normalized if it is differentiable and

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)^2] = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi'(x)^2] = 1. \quad (8)$$

The definition is inspired by the following theorem which shows the fundamental relationship between a function and its derivative under Gaussian measure.

**Theorem 1 (Gaussian–Poincaré Inequality).** If function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  is differentiable and has bounded  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)^2]$  and  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi'(x)^2]$ , then

$$\text{Var}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)] \leq \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi'(x)^2]. \quad (9)$$

The proof can be found in Bogachev (1998). Note that there is an implicit assumption that the input is approximately Gaussian for a Gaussian–Poincaré normalized (GPN) function. Even though this is standard in the literature (Klambauer et al., 2017; Pennington et al., 2017; Schoenholz, Gilmer, Ganguli, & Sohl-Dickstein, 2017), we will rigorously prove that this assumption is valid when orthogonal weight matrices are used in (1). Next, we state a property of GPN functions.

**Proposition 2.** Function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  is Gaussian–Poincaré normalized and  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)] = 0$  if and only if  $\phi(x) = x$  or  $\phi(x) = -x$ .

This result indicates that any nonlinear function with zero mean under Gaussian distribution (e.g., Tanh and SELU) is not GPN. Now we show that a large class of activation functions can be converted into their respective GPN versions using an affine transformation.

**Proposition 3.** For any differentiable function  $\phi: \mathbb{R} \rightarrow \mathbb{R}$  with non-zero and bounded  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)^2]$  and  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi'(x)^2]$ , there exist two constants  $a \in \mathbb{R}$  and  $b \in \mathbb{R}$  such that  $a\phi(x) + b$  is Gaussian–Poincaré normalized.



**Table 1**  
Constants for Gaussian–Poincaré normalization of activation functions.

	Tanh	ReLU	LeakyReLU	ELU	SELU	GELU
$a$	1.4674	1.4142	1.4141	1.2234	0.9660	1.4915
$b$	0.3885	0.0000	0.0000	0.0742	0.2585	−0.9097

To obtain  $a$  and  $b$ , one can use numerical procedure to compute the values of  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi'(x)^2]$ ,  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)^2]$  and  $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[\phi(x)]$  and then solve the quadratic equations

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[a^2 \phi'(x)^2] = 1, \quad (10)$$

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0,1)}[(a\phi(x) + b)^2] = 1. \quad (11)$$

We computed  $a$  and  $b$  (not unique) for several common activation functions (Clevert, Unterthiner, & Hochreiter, 2016; Hendrycks & Gimpel, 2016; Klambauer et al., 2017; Maas, Hannun, & Ng, 2013; Nair & Hinton, 2010) with their default hyperparameters<sup>1</sup> and the results are listed in Table 1. Note that ReLU, LeakyReLU and SELU are not differentiable at  $x = 0$  but they can be regarded as approximations of their smooth counterparts. We ignore such point and evaluate the integrals for  $x \in (-\infty, 0) \cup (0, \infty)$ .

With the orthogonal constraint on the weight matrices and the Gaussian–Poincaré normalization on the activation function, we prove that bidirectional self-normalization is achievable with high probability under mild conditions in the next subsection.

### 2.3. Norm-preservation theorems

The bidirectional self-normalization may not be achievable precisely in general unless the neural network is a linear one. Therefore, we investigate the properties of neural networks in a probabilistic framework. The random matrix theory and the high-dimensional probability theory allow us to characterize the behaviors of a large class of neural networks by its mean behavior, which is significantly simpler to analyze. Therefore, we study neural networks of random weights whose properties may shed light on the trainability of neural networks in practice.

First, we need a probabilistic version of the vector norm constraint.

**Definition 4 (Thin-Shell Concentration).** Random vector  $\mathbf{x} \in \mathbb{R}^d$  is thin-shell concentrated if for any  $\epsilon > 0$

$$\mathbb{P}\left\{\left|\frac{1}{d}\|\mathbf{x}\|_2^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (12)$$

as  $d \rightarrow \infty$ .

The definition is modified from the one in Bobkov (2003). Examples of thin-shell concentrated distributions include standard multivariate Gaussian and any distribution on the  $d$ -dimensional sphere of radius  $\sqrt{d}$ .

To prove the main results, *i.e.*, the norm-preservation theorems, we require the following assumptions.

#### Assumptions.

1. Random vector  $\mathbf{x} \in \mathbb{R}^d$  is thin-shell concentrated.
2. Random orthogonal matrix  $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_d)^T$  is uniformly distributed.
3. Function  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  is Gaussian–Poincaré normalized.
4. Function  $\phi$  and its derivative are Lipschitz continuous.

<sup>1</sup> We use  $\phi(x) = \max(0, x) + 0.01 \min(0, x)$  for LeakyReLU,  $\phi(x) = \max(0, x) + \min(0, \exp(x) - 1)$  for ELU and  $\phi(x) = x/(1 + \exp(-1.702x))$  for GELU.

The above assumptions are not restrictive. For Assumption 1, one can always normalize the input vectors of a neural network. For Assumption 2, orthogonal constraint or its relaxation has already been employed in neural network training (Brock, Lim, Ritchie, & Weston, 2017). Note, in Assumption 2, uniformly distributed means that  $\mathbf{W}$  is distributed under Haar measure, which is the unique rotation invariant probability measure on orthogonal matrix group. We refer the reader to Meckes (2019) for details. Furthermore, all the activation functions or their smooth counterparts listed in Table 1 satisfy Assumptions 3 and 4.

With the above assumptions, we can prove the following norm-preservation theorems.

**Theorem 2 (Forward Norm-Preservation).**

$$(\phi(\mathbf{w}_1^T \mathbf{x}), \dots, \phi(\mathbf{w}_d^T \mathbf{x})) \quad (13)$$

is thin-shell concentrated.

This result shows the transformation (orthogonal matrix followed by the GPN activation function) can preserve the norm of its input with high probability. Since the output is thin-shell concentrated, it serves as the input for the next layer and so on. Hence, the forward pass can preserve the norm of its input in each layer along the forward path when  $d$  is sufficiently large.

**Theorem 3 (Backward Norm-Preservation).** Let

$$\mathbf{D} = \text{diag}(\phi'(\mathbf{w}_1^T \mathbf{x}), \dots, \phi'(\mathbf{w}_d^T \mathbf{x})) \quad (14)$$

and  $\mathbf{y} \in \mathbb{R}^d$  be a fixed vector with bounded  $\|\mathbf{y}\|_\infty$ . Then for any  $\epsilon > 0$

$$\mathbb{P}\left\{\frac{1}{d}\|\mathbf{D}\mathbf{y}\|_2^2 - \|\mathbf{y}\|_2^2 \geq \epsilon\right\} \rightarrow 0 \quad (15)$$

as  $d \rightarrow \infty$ .

This result shows that the multiplication by the diagonal matrix  $\mathbf{D}$  preserves the norm of its input with high probability. Since orthogonal matrix  $\mathbf{W}$  also preserves the norm of its input, when the gradient error signal is propagated backwards as in (3), the norm is preserved in each layer along the backward path when  $d$  is sufficiently large.

Hence, combining Theorems 2 and 3, we proved that bidirectional self-normalization is achievable with high probability if the neural network is wide enough and the conditions in the Assumptions are satisfied. Then by Proposition 1, the vanishing/exploding gradients problem disappears with high probability.

**Sketch of the proofs.** The proofs of Theorems 2 and 3 are mainly based on a phenomenon in high-dimensional probability spaces, concentration of measure. We refer the reader to Vershynin (2018) for an introduction to the subject. Briefly, it can be shown that for some high-dimensional probability distributions, most mass is concentrated around certain range. For example, while most mass of a low-dimensional standard multivariate Gaussian distribution is concentrated around the center, most mass of a high-dimensional standard multivariate Gaussian distribution is concentrated around a thin-shell. Based on this phenomenon, it can be shown that random vector  $\mathbf{W}\mathbf{x}$  in high dimensions is approaching a random vector uniformly distributed on a sphere. Then the random vector uniformly distributed on a high-dimensional sphere is approximately Gaussian. Then the Gaussian random variables transformed by Lipschitz and GPN functions are subgaussian with unit variance. And the random vector of the subgaussian random variables has the concentration of norm property in high dimensions. Each of these steps is rigorously proved in Appendix.

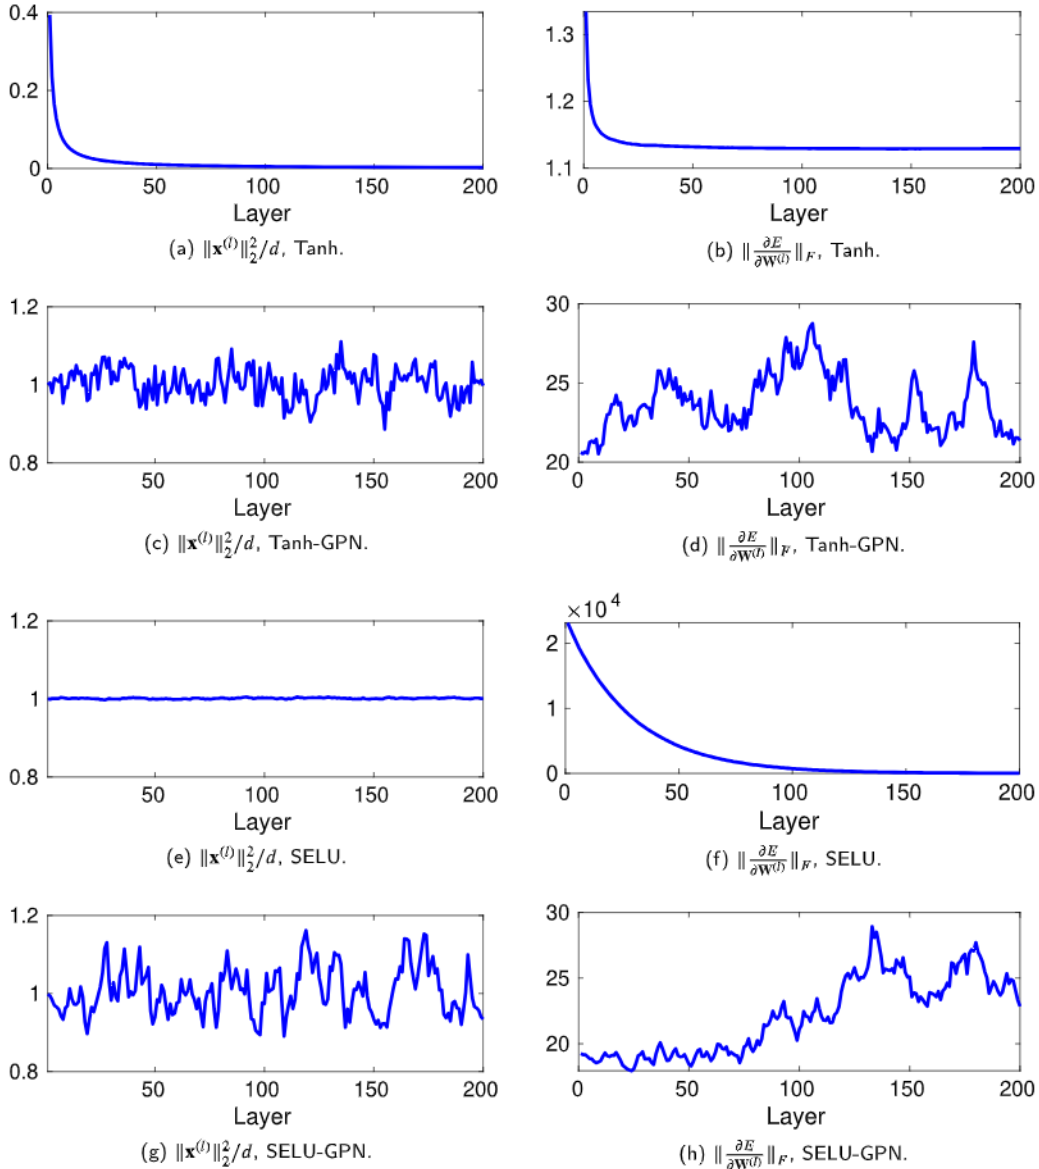


Fig. 1. Results on synthetic data with different activation functions. “-GPN” denotes the function is Gaussian–Poincaré normalized.  $\|\mathbf{x}^{(l)}\|_2$  denotes the  $l_2$  norm of the outputs of the  $l$ th layer.  $d$  denotes the width.  $\|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F$  is the Frobenius norm of the gradient of the weight matrix in the  $l$ th layer.

### 3. Experiments

We verify our theory on both synthetic and real-world data. In short, while very deep neural networks with non-GPN activations show vanishing/exploding gradients, GPN versions show stable gradients and improved trainability in both synthetic and real data. Furthermore, compared to dynamical isometry theory, BSNNs do not exhibit pseudo-linearity and maintain nonlinear functionality.

#### 3.1. Synthetic data

We create synthetic data to test the norm-preservation properties of the neural networks. The input  $\mathbf{x}^{(1)}$  is 500 data points of random standard Gaussian vectors of 500 dimensions. The gradient error  $\partial E / \partial \mathbf{x}^{(L+1)}$  is also random standard Gaussian vector of 500 dimensions. All the neural networks have depth 200. All the weight matrices are random orthogonal matrices uniformly generated. No training is performed.

In Fig. 1, we show the average norm of inputs and gradients of the neural networks of width 500. The standard deviation is not large and therefore not plotted for visual clarity. From the results, we can see that with GPN, the vanishing/exploding gradients problem is eliminated to large extent. The neural network with Tanh activation function does not show the vanishing/exploding gradients problem either. However,  $\|\mathbf{x}^{(l)}\|$  is close to zero for large  $l$  and each layer is close to a linear one since  $\text{Tanh}(x) \approx x$  when  $x \approx 0$  (pseudo-linearity), for which dynamical isometry is achieved.

One might wonder if bidirectional self-normalization has the same effect as dynamical isometry in solving the vanishing/exploding gradients problem, that is, to make the neural network close to an orthogonal matrix. To answer this question, we show the histogram of  $\phi'(h_i^{(l)})$  in Fig. 2. If the functionality of a neural network is close to an orthogonal matrix, since the weight matrices are orthogonal, then the values of  $\phi'(h_i^{(l)})$  would concentrate around one (Fig. 2(a)), which is not the case for BSNNs

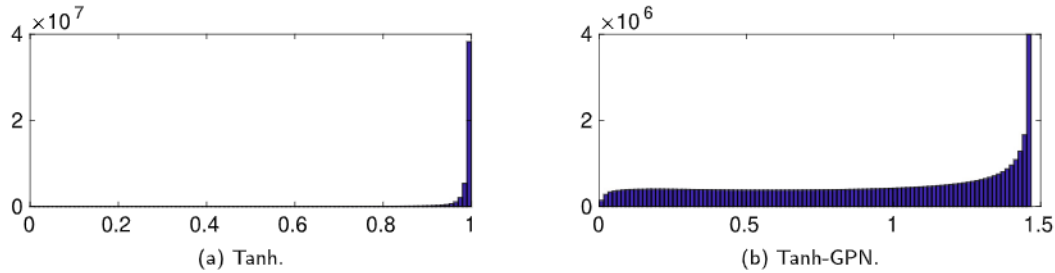


Fig. 2. Histogram of  $\phi'(h_i^{(l)})$ . The values of  $\phi'(h_i^{(l)})$  are accumulated for all units, all layers and all samples in the histogram.

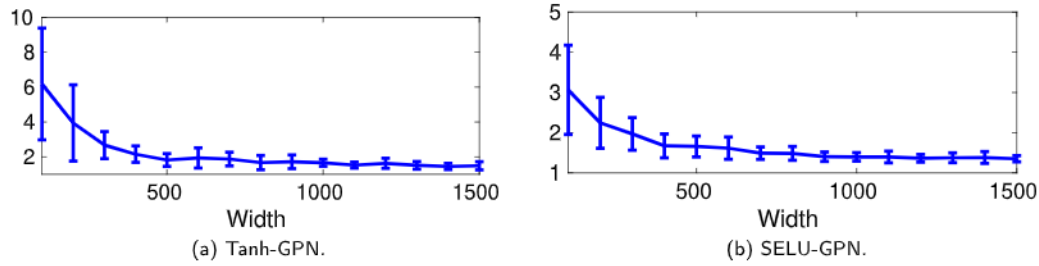


Fig. 3. Gradient norm ratio for different layer width on synthetic data. The ratio is  $\max_i \|\frac{\partial E}{\partial w_i}\|_F / \min_i \|\frac{\partial E}{\partial w_i}\|_F$ . The width ranges from 100 to 1500. The error bars show standard deviation.

Table 2

Accuracy (percentage) of neural networks of depth 200 with different activation functions on real-world data. The numbers in parenthesis denote the results when batch normalization is applied before the activation function.

	MNIST		CIFAR-10	
	Train	Test	Train	Test
Tanh	99.05 (87.39)	<b>96.57</b> (89.32)	80.84 (27.90)	<b>42.71</b> (29.32)
Tanh-GPN	<b>99.81</b> (84.93)	95.54 (87.11)	<b>96.39</b> (25.13)	40.95 (26.58)
ReLU	11.24 (11.24)	11.35 (11.42)	10.00 (10.00)	10.00 (10.00)
ReLU-GPN	<b>33.28</b> (11.42)	<b>28.13</b> (11.34)	<b>46.60</b> (10.09)	<b>34.96</b> (9.96)
LeakyReLU	11.24 (11.24)	11.35 (11.63)	10.00 (10.21)	10.00 (10.06)
LeakyReLU-GPN	<b>43.17</b> (11.19)	<b>49.28</b> (11.66)	<b>51.85</b> (9.89)	<b>39.38</b> (10.00)
ELU	99.06 (98.24)	95.41 ( <b>97.48</b> )	80.73 (42.39)	<b>45.76</b> (44.16)
ELU-GPN	<b>100.00</b> (97.86)	96.56 (96.69)	<b>99.37</b> (43.35)	43.12 (44.36)
SELU	99.86 (97.82)	97.33 (97.38)	29.23 (46.47)	29.55 (45.88)
SELU-GPN	<b>99.92</b> (97.91)	96.97 ( <b>97.39</b> )	<b>98.24</b> (47.74)	<b>45.90</b> (45.52)
GELU	11.24 (12.70)	11.35 (10.28)	10.00 (10.43)	10.00 (10.00)
GELU-GPN	<b>97.67</b> (11.22)	<b>95.82</b> (9.74)	<b>90.51</b> (10.00)	<b>36.94</b> (10.00)

(Fig. 2(b)). This shows that BSNNs do not suffer from the vanishing/exploding gradients problem while exhibiting nonlinear functionality.

In Fig. 3, we show the gradient norm of BSNNs with varying width. Notice, as the width increases, the norm of gradient in each layer of the neural network becomes more equalized, as predicted by our theory.

### 3.2. Real-world data

We run experiments on real-world image datasets MNIST (LeCun, Bottou, Bengio, & Haffner, 1998) and CIFAR-10 (Krizhevsky, Hinton, et al., 2009). The MNIST dataset contains 60 000 training images and 10 000 testing images in 10 classes. Each image is gray-scaled and has size  $28 \times 28$ . The CIFAR-10 dataset contains 50 000 training images and 10 000 testing images in 10 classes. Each image is colored and has size  $32 \times 32$ . The neural networks have width 500 and depth 200 (plus one unconstrained layer at bottom and one at top to fit the dimensionality of the input and output). We use stochastic gradient descent of momentum 0.5

with mini-batch size 64 and learning rate 0.0001. The training is run for 50 epochs for MNIST and 100 epochs for CIFAR-10. We do not use data augmentation. Since it is computationally expensive to enforce the orthogonality constraint, we simply constrain each row of the weight matrix to have  $l_2$  norm one as a relaxation of orthogonality by the following parametrization  $\mathbf{W} = (\mathbf{v}_1/\|\mathbf{v}_1\|_2, \dots, \mathbf{v}_d/\|\mathbf{v}_d\|_2)^T$  and optimize  $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_d)^T$  as an unconstrained problem.

We summarize the results in Table 2. We can see that, for activation functions ReLU, LeakyReLU and GELU, the neural networks are not trainable. But once these functions are GPN, the neural network can be trained. GPN activation functions consistently outperform their unnormalized counterparts in terms of the trainability, as the training accuracy is increased, but not necessarily generalization ability.

We show the test accuracy during training in Fig. 4, from which we can see the training is accelerated when SELU is GPN. ReLU, LeakyReLU and GELU, if not GPN, are completely untrainable due to the vanished gradients.

We observe that batch normalization leads to gradient explosion when combining with any of the activation functions. This confirms the claim of Philipp et al. (2018) and Yang et al. (2019) that batch normalization does not solve the vanishing/exploding gradients problem. On the other hand, without batch normalization the neural network with any GPN activation function has stable gradient magnitude throughout training. This indicates that BSNNs can dispense batch normalization and therefore avoid its shortcomings.

### 4. Discussion

We compare our theory to several most relevant theories in literature. A key distinguishing feature of our theory is that we provide rigorous proofs of the conditions under which the vanishing/exploding gradients problem disappears. To the best of our knowledge, this is the first time that the problem is provably solved for nonlinear neural networks.

Self-normalizing neural networks enforce zero mean and unit variance for the output of each unit with the SELU activation function (Klambauer et al., 2017). However, as pointed out in Philipp



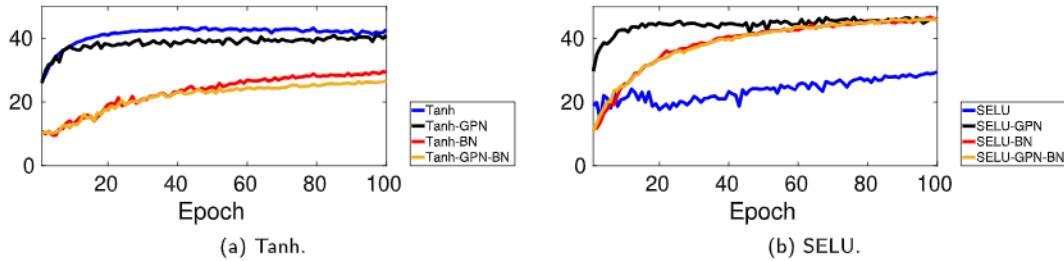


Fig. 4. Test accuracy (percentage) during training on CIFAR-10. “-BN” denotes that batch normalization is applied before the activation function.

et al. (2018), only constraining forward signal propagation does not solve the vanishing/exploding gradients problem since the norm of the backward signal can grow or shrink. In Philipp et al. (2018) and our experiments, SELU is indeed shown to cause gradient exploding. To solve the problem, the signal propagation in both directions need to be constrained, as in our theory. However, our work lacks of the attractive fixed-point property of SELU (Klambauer et al., 2017), which deserves the future investigation.

Our theory is largely developed from the deep signal propagation theory (Poole, Lahiri, Raghu, Sohl-Dickstein, & Ganguli, 2016; Schoenholz et al., 2017). Both theories require  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = 1$ . However, ours also requires the quantity  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2]$  to be one while in Poole et al. (2016) and Schoenholz et al. (2017) it can be an arbitrary positive number. We emphasize that it is desirable to enforce  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = 1$  to avoid trivial solutions. For example, if  $\phi(x) = \text{Tanh}(\epsilon x)$  with  $\epsilon \approx 0$ , then  $\phi(\epsilon x) \approx \epsilon x$  and the neural network becomes essentially a linear one for which depth is unnecessary (pseudo-linearity Philipp et al., 2018). This is observed in Fig. 1(a). Moreover, in Poole et al. (2016) and Schoenholz et al. (2017) the signal propagation analysis is done based on random weights under i.i.d. Gaussian distribution whereas we proved how one can solve the vanishing/exploding gradients problem assuming the weight matrices are orthogonal and uniformly distributed under Haar measure.

Dynamical isometry theory (Pennington et al., 2017) enforces the Jacobian matrix of the input-output function of a neural network to have all singular values close to one. Since the weight matrices are constrained to be orthogonal, it is equivalent to enforce each  $\mathbf{D}^{(l)}$  in (3) to be close to the identity matrix, which implies the functionality of neural network at initialization is close to an orthogonal matrix (pseudo-linearity). This indeed enables trainability since linear neural networks with orthogonal weight matrices do not suffer from the vanishing/exploding gradients problem. As neural networks need to learn a nonlinear input-output functionality to solve certain tasks, during training the weights of a neural network are unconstrained so that the neural network would move to a nonlinear region where the vanishing/exploding gradients problem might return. In our theory, although the orthogonality of weight matrices is also required, we approach the problem from a different perspective. We do not encourage the linearity at initialization. The neural network can be initialized to be nonlinear and stay nonlinear during the training even when the weights are constrained. This is shown in Section 3.1.

Our work is also related to weight normalization (Salimans & Kingma, 2016) as we relax the orthogonality for computational efficiency during training.

Although random matrices of size  $N \times N$  whose elements are i.i.d sampled from Gaussian  $\mathcal{N}(0, \frac{1}{N})$  are orthogonal in expectation, their condition number increases as  $N \rightarrow \infty$ , which means the sampled random Gaussian matrices do not enjoy the strict norm-preserving properties as orthogonal matrices. Therefore, the orthogonality is crucial in our theory.

A limitation of our empirical evaluation is that we reported the results on real-world datasets for only single trials of training in Table 2, instead of running the training for multiples trials under different random seeds and providing error bars. This is mainly due to that the overall training takes over 100 GPU hours and the limited computational resources we have. However, we provide the code in the online material for reproducing our experiments.

Future work includes extending our theory to more sophisticated networks such as convolutional architectures as well as investigating the generalization capabilities of BSNNs.

### Acknowledgements

This work was funded by the Australian Research Council Centre of Excellence for Robotic Vision (CE140100016). Yao Lu was supported by a CSIRO/Data61 scholarship.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

The data is publicly available online.

### Appendix. Proofs

**Proposition 1.** If a neural network is bidirectionally self-normalizing, then

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(1)}} \right\|_F = \dots = \left\| \frac{\partial E}{\partial \mathbf{W}^{(L)}} \right\|_F. \quad (16)$$

**Proof.** For each  $l$ , we have

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right\|_F = \sqrt{\text{trace} \left( \frac{\partial E}{\partial \mathbf{W}^{(l)}} \left( \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right)^T \right)} \quad (17)$$

$$= \sqrt{\text{trace}(\mathbf{y}^{(l)}(\mathbf{x}^{(l)})^T \mathbf{x}^{(l)}(\mathbf{y}^{(l)})^T)} \quad (18)$$

$$= \sqrt{\text{trace}((\mathbf{x}^{(l)})^T \mathbf{x}^{(l)}(\mathbf{y}^{(l)})^T \mathbf{y}^{(l)})} \quad (19)$$

$$= \sqrt{(\mathbf{x}^{(l)})^T \mathbf{x}^{(l)}} \sqrt{(\mathbf{y}^{(l)})^T \mathbf{y}^{(l)}} \quad (20)$$

$$= \|\mathbf{x}^{(l)}\|_2 \|\mathbf{y}^{(l)}\|_2. \quad (21)$$

By the definition of bidirectional self-normalization, we have

$$\frac{\partial E}{\partial \mathbf{W}^{(1)}} \Big|_F = \dots = \frac{\partial E}{\partial \mathbf{W}^{(L)}} \Big|_F. \quad (22)$$

**Proposition 2.** Function  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  is Gaussian-Poincaré normalized and  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] = 0$  if and only if  $\phi(x) = x$  or  $\phi(x) = -x$ .

**Proof.**  $\phi(x)$  and  $\phi'(x)$  can be expanded in terms of Hermite polynomials since  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] < \infty$  and  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] < \infty$ . Let the Hermite polynomial of degree  $k$  be

$$H_k(x) = \frac{(-1)^k}{\sqrt{k!}} \exp\left(\frac{x^2}{2}\right) \frac{d^k}{dx^k} \exp\left(-\frac{x^2}{2}\right) \quad (23)$$

and due to  $H'_k(x) = \sqrt{k}H_{k-1}(x)$ , we have

$$\phi(x) = \sum_{k=0}^{\infty} a_k H_k(x), \quad (24)$$

$$\phi'(x) = \sum_{k=1}^{\infty} \sqrt{k} a_k H_{k-1}(x) \quad (25)$$

where  $a_k$  denotes the coefficient that

$$a_k = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[H_k(x)\phi(x)]. \quad (26)$$

Since  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] = 0$ , we have

$$a_0 = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[H_0(x)\phi(x)] \quad (27)$$

$$= \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] \quad (28)$$

$$= 0. \quad (29)$$

Since

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = 1 \quad (30)$$

and Hermite polynomials are orthonormal, we have

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \sum_{k=1}^{\infty} a_k^2 \quad (31)$$

$$= \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] \quad (32)$$

$$= \sum_{k=1}^{\infty} k a_k^2 = 1. \quad (33)$$

Therefore, we have

$$\sum_{k=1}^{\infty} k a_k^2 - \sum_{k=1}^{\infty} a_k^2 = 0 \quad (34)$$

that is

$$\sum_{k=2}^{\infty} (k-1) a_k^2 = 0. \quad (35)$$

Since each term in  $\sum_{k=2}^{\infty} (k-1) a_k^2$  is nonnegative, the only solution is  $a_k = 0$  for  $k \geq 2$ . And since  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = a_1^2 = 1$ , we have  $a_1 = \pm 1$ . Hence,  $\phi(x) = \pm H_1(x) = \pm x$ .

On the other hand, if  $\phi(x) = \pm x$ , then

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[x^2] = 1 \quad (36)$$

and

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[1] = 1. \quad (37)$$

**Proposition 3.** For any differentiable function  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  with non-zero and bounded  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2]$  and  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2]$ , there exist two constants  $a \in \mathbb{R}$  and  $b \in \mathbb{R}$  such that  $a\phi(x) + b$  is Gaussian–Poincaré normalized.

**Proof.** Let  $\varphi(x) = \phi(x) + c$  where  $c \in \mathbb{R}$  is a constant. We first prove that there exists  $c$  such that  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\varphi(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\varphi'(x)^2]$  and then prove the proposition. Let

$$\psi(c) = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\varphi(x)^2] \quad (38)$$

$$- \mathbb{E}_{x \sim \mathcal{N}(0,1)}[(\phi'(x))^2] \quad (39)$$

$$= \text{Var}_{x \sim \mathcal{N}(0,1)}[\varphi(x)] \quad (40)$$

$$+ (\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\varphi(x)])^2 \quad (41)$$

$$- \mathbb{E}_{x \sim \mathcal{N}(0,1)}[(\phi'(x))^2] \quad (42)$$

$$= \text{Var}_{x \sim \mathcal{N}(0,1)}[\phi(x)] \quad (43)$$

$$+ (\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] + c)^2 \quad (44)$$

$$- \mathbb{E}_{x \sim \mathcal{N}(0,1)}[(\phi'(x))^2]. \quad (45)$$

Therefore,  $\psi(c)$  is a quadratic function of  $c$ . We also have  $\psi(c) > 0$  as  $c \rightarrow \infty$  and  $\psi(-\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)]) \leq 0$  due to Gaussian–Poincaré inequality. Hence, there exists  $c$  for which  $\psi(c) = 0$  such that  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[(\phi(x) + c)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2]$ . Let  $a = (\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2])^{-1/2}$  and  $b = ac$ , we have  $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[(a\phi(x) + b)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[(a\phi'(x))^2] = 1$ .

The proof is largely due to Eldredge (2020) with minor modification in here.

**Assumptions.**

1. Random vector  $\mathbf{x} \in \mathbb{R}^d$  is thin-shell concentrated.
2. Random orthogonal matrix  $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_d)^T$  is uniformly distributed.
3. Function  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  is Gaussian–Poincaré normalized.
4. Function  $\phi$  and its derivative are Lipschitz continuous.

**Theorem 2** (Forward Norm-Preservation).

$$(\phi(\mathbf{w}_1^T \mathbf{x}), \dots, \phi(\mathbf{w}_d^T \mathbf{x})) \quad (46)$$

is thin-shell concentrated.

**Theorem 3** (Backward Norm-Preservation). Let

$$\mathbf{D} = \text{diag}(\phi'(\mathbf{w}_1^T \mathbf{x}), \dots, \phi'(\mathbf{w}_d^T \mathbf{x})) \quad (47)$$

and  $\mathbf{y} \in \mathbb{R}^d$  be a fixed vector with bounded  $\|\mathbf{y}\|_{\infty}$ . Then for any  $\epsilon > 0$

$$\mathbb{P}\left\{\frac{1}{d} \|\mathbf{D}\mathbf{y}\|_2^2 - \|\mathbf{y}\|_2^2 \geq \epsilon\right\} \rightarrow 0 \quad (48)$$

as  $d \rightarrow \infty$ .

*Notations.*  $\mathbb{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 = 1\}$ .  $\mathbb{O}(d)$  is the orthogonal matrix group of size  $d$ .  $\mathbf{1}_{\{\cdot\}}$  denotes the indicator function.  $\mathbf{0}_d$  denotes the vector of dimension  $d$  and all elements equal to zero.  $\mathbf{I}_d$  denotes the identity matrix of size  $d \times d$ .

**Lemma 1.** If random variable  $x \sim \mathcal{N}(0, 1)$  and function  $f : \mathbb{R} \rightarrow \mathbb{R}$  is Lipschitz continuous, then random variable  $f(x)$  is sub-Gaussian.

**Proof.** Due to the Gaussian concentration theorem (Theorem 5.2.2 in Vershynin (2018)), we have

$$\|f(x) - \mathbb{E}[f(x)]\|_{\psi_2} \leq CK \quad (49)$$

where  $\|\cdot\|_{\psi_2}$  denotes sub-Gaussian norm,  $C$  is a constant and  $K$  is the Lipschitz constant of  $f$ . This implies  $f(x) - \mathbb{E}[f(x)]$  is sub-Gaussian (Proposition 2.5.2 in Vershynin (2018)). Therefore  $f(x)$  is sub-Gaussian (Lemma 2.6.8 in Vershynin (2018)).

**Lemma 2.** Let  $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$  be a random vector that each coordinate  $x_i$  is independent and sub-Gaussian and  $\mathbb{E}[x_i^2] = 1$ . Let  $\mathbf{y} = (y_1, \dots, y_d) \in \mathbb{R}^d$  be a fixed vector with bounded  $\|\mathbf{y}\|_{\infty}$ . Then

$$\mathbb{P}\left\{\frac{1}{d} \left| \sum_i x_i^2 y_i^2 - \sum_i y_i^2 \right| \geq \epsilon\right\} \rightarrow 0 \quad (50)$$

as  $d \rightarrow \infty$ .



**Proof.** Since  $y_i x_i$  is sub-Gaussian, then  $y_i^2 x_i^2$  is sub-exponential (Lemma 2.7.6 in Vershynin (2018)). Since  $\mathbb{E}[y_i^2 x_i^2] = y_i^2 \mathbb{E}[x_i^2] = y_i^2$ ,  $y_i^2 x_i^2 - y_i^2$  is sub-exponential with zero mean (Exercise 2.7.10 in Vershynin (2018)). Applying Bernstein's inequality (Corollary 2.8.3 in Vershynin (2018)), we proved the lemma.

**Lemma 3.** Let  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ . Then for any  $0 < \delta < 1$

$$\mathbb{P}\{\mathbf{z} \in \mathbb{R}^d : (1 - \delta)\sqrt{d} \leq \|\mathbf{z}\|_2 \leq (1 + \delta)\sqrt{d}\} \quad (51)$$

$$\geq 1 - 2 \exp(-d\delta^2). \quad (52)$$

**Proof.** See Theorem 1.2 in Alberts and Khoshnevisan (2018).

**Lemma 4.** Let  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ . Then  $\mathbf{z}/\|\mathbf{z}\|_2$  is uniformly distributed on  $\mathbb{S}^{d-1}$ .

**Proof.** See Dawkins (2016).

**Lemma 5.** Let  $\mathbf{z} = (z_1, \dots, z_d) \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ ,  $\mathbf{a} = (a_1, \dots, a_d)$  be a fixed vector with bounded  $\|\mathbf{a}\|_\infty$  and  $f : \mathbb{R} \rightarrow \mathbb{R}$  be a continuous function. Then for any  $\epsilon > 0$

$$\mathbb{P}\left\{\frac{1}{d} \left| \sum_i a_i f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - \sum_i a_i f(z_i) \right| \geq \epsilon \right\} \rightarrow 0 \quad (53)$$

as  $d \rightarrow \infty$ .

**Proof.** To prove this proposition, first note that since

$$\frac{1}{d} \left| \sum_i a_i f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - \sum_i a_i f(z_i) \right| \quad (54)$$

$$\leq \frac{1}{d} \sum_i |a_i| \cdot |f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - f(z_i)|, \quad (55)$$

if

$$\mathbb{P}\left\{\frac{1}{d} \sum_i |a_i| \cdot |f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - f(z_i)| \geq \epsilon \right\} \rightarrow 0, \quad (56)$$

as  $d \rightarrow \infty$ , then

$$\mathbb{P}\left\{\frac{1}{d} \left| \sum_i a_i f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - \sum_i a_i f(z_i) \right| \geq \epsilon \right\} \rightarrow 0. \quad (57)$$

Therefore, we only need to prove (56). For  $0 < \delta < 1$ , let

$$A = \left\{ \mathbf{z} \in \mathbb{R}^d : \frac{1}{d} \sum_i |a_i| \cdot |f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - f(z_i)| \geq \epsilon \right\}, \quad (58)$$

$$\mathcal{U}_\delta = \left\{ \mathbf{z} \in \mathbb{R}^d : (1 - \delta)\sqrt{d} \leq \|\mathbf{z}\|_2 \leq (1 + \delta)\sqrt{d} \right\}. \quad (59)$$

Then

$$\mathbb{P}\left\{\frac{1}{d} \sum_i |a_i| \cdot |f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - f(z_i)| \geq \epsilon \right\} \quad (60)$$

$$= \int_{\mathbb{R}^d} \mathbf{1}_{\{z \in A\}} \gamma(\mathbf{z}) d\mathbf{z} \quad (61)$$

$$= \int_{\mathbb{R}^d \setminus \mathcal{U}_\delta} \mathbf{1}_{\{z \in A\}} \gamma(\mathbf{z}) d\mathbf{z} + \int_{\mathcal{U}_\delta} \mathbf{1}_{\{z \in A\}} \gamma(\mathbf{z}) d\mathbf{z} \quad (62)$$

where  $\gamma(\mathbf{z})$  denotes the density function of  $\mathbf{z}$ .

Let  $\delta = d^{-1/4}$ . From Lemma 3, we have, as  $d \rightarrow \infty$ ,

$$\int_{\mathbb{R}^d \setminus \mathcal{U}_\delta} \mathbf{1}_{\{z \in A\}} \gamma(\mathbf{z}) d\mathbf{z} \quad (63)$$

$$\leq \int_{\mathbb{R}^d \setminus \mathcal{U}_\delta} \gamma(\mathbf{z}) d\mathbf{z} = 1 - \mathbb{P}\{\mathbf{z} \in \mathcal{U}_\delta\} \quad (64)$$

$$\leq 2 \exp(-d\delta^2) \rightarrow 0. \quad (65)$$

For  $\mathbf{z} \in \mathcal{U}_\delta$  and  $\delta = d^{-1/4}$ , we have  $\|\mathbf{z}\|_2 \rightarrow \sqrt{d}$ ,  $\sqrt{d}/\|\mathbf{z}\|_2 z_i \rightarrow z_i$  and therefore  $f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) \rightarrow f(z_i)$  as  $d \rightarrow \infty$ . Since  $|a_i|$  is bounded, we have  $\frac{1}{d} \sum_i |a_i| \cdot |f(\sqrt{d}/\|\mathbf{z}\|_2 z_i) - f(z_i)| \rightarrow 0$  and therefore  $\int_{\mathcal{U}_\delta} \mathbf{1}_{\{z \in A\}} d\mathbf{z} \rightarrow 0$ , as  $d \rightarrow \infty$ .

**Lemma 6.** Let random matrix  $\mathbf{W}$  be uniformly distributed on  $\mathbb{O}(d)$  random vector  $\theta$  be uniformly distributed on  $\mathbb{S}^{d-1}$  and random vector  $\mathbf{x} \in \mathbb{R}^d$  be thin-shell concentrated. Then  $\mathbf{W}\mathbf{x} \rightarrow \sqrt{d}\theta$  as  $d \rightarrow \infty$ .

**Proof.** Let  $\mathbf{y} \in \mathbb{R}^d$  be any vector with  $\|\mathbf{y}\|_2 = \sqrt{d}$  and  $\mathbf{a} = (\sqrt{d}, 0, \dots, 0) \in \mathbb{R}^d$ . Since  $\mathbf{W}$  is uniformly distributed,  $\mathbf{W}\mathbf{y}$  has the same distribution as  $\mathbf{W}\mathbf{a}$ .  $\mathbf{W}\mathbf{a}$  is the first column of  $\sqrt{d}\mathbf{W}$ , which is equivalent to random vector  $\sqrt{d}\theta$  (Meckes, 2019). Since  $\mathbf{x}$  is thin-shell concentrated,  $\mathbf{x} \rightarrow \sqrt{d}/\|\mathbf{x}\|_2 \mathbf{x} = \mathbf{y}$  and therefore  $\mathbf{W}\mathbf{x} \rightarrow \sqrt{d}\theta$  as  $d \rightarrow \infty$ .

**Proof of Theorem 2.** Let  $\mathbf{z} = (z_1, \dots, z_d) \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ . Due to Lemma 1, random variable  $\phi(z_i)$  is sub-Gaussian. Since  $\phi$  is Gaussian-Poincaré normalized,  $\mathbb{E}_{z_i \sim \mathcal{N}(0,1)}[\phi(z_i)^2] = 1$ . Applying Lemma 2 with each  $y_i = 1$ , we have for  $\epsilon > 0$

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i \phi(z_i)^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (66)$$

as  $d \rightarrow \infty$ .

Due to Lemmas 4 and 5 (with each  $a_i = 1$ ), for random vector  $\theta = (\theta_1, \dots, \theta_d)$  uniformly distributed on  $\mathbb{S}^{d-1}$ , we have

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i \phi(\sqrt{d}\theta_i)^2 - \frac{1}{d} \sum_i \phi(z_i)^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (67)$$

and therefore

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i \phi(\sqrt{d}\theta_i)^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (68)$$

as  $d \rightarrow \infty$ .

Then from Lemma 6, we have  $\mathbf{W}\mathbf{x} \rightarrow \sqrt{d}\theta$  and therefore

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i \phi(\mathbf{w}_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (69)$$

as  $d \rightarrow \infty$ .

**Proof of Theorem 3.** Let  $\mathbf{z} = (z_1, \dots, z_d) \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$ . Due to Lemma 1, random variable  $\phi'(z_i)$  is sub-Gaussian. Since  $\phi$  is Gaussian-Poincaré normalized,  $\mathbb{E}_{z_i \sim \mathcal{N}(0,1)}[\phi'(z_i)^2] = 1$ . Applying Lemma 2, we have for  $\epsilon > 0$

$$\mathbb{P}\left\{\frac{1}{d} \left| \sum_i y_i^2 \phi'(z_i)^2 - y_i^2 \right| \geq \epsilon\right\} \rightarrow 0 \quad (70)$$

as  $d \rightarrow \infty$ .

Due to Lemmas 4 and 5 (with each  $a_i = y_i^2$ ), for random vector  $\theta = (\theta_1, \dots, \theta_d)$  uniformly distributed on  $\mathbb{S}^{d-1}$ , we have

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i y_i^2 \phi'(\sqrt{d}\theta_i)^2 - \frac{1}{d} \sum_i y_i^2 \phi'(z_i)^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (71)$$

and therefore

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i y_i^2 \phi'(\sqrt{d}\theta_i)^2 - y_i^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (72)$$

as  $n \rightarrow \infty$ .

Then from Lemma 6, we have  $\mathbf{W}\mathbf{x} \rightarrow \sqrt{d}\theta$  and therefore

$$\mathbb{P}\left\{\left|\frac{1}{d} \sum_i y_i^2 \phi'(\mathbf{w}_i^T \mathbf{x})^2 - y_i^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (73)$$

as  $d \rightarrow \infty$ .

## References

- Alberts, T., & Khoshnevisan, D. (2018). Calculus on Gauss space: An introduction to Gaussian analysis.
- Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*.
- Bobkov, S. G. (2003). On concentration of distributions of random weighted sums. *The Annals of Probability*.
- Bogachev, V. I. (1998). *Gaussian measures*. American Mathematical Society.
- Brock, A., Lim, T., Ritchie, J. M., & Weston, N. (2017). Neural photo editing with introspective adversarial networks. In *International conference on learning representations*.
- Ciregan, D., Meier, U., & Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. In *IEEE conference on computer vision and pattern recognition*.
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (elus). In *International conference on learning representations*.
- Dawkins, J. (2016). Normalized vector of Gaussian variables is uniformly distributed on the sphere. In *Math stackexchange*. <https://math.stackexchange.com/q/1864563>.
- Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*.
- Eldredge, N. (2020). When do  $\phi^2$  and  $\phi'^2$  have the same expectation under a Gaussian random variable? MathOverflow, <https://mathoverflow.net/q/351553>.
- Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feed-forward neural networks. In *International conference on artificial intelligence and statistics*.
- Glorot, X., Bordes, A., & Bengio, Y. (2011). Deep sparse rectifier neural networks. In *International conference on artificial intelligence and statistics*.
- Graves, A., Mohamed, A.-r., & Hinton, G. (2013). Speech recognition with deep recurrent neural networks. In *IEEE international conference on acoustics, speech and signal processing*.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*.
- Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). arXiv preprint [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- Hochreiter, S. (1991). *Untersuchungen zu dynamischen neuronalen Netzen: Diploma*, Technische Universität München.
- Ioffe, S. (2017). Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. *Advances in Neural Information Processing Systems*.
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations*.
- Klambauer, G., Unterthiner, T., Mayr, A., & Hochreiter, S. (2017). Self-normalizing neural networks. *Advances in Neural Information Processing Systems*.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*.
- Maas, A. L., Hannun, A. Y., & Ng, A. Y. (2013). Rectifier nonlinearities improve neural network acoustic models. In *International conference on machine learning*.
- Meckes, E. S. (2019). *The random matrix theory of the classical compact groups*. Cambridge University Press.
- Mishkin, D., & Matas, J. (2016). All you need is a good init. In *International conference on learning representations*.
- Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *International conference on machine learning*.
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *International conference on machine learning*.
- Pennington, J., Schoenholz, S., & Ganguli, S. (2017). Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. *Advances in Neural Information Processing Systems*.
- Philipp, G., Song, D., & Carbonell, J. G. (2018). The exploding gradient problem demystified-definition, prevalence, impact, origin, tradeoffs, and solutions. arXiv preprint [arXiv:1712.05577](https://arxiv.org/abs/1712.05577).
- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., & Ganguli, S. (2016). Exponential expressivity in deep neural networks through transient chaos. *Advances in Neural Information Processing Systems*.
- Salimans, T., & Kingma, D. P. (2016). Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in Neural Information Processing Systems*.
- Saxe, A. M., McClelland, J. L., & Ganguli, S. (2014). Exact solutions to the non-linear dynamics of learning in deep linear neural networks. In *International conference on learning representations*.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., & Sohl-Dickstein, J. (2017). Deep information propagation. In *International conference on learning representations*.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*.
- Tieleman, T., & Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. In *Neural networks for machine learning*. COURSERA.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*. Cambridge University Press.
- Yang, G., Pennington, J., Rao, V., Sohl-Dickstein, J., & Schoenholz, S. S. (2019). A mean field theory of batch normalization. In *International conference on learning representations*.