
Bidirectional Self-Normalizing Neural Networks

Yao Lu Stephen Gould Thalaiyasingam Ajanthan
Australian National University
{yao.lu,stephen.gould,thalaiyasingam.ajanthan}@anu.edu.au

Abstract

The problem of exploding and vanishing 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 exploding/vanishing gradient 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.

1 Introduction

Neural networks have brought unprecedented performance in various artificial intelligence tasks [10, 17, 28]. 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 gradient exploding/vanishing problem [9, 12, 22, 24]. 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 [7, 15, 29] 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.

To enable the trainability of deep neural networks, batch normalization [14] 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 inserted between the linear and nonlinear layers. It is reported that batch normalization can accelerate neural network training significantly [14]. However, batch normalization does not solve the gradient exploding/vanishing problem [24]. Indeed it is proved that batch normalization can actually worsen the problem [32]. Besides, batch normalization requires separate training and testing phases and is ineffective when the mini-batch size is small [13]. The shortcomings of batch normalization motivate us to search for a more principled and generic approach to solve the gradient exploding/vanishing problem.

Alternatively, self-normalizing neural networks [16] and dynamical isometry theory [23] were proposed to combat gradient exploding/vanishing problem. In self-normalizing neural networks, the output of each network unit is constrained to have zero mean and unit variance. Based on this motivation, a new activation function, scaled exponential linear unit (SELU), was devised. In dynamical 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 network to be close to an orthogonal matrix. While the two theories dispense batch normalization, it is shown that SELU still suffers from exploding/vanishing gradient problem and dynamical isometry restricts the functionality of the network to be close to linear (pseudo-linearity) [24].

In this paper, we follow the above line of research to investigate neural network trainability. Our contributions are three-fold: First, we introduce bidirectional self-normalizing neural network (BSNN) 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 gradient exploding/vanishing problem disappears with high probability in BSNNs if the width of each layer is sufficiently large. Third, with experiments on synthetic and real-world data, we confirm that BSNNs solve the gradient vanishing/exploding problem to large extent while maintaining nonlinear functionality.

2 Theory

In this section, we introduce self-normalizing neural network (BSNN) 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}^n to \mathbb{R}^n 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 $\mathbf{W}^{(l)} \in \mathbb{R}^{n \times n}$, $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is a differentiable function applied element-wise to a vector, $\mathbf{x}^{(1)}$ is the input and $\mathbf{x}^{(L+1)}$ is the output.

Under this definition, ϕ is called the activation function, $\{\mathbf{W}^{(l)}\}_{l=1}^L$ are called the parameters, n 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.* [23] 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 $\mathbf{D}^{(l)}$ be the diagonal matrix with diagonal elements $D_{ii}^{(l)} = \phi'(h_i^{(l)})$, where ϕ' denotes the derivative of ϕ . Now, the error signal is back propagated via

$$\mathbf{d}^{(L+1)} = \frac{\partial E}{\partial \mathbf{x}^{(L+1)}}, \quad \mathbf{d}^{(l)} = \mathbf{D}^{(l)} \left(\prod_{k=l+1}^L \mathbf{W}^{(k)} \mathbf{D}^{(k)} \right) \mathbf{d}^{(L+1)}, \quad (2)$$

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

$$\frac{\partial E}{\partial \mathbf{W}^{(l)}} = \mathbf{d}^{(l)} (\mathbf{x}^{(l)})^T. \quad (3)$$

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

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

$$\|\mathbf{x}^{(1)}\|_2 = \|\mathbf{x}^{(2)}\|_2 = \dots = \|\mathbf{x}^{(L+1)}\|_2 = \sqrt{n}, \quad (4)$$

$$\|\mathbf{d}^{(1)}\|_2 = \|\mathbf{d}^{(2)}\|_2 = \dots = \|\mathbf{d}^{(L+1)}\|_2 = \left\| \frac{\partial E}{\partial \mathbf{x}^{(L+1)}} \right\|_2. \quad (5)$$

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

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(1)}} \right\|_F = \left\| \frac{\partial E}{\partial \mathbf{W}^{(2)}} \right\|_F = \dots = \left\| \frac{\partial E}{\partial \mathbf{W}^{(L)}} \right\|_F. \quad (6)$$

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 networks, this guarantees bidirectional self-normalization and its further benefits are discussed in Saxe *et al.* [26]. Even for nonlinear neural networks, orthogonal constraints are shown to improve the trainability with proper scaling [20, 23].

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 that captures the relationship between a function and its derivative.

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

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = 1. \quad (7)$$

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, Bogachev [4] Theorem 1.6.4). *If function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable with bounded $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2]$ and $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2]$, then*

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

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 [16, 23, 27], 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 $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] = 0$ 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 $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2]$ and $\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2]$, there exist two constants a and b such that $a\phi(x) + b$ is Gaussian-Poincaré normalized.*

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

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[a^2\phi'(x)^2] = 1, \quad (9)$$

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

We computed a and b (not unique) for several common activation functions with their default hyperparameters¹ 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 $x \in (-\infty, 0) \cup (0, \infty)$.

	Tanh	ReLU	LeakyReLU [18]	ELU [6]	SELU [16]	GELU [11]
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

Table 1: Constants for Gaussian-Poincaré normalization of activation functions.

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.

¹We use $\alpha = 0.01$ for LeakyReLU, $\alpha = 1$ for ELU and $\phi(x) = x/(1 + \exp(-1.702x))$ for GELU.

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 $\mathbf{x} \in \mathbb{R}^n$ is thin-shell concentrated if for any $\epsilon > 0$*

$$\mathbb{P}\left\{\left|\frac{1}{n}\|\mathbf{x}\|_2^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (11)$$

as $n \rightarrow \infty$.

The definition is modified from the one in Bobkov [3]. Examples of thin-shell concentrated distributions include standard multivariate Gaussian and any distribution on the n -dimensional unit sphere scaled by \sqrt{n} .

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

Assumptions.

1. *Random vector $\mathbf{x} \in \mathbb{R}^n$ is thin-shell concentrated.*
2. *Random orthogonal matrix $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n)^T$ is uniformly distributed.*
3. *Function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is Gaussian-Poincaré normalized.*
4. *Function ϕ and its derivative are Lipschitz continuous.*

The above assumptions are not restrictive. For Assumption 1, one can always normalize the input vectors of a neural networks. For Assumption 2, orthogonal constraint or its relaxation has already been employed in neural network training [5]. 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 [19] 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). *Random vector*

$$(\phi(\mathbf{w}_1^T \mathbf{x}), \phi(\mathbf{w}_2^T \mathbf{x}), \dots, \phi(\mathbf{w}_n^T \mathbf{x})) \quad (12)$$

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 n is sufficiently large.

Theorem 3 (Backward Norm-Preservation). *Let \mathbf{D} be the diagonal matrix whose diagonal elements are $D_{ii} = \phi'(\mathbf{w}_i^T \mathbf{x})$ and $\mathbf{y} \in \mathbb{R}^n$ be a fixed vector with bounded $\|\mathbf{y}\|_\infty$. Then for any $\epsilon > 0$*

$$\mathbb{P}\left\{\frac{1}{n}\|\mathbf{D}\mathbf{y}\|_2^2 - \|\mathbf{y}\|_2^2 \geq \epsilon\right\} \rightarrow 0 \quad (13)$$

as $n \rightarrow \infty$.

This result shows that the multiplication by \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 (2), the norm is preserved in each layer along the backward path when n is sufficient 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 gradient exploding/vanishing 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 theory, concentration of measure. We refer the reader to Vershynin [31] 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. Furthermore, the random variables transformed by Lipschitz functions are also concentrated around certain values. Using this phenomenon, it can be shown that rows $\{\mathbf{w}_i\}$ of a random orthogonal matrix \mathbf{W} in high dimension are approximately independent random unit vectors and the inner product $\mathbf{w}_i^T \mathbf{x}$ for thin-shell concentrated vector \mathbf{x} can be shown to be approximately Gaussian. Then from the assumptions that ϕ is GPN and ϕ and ϕ' are Lipschitz continuous, the proofs follow. Each of these steps is rigorously proved in the Appendix.

3 Experiments

We verify our theory on both synthetic and real-world data. More experimental results can be found in the Appendix. 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 dimension. The gradient error $\partial E / \partial \mathbf{x}^{L+1}$ is also random standard Gaussian vector of 500 dimension. All the neural networks have depth 200. All the weight matrices are random orthogonal matrices uniformly generated. No training is performed.

In Figure 1, we show the norm of inputs and gradient of the neural networks of width 500. From the results, we can see that with GPN, the gradient exploding/vanishing problem is eliminated to large extent. The neural network with Tanh activation function does not show gradient exploding/vanishing 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 gradient exploding/vanishing 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 Figure 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 (Figure 2 (a)), which is not the case for BSNNs (Figure 2 (b)). This shows that BSNNs do not suffer from the gradient vanishing/explosion problem while exhibiting nonlinear functionality.

In Figure 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 and CIFAR-10. 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 $\mathbf{W} = (\mathbf{v}_1 / \|\mathbf{v}_1\|_2, \mathbf{v}_2 / \|\mathbf{v}_2\|_2, \dots, \mathbf{v}_n / \|\mathbf{v}_n\|_2)^T$ and optimize $\mathbf{V} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)^T$ 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 Figure 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 gradient (see Appendix).

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.* [24] and Yang *et al.* [32] that batch normalization does not solve the gradient exploding/vanishing problem. On the other hand, without batch normalization the neural network with any GPN activation function has stable gradient magnitude throughout training (see Appendix). This indicates that BSNNs can dispense batch normalization and therefore avoid its shortcomings.

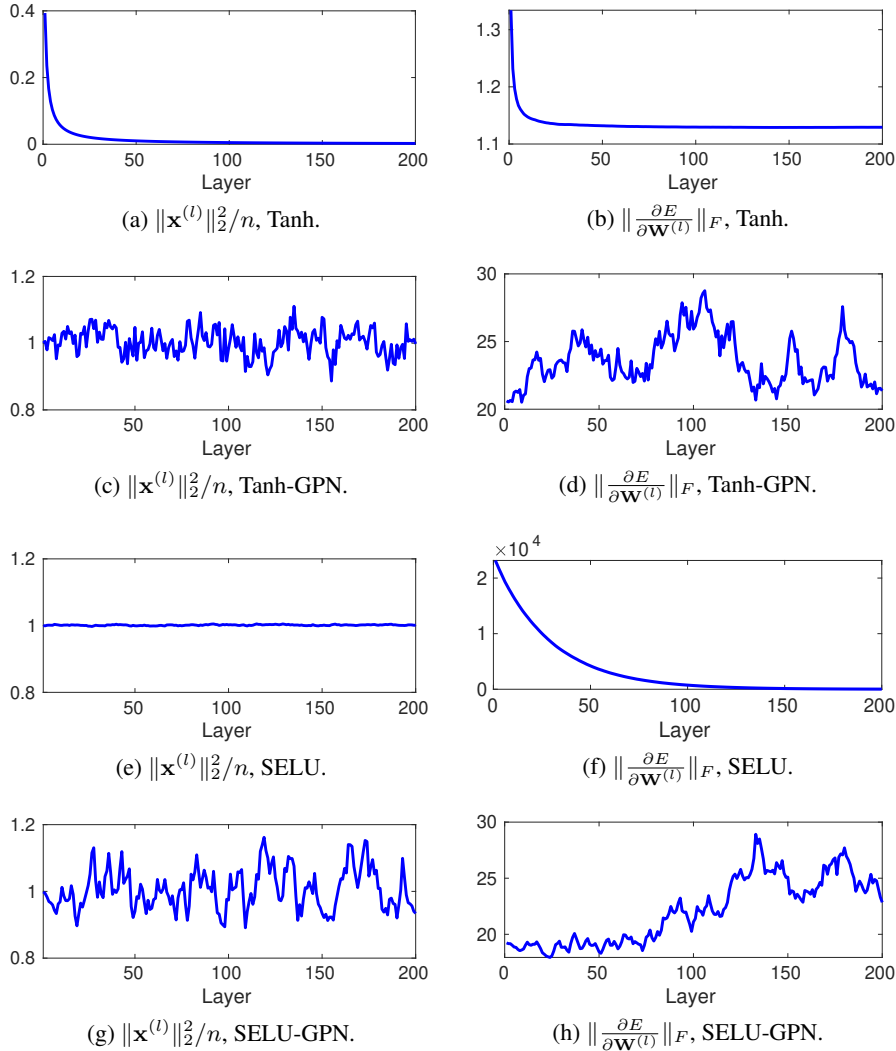


Figure 1: Results on synthetic data with different activation functions. “-GPN” denotes the function is Gaussian-Poincaré normalized. $\|\mathbf{x}^{(l)}\|_2$ denotes the l_2 norm of the outputs of the l -th layer. n denotes the width. $\|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F$ is the Frobenius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




Figure 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.




Figure 3: Gradient norm ratio for different layer width on synthetic data. The ratio is $\max_l \|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F / \min_l \|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F$. The width ranges from 100 to 1500. The error bars show standard deviation.

	MNIST		CIFAR-10	
	Train	Test	Train	Test
Tanh	99.05 (87.39)	96.57 (89.32)	80.84 (27.90)	42.71 (29.32)
Tanh-GPN	99.81 (84.93)	95.54 (87.11)	96.39 (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	33.28 (11.42)	28.13 (11.34)	46.60 (10.09)	34.96 (9.96)
LeakyReLU	11.24 (11.24)	11.35 (11.63)	10.00 (10.21)	10.00 (10.06)
LeakyReLU-GPN	43.17 (11.19)	49.28 (11.66)	51.85 (9.89)	39.38 (10.00)
ELU	99.06 (98.24)	95.41 (97.48)	80.73 (42.39)	45.76 (44.16)
ELU-GPN	100.00 (97.86)	96.56 (96.69)	99.37 (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	99.92 (97.91)	96.97 (97.39)	98.24 (47.74)	45.90 (45.52)
GELU	11.24 (12.70)	11.35 (10.28)	10.00 (10.43)	10.00 (10.00)
GELU-GPN	97.67 (11.22)	95.82 (9.74)	90.51 (10.00)	36.94 (10.00)

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.




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

4 Related Work

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 exploding/vanishing 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 [16]. However, as pointed out in Philipp *et al.* [24], only constraining forward signal propagation does not solve the gradient exploding/vanishing problem since the norm of the backward signal can grow or shrink. In Philipp *et al.* [24] and our experiments, SELU is indeed shown to cause gradient exploding. To solve the gradient exploding/vanishing problem, the signal propagation in both directions need to be constrained, as in our theory.

Our theory is largely developed from the deep signal propagation theory [25, 27]. 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.* [25, 27] 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 [24]). This is observed in Figure 1 (a). Moreover, in Poole *et al.* [25, 27] the signal propagation analysis is done based on random weights under i.i.d. Gaussian distribution whereas we proved how one can solve gradient vanishing/exploding problem assuming the weight matrices are orthogonal and uniformly distributed under Haar measure.

Dynamical isometry theory [23] 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 (2) 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 gradient exploding/vanishing 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 gradient exploding/vanishing 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 §3.1.

5 Conclusion

In this paper, we have introduced bidirectional self-normalizing neural network (BSNN) which constrains both forward and backward signal propagation using a new class of Gaussian-Poincaré normalized activation functions and orthogonal weight matrices. BSNNs are not restrictive in the sense that many commonly used activation functions can be Gaussian-Poincaré normalized. We have rigorously proved that gradient vanishing/exploding problem disappears in BSNNs with high probability under mild conditions. Experiments on synthetic and real-world data confirm the validity of our theory and demonstrate that BSNNs have excellent trainability without batch normalization. Currently, the theoretical analysis is limited to same width, fully-connected neural networks. Future work includes extending our theory to more sophisticated networks such as convolutional architectures as well as investigating the generalization capabilities of BSNNs.

References

- [1] T. Alberts and D. Khoshnevisan. *Calculus on Gauss Space: An Introduction to Gaussian Analysis*. 2018.
- [2] K. Ball. An elementary introduction to modern convex geometry. *Flavors of Geometry*, 1997.
- [3] S. G. Bobkov. On concentration of distributions of random weighted sums. *Annals of Probability*, 2003.
- [4] V. I. Bogachev. *Gaussian Measures*. American Mathematical Society, 1998.
- [5] A. Brock, T. Lim, J. M. Ritchie, and N. Weston. Neural photo editing with introspective adversarial networks. *International Conference on Learning Representations*, 2017.
- [6] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *International Conference on Learning Representations*, 2016.
- [7] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 2011.
- [8] N. Eldredge. When do ϕ^2 and ϕ'^2 have the same expectation under a gaussian random variable? MathOverflow. <https://mathoverflow.net/q/351553>.
- [9] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *International Conference on Artificial Intelligence and Statistics*, 2010.
- [10] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013.
- [11] D. Hendrycks and K. Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [12] S. Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. *Diploma, Technische Universität München*, 1991.
- [13] S. Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. *Advances in Neural Information Processing Systems*, 2017.
- [14] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning*, 2015.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- [16] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter. Self-normalizing neural networks. *Advances in Neural Information Processing Systems*, 2017.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 2012.
- [18] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. *International Conference on Machine Learning*, 2013.
- [19] E. S. Meckes. *The Random Matrix Theory of the Classical Compact Groups*. Cambridge University Press, 2019.
- [20] D. Mishkin and J. Matas. All you need is a good init. *International Conference on Learning Representations*, 2016.
- [21] K. R. Parthasarathy. *Probability Measures on Metric Spaces*. American Mathematical Society, 1967.
- [22] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. *International Conference on Machine Learning*, 2013.

- [23] J. Pennington, S. Schoenholz, and S. Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. *Advances in Neural Information Processing Systems*, 2017.
- [24] G. Philipp, D. Song, and J. G. Carbonell. The exploding gradient problem demystified-definition, prevalence, impact, origin, tradeoffs, and solutions. *arXiv preprint arXiv:1712.05577*, 2018.
- [25] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli. Exponential expressivity in deep neural networks through transient chaos. *Advances in Neural Information Processing Systems*, 2016.
- [26] A. M. Saxe, J. L. McClelland, and S. Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *International Conference on Learning Representations*, 2014.
- [27] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein. Deep information propagation. *International Conference on Learning Representations*, 2017.
- [28] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 2017.
- [29] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 2012.
- [30] A. W. Van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- [31] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018.
- [32] G. Yang, J. Pennington, V. Rao, J. Sohl-Dickstein, and S. S. Schoenholz. A mean field theory of batch normalization. *International Conference on Learning Representations*, 2019.

Appendix A Proofs

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

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(1)}} \right\|_F = \left\| \frac{\partial E}{\partial \mathbf{W}^{(2)}} \right\|_F = \dots = \left\| \frac{\partial E}{\partial \mathbf{W}^{(L)}} \right\|_F. \quad (14)$$

Proof. For each l , we have

$$\left\| \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right\|_F = \sqrt{\text{trace} \left(\frac{\partial E}{\partial \mathbf{W}^{(l)}} \left(\frac{\partial E}{\partial \mathbf{W}^{(l)}} \right)^T \right)} \quad (15)$$

$$= \sqrt{\text{trace}(\mathbf{d}^{(l)} (\mathbf{x}^{(l)})^T \mathbf{x}^{(l)} (\mathbf{d}^{(l)})^T)} \quad (16)$$

$$= \sqrt{\text{trace}((\mathbf{x}^{(l)})^T \mathbf{x}^{(l)} (\mathbf{d}^{(l)})^T \mathbf{d}^{(l)})} \quad (17)$$

$$= \sqrt{(\mathbf{x}^{(l)})^T \mathbf{x}^{(l)}} \sqrt{(\mathbf{d}^{(l)})^T \mathbf{d}^{(l)}} \quad (18)$$

$$= \sqrt{n} \left\| \frac{\partial E}{\partial \mathbf{x}^{(L+1)}} \right\|_2. \quad (19)$$

□

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

$$H_k(x) = \frac{(-1)^k}{\sqrt{k!}} \exp\left(\frac{x^2}{2}\right) \frac{d^k}{dx^k} \exp\left(-\frac{x^2}{2}\right) \quad (20)$$

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

$$\phi(x) = \sum_{k=0}^{\infty} a_k H_k(x), \quad (21)$$

$$\phi'(x) = \sum_{k=1}^{\infty} \sqrt{k} a_k H_{k-1}(x). \quad (22)$$

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

$$a_0 = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[H_0(x)\phi(x)] \quad (23)$$

$$= \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)] \quad (24)$$

$$= 0. \quad (25)$$

Since

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = 1 \quad (26)$$

and Hermite polynomials are orthonormal, we have

$$\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi(x)^2] = \sum_{k=1}^{\infty} a_k^2 = \mathbb{E}_{x \sim \mathcal{N}(0,1)}[\phi'(x)^2] = \sum_{k=1}^{\infty} k a_k^2 = 1. \quad (27)$$

Therefore, we have

$$\sum_{k=1}^{\infty} k a_k^2 - \sum_{k=1}^{\infty} a_k^2 = 0 \quad (28)$$

that is

$$\sum_{k=2}^{\infty} (k-1) a_k^2 = 0. \quad (29)$$

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, $\phi(x) = \pm H_1(x) = \pm x$. □

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 and b such that $a\phi(x) + b$ is Gaussian-Poincaré normalized.

Proof. Let $\varphi(x) = \phi(x) + c$. Then let

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

$$= \text{Var}_{x \sim \mathcal{N}(0,1)}[\varphi(x)] + (\mathbb{E}_{x \sim \mathcal{N}(0,1)}[\varphi(x)])^2 - \mathbb{E}_{x \sim \mathcal{N}(0,1)}[(\phi'(x))^2] \quad (31)$$

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

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

The proof is largely due to [8] with minor modification in here.

Assumptions.

1. Random vector $\mathbf{x} \in \mathbb{R}^n$ is thin-shell concentrated.
2. Random orthogonal matrix $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n)^T$ is uniformly distributed.
3. Function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is Gaussian-Poincaré normalized.
4. Function ϕ and its derivative are Lipschitz continuous.

Theorem 2 (Forward Norm-Preservation). Random vector

$$(\phi(\mathbf{w}_1^T \mathbf{x}), \phi(\mathbf{w}_2^T \mathbf{x}), \dots, \phi(\mathbf{w}_n^T \mathbf{x})) \quad (33)$$

is thin-shell concentrated.

Theorem 3 (Backward Norm-Preservation). Let \mathbf{D} be the diagonal matrix whose diagonal elements are $D_{ii} = \phi'(\mathbf{w}_i^T \mathbf{x})$ and $\mathbf{y} \in \mathbb{R}^n$ be a fixed vector with bounded $\|\mathbf{y}\|_\infty$. Then for any $\epsilon > 0$

$$\mathbb{P}\left\{\left|\frac{1}{n}\|\mathbf{D}\mathbf{y}\|_2^2 - \|\mathbf{y}\|_2^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (34)$$

as $n \rightarrow \infty$.

Remark. To prove Theorem 2 and 3, we need the following lemmas. Lemma 1, 2, 3 and 4 are existing results. Lemma 5 and 6 follow immediately from existing results. Lemma 7, 8, 9 and 10 are our contribution. The proofs require the basic knowledge of Lebesgue integral and measure theory.

Notations. $\mathbb{S}^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$, $\mathcal{S}(n) = \overbrace{\mathbb{S}^{n-1} \times \mathbb{S}^{n-1} \times \dots \times \mathbb{S}^{n-1}}^{n \text{ factors}}$ and $\mathbb{O}(n)$ is the orthogonal matrix group of size n . $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. \mathbf{I}_n denotes the identity matrix of size $n \times n$.

Lemma 1. Let (M, d) be a metric space where M is a set and d is the metric function. Let \mathcal{B} be the Borel algebra of M . Then for any $A \in \mathcal{B}$, any probability measure μ of M and any $\epsilon > 0$, there exists an open set G such that $A \subseteq G$ and $\mu(G \setminus A) < \epsilon$.

See [21] (Theorem 1.2 in page 27) for a proof.

Lemma 2. Let (M, d) be a metric space where M is a set and d is the metric function. For every open set $G \subseteq M$, there exists a sequence of bounded and Lipschitz function f_k such that $0 \leq f_k \leq 1$ and $f_k \rightarrow \mathbf{1}_{\{x \in G\}}$ as $k \rightarrow \infty$.

Proof. Let $f_k(x) = \min(kd(x, G^c), 1)$ where G^c denotes the complement of G and $d(x, G^c) = \min_{y \in G^c} d(x, y)$. \square

This constructive proof is from [30] (Lemma 2.2).

Lemma 3. *If random vector $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are independently and uniformly distributed in \mathbb{S}^{n-1} , then for any $0 < \epsilon < 1$*

$$\mathbb{P}\{|\boldsymbol{\theta}_1^T \boldsymbol{\theta}_2| \geq \epsilon\} \leq 2 \exp\left(-\frac{n\epsilon^2}{2}\right). \quad (35)$$

The proof follows from [2] (Lemma 2.2).

Lemma 4. *If random vector $\mathbf{v} \in \mathbb{R}^n$ is standard multivariate Gaussian, then for any $\epsilon > 0$*

$$\mathbb{P}\{\mathbf{v} \in \mathbb{R}^n : (1 - \epsilon)\sqrt{n} \leq \|\mathbf{v}\|_2 \leq (1 + \epsilon)\sqrt{n}\} \geq 1 - 2 \exp(-n\epsilon^2). \quad (36)$$

See [1] (Theorem 1.2) for a proof.

Lemma 5. *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 [31]), we have

$$\|f(x) - \mathbb{E}[f(x)]\|_{\psi_2} \leq CK \quad (37)$$

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 [31]). Therefore $f(x)$ is sub-gaussian (Lemma 2.6.8 in [31]). \square

Lemma 6. *Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ be a random vector that each coordinate x_i is independent and sub-gaussian and $\mathbb{E}[x_i^2] = 1$. Let $\mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$ be a fixed vector with bounded $\|\mathbf{y}\|_\infty$. Then*

$$\mathbb{P}\left\{\frac{1}{n} \left| \sum_i x_i^2 y_i^2 - \sum_i y_i^2 \right| \geq \epsilon\right\} \rightarrow 0 \quad (38)$$

as $n \rightarrow \infty$.

Proof. Since $y_i x_i$ is sub-gaussian, then $y_i^2 x_i^2$ is sub-exponential (Lemma 2.7.6 in [31]). 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 [31]). Applying Bernstein's inequality (Corollary 2.8.3 in [31]), we proved the lemma. \square

Lemma 7. *If random vector $\mathbf{x} \in \mathbb{R}^n$ is thin-shell concentrated, random vector $\boldsymbol{\theta}$ is uniformly distributed in \mathbb{S}^{n-1} and $z \sim \mathcal{N}(0, 1)$, then for any $A \in \mathcal{B}$ where \mathcal{B} is the Borel algebra of \mathbb{R}*

$$|\mathbb{P}\{\mathbf{x}^T \boldsymbol{\theta} \in A\} - \mathbb{P}\{z \in A\}| \rightarrow 0 \quad (39)$$

as $n \rightarrow \infty$.

Proof. First, due to Lemma 1, for any $\epsilon > 0$ there exists an open set G such that $A \subseteq G$ and

$$|\mathbb{P}\{\mathbf{x}^T \boldsymbol{\theta} \in G\} - \mathbb{P}\{\mathbf{x}^T \boldsymbol{\theta} \in A\}| < \epsilon, \quad (40)$$

$$|\mathbb{P}\{z \in G\} - \mathbb{P}\{z \in A\}| < \epsilon. \quad (41)$$

Next, due to Lemma 2 and the Bounded Convergence Theorem, there exists a sequence of bounded and Lipschitz continuous function $f_k : \mathbb{R} \rightarrow \mathbb{R}$ such that $0 \leq f_k \leq 1$

$$|\mathbb{E}[f_k(\mathbf{x}^T \boldsymbol{\theta})] - \mathbb{P}\{\mathbf{x}^T \boldsymbol{\theta} \in G\}| = \left| \int_{\mathbb{R}} f_k(\mathbf{x}^T \boldsymbol{\theta}) d\mu(\mathbf{x}^T \boldsymbol{\theta}) - \int_{\mathbb{R}} \mathbf{1}_{\{\mathbf{x}^T \boldsymbol{\theta} \in G\}} d\mu(\mathbf{x}^T \boldsymbol{\theta}) \right| \quad (42)$$

$$\leq \int_{\mathbb{R}} |f_k(\mathbf{x}^T \boldsymbol{\theta}) - \mathbf{1}_{\{\mathbf{x}^T \boldsymbol{\theta} \in G\}}| d\mu(\mathbf{x}^T \boldsymbol{\theta}) \rightarrow 0 \quad (43)$$

where μ denotes the probability measure for $\mathbf{x}^T \boldsymbol{\theta}$, as $k \rightarrow \infty$. Similarly, $|\mathbb{E}[f_k(z)] - \mathbb{P}\{z \in G\}| \rightarrow 0$ as $k \rightarrow \infty$.

Thus, if for any bounded and Lipschitz continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$

$$|\mathbb{E}[f(\mathbf{x}^T \boldsymbol{\theta})] - \mathbb{E}[f(z)]| \rightarrow 0 \quad (44)$$

as $n \rightarrow \infty$, then the lemma is proved.

Let $\mathbf{y} = \sqrt{n}/\|\mathbf{x}\|_2 \mathbf{x}$ and K be the Lipschitz constant of f . Since $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is thin-shell concentrated, we have

$$|f(\mathbf{x}^T \boldsymbol{\theta}) - f(\mathbf{y}^T \boldsymbol{\theta})| \leq K |\mathbf{x}^T \boldsymbol{\theta} - \sqrt{n}/\|\mathbf{x}\|_2 \mathbf{x}^T \boldsymbol{\theta}| \rightarrow 0 \quad (45)$$

and therefore $|\mathbb{E}[f(\mathbf{x}^T \boldsymbol{\theta})] - \mathbb{E}[f(\mathbf{y}^T \boldsymbol{\theta})]| \rightarrow 0$ as $n \rightarrow \infty$.

It is a known result that $\boldsymbol{\theta}$ is equivalent to $\mathbf{v}/\|\mathbf{v}\|_2$ for $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Define the following set for $0 < \epsilon < 1$

$$\mathcal{U}_\epsilon = \{\mathbf{v} \in \mathbb{R}^n : (1 - \epsilon)\sqrt{n} \leq \|\mathbf{v}\|_2 \leq (1 + \epsilon)\sqrt{n}\}. \quad (46)$$

We can then write $\mathbb{E}_\theta[f(\mathbf{y}^T \boldsymbol{\theta})]$ as

$$\int_{\mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) + \int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) \quad (47)$$

where $\gamma_n(\mathbf{v})$ denotes the standard Gaussian measure of dimension n .

We know from Lemma 4 that

$$\int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} d\gamma_n(\mathbf{v}) \leq 2 \exp(-n\epsilon^2). \quad (48)$$

Let $\epsilon = n^{-1/4}$. Since f is bounded, let $M = \sup |f|$. We have

$$\left| \int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) - \int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\sqrt{n}) d\gamma_n(\mathbf{v}) \right| \quad (49)$$

$$\leq \int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} |f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) - f(\mathbf{y}^T \mathbf{v}/\sqrt{n})| d\gamma_n(\mathbf{v}) \quad (50)$$

$$\leq 2M \int_{\mathbb{R}^n \setminus \mathcal{U}_\epsilon} d\gamma_n(\mathbf{v}) \rightarrow 0 \quad (51)$$

as $n \rightarrow \infty$.

For $\mathbf{v} \in \mathcal{U}_\epsilon$, since

$$\frac{1}{\sqrt{n}} \frac{1}{1 + \epsilon} \leq \frac{1}{\|\mathbf{v}\|_2} \leq \frac{1}{\sqrt{n}} \frac{1}{1 - \epsilon} \quad (52)$$

we have

$$-\frac{1}{\sqrt{n}} \frac{\epsilon}{1 + \epsilon} \leq \frac{1}{\|\mathbf{v}\|_2} - \frac{1}{\sqrt{n}} \leq \frac{1}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon}. \quad (53)$$

Then

$$\left| \frac{1}{\|\mathbf{v}\|_2} - \frac{1}{\sqrt{n}} \right| \leq \frac{1}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon} \quad (54)$$

and

$$\left| \frac{\mathbf{y}^T \mathbf{v}}{\|\mathbf{v}\|_2} - \frac{\mathbf{y}^T \mathbf{v}}{\sqrt{n}} \right| = \left| \frac{1}{\|\mathbf{v}\|_2} - \frac{1}{\sqrt{n}} \right| |\mathbf{y}^T \mathbf{v}| \leq \frac{1}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon} |\mathbf{y}^T \mathbf{v}|. \quad (55)$$

Hence,

$$|f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) - f(\mathbf{y}^T \mathbf{v}/\sqrt{n})| \leq \frac{K}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon} |\mathbf{y}^T \mathbf{v}| \quad (56)$$

and

$$\left| \int_{\mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) - \int_{\mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v}/\sqrt{n}) d\gamma_n(\mathbf{v}) \right| \quad (57)$$

$$\leq \int_{\mathcal{U}_\epsilon} |f(\mathbf{y}^T \mathbf{v}/\|\mathbf{v}\|_2) - f(\mathbf{y}^T \mathbf{v}/\sqrt{n})| d\gamma_n(\mathbf{v}) \quad (58)$$

$$\leq \frac{K}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon} \int_{\mathcal{U}_\epsilon} |\mathbf{y}^T \mathbf{v}| d\gamma_n(\mathbf{v}) \quad (59)$$

$$\leq \frac{K}{\sqrt{n}} \frac{\epsilon}{1 - \epsilon} \int_{\mathbb{R}^n} |\mathbf{y}^T \mathbf{v}| d\gamma_n(\mathbf{v}). \quad (60)$$

Since $\|\mathbf{y}\|_2 = \sqrt{n}$ and $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, random variable $\mathbf{y}^T \mathbf{v}$ is equivalent to $\sqrt{n}z$ and $|\sqrt{n}z|$ has a half-normal distribution. Using the known result of the mean of half-normal distribution, we have

$$\int_{\mathbb{R}^n} |\mathbf{y}^T \mathbf{v}| d\gamma_n(\mathbf{v}) = \mathbb{E}[|\sqrt{n}z|] = \frac{\sqrt{n}\sqrt{2}}{\sqrt{\pi}}. \quad (61)$$

Hence,

$$\left| \int_{\mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v} / \|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) - \int_{\mathcal{U}_\epsilon} f(\mathbf{y}^T \mathbf{v} / \sqrt{n}) d\gamma_n(\mathbf{v}) \right| \quad (62)$$

$$\leq \frac{K}{\sqrt{n}} \frac{\epsilon}{1-\epsilon} \frac{\sqrt{n}\sqrt{2}}{\sqrt{\pi}} = \frac{\epsilon}{1-\epsilon} \frac{K\sqrt{2}}{\sqrt{\pi}}. \quad (63)$$

For $\epsilon = n^{-1/4}$, we have

$$\left| \int_{\mathbb{R}^n} f(\mathbf{y}^T \mathbf{v} / \|\mathbf{v}\|_2) d\gamma_n(\mathbf{v}) - \int_{\mathbb{R}^n} f(\mathbf{y}^T \mathbf{v} / \sqrt{n}) d\gamma_n(\mathbf{v}) \right| \rightarrow 0 \quad (64)$$

as $n \rightarrow \infty$. Also since random variable $\mathbf{y}^T \mathbf{v} / \sqrt{n}$ is equivalent to $z \sim \mathcal{N}(0, 1)$, we have

$$\int_{\mathbb{R}^n} f(\mathbf{y}^T \mathbf{v} / \sqrt{n}) d\gamma_n(\mathbf{v}) = \mathbb{E}_{z \sim \mathcal{N}(0,1)}[f(z)]. \quad (65)$$

Therefore,

$$|\mathbb{E}[f(\mathbf{x}^T \boldsymbol{\theta})] - \mathbb{E}[f(z)]| \rightarrow 0 \quad (66)$$

as $n \rightarrow \infty$. \square

Lemma 8. *If random matrix $\boldsymbol{\Theta}$ is uniformly distributed in $\mathcal{S}(n)$ and let $\mathbf{U}\mathbf{D}\mathbf{V}^T = \boldsymbol{\Theta}$ be its singular value decomposition, then $P(\boldsymbol{\Theta}) = \mathbf{U}\mathbf{V}^T$ is uniformly distributed in $\mathbb{O}(n)$.*

Proof. By the property of singular value decomposition, we have $P(\boldsymbol{\Theta}) = \operatorname{argmin}_{\mathbf{Q} \in \mathbb{O}(n)} \|\mathbf{Q} - \boldsymbol{\Theta}\|_F$. Therefore, for any $\mathbf{R} \in \mathbb{O}(n)$,

$$P(\boldsymbol{\Theta}\mathbf{R}) = \operatorname{argmin}_{\mathbf{Q} \in \mathbb{O}(n)} \|\mathbf{Q} - \boldsymbol{\Theta}\mathbf{R}\|_F \quad (67)$$

$$= P(\boldsymbol{\Theta})\mathbf{R}. \quad (68)$$

Let μ be the uniform probability measure for $\boldsymbol{\Theta}$ and ν be the probability measure for $P(\boldsymbol{\Theta})$. For any $A \subseteq \mathbb{O}(n)$ and any $\mathbf{R} \in \mathbb{O}(n)$, let

$$\mathbf{A}\mathbf{R} = \{\mathbf{Q}\mathbf{R} \in \mathbb{O}(n) : \mathbf{Q} \in A\}, \quad (69)$$

$$P^{-1}(A) = \{\boldsymbol{\Theta} \in \mathcal{S}(n) : P(\boldsymbol{\Theta}) \in A\}, \quad (70)$$

$$P^{-1}(A)\mathbf{R} = \{\boldsymbol{\Theta}\mathbf{R} \in \mathcal{S}(n) : \boldsymbol{\Theta} \in P^{-1}(A)\}, \quad (71)$$

then

$$\nu(A) = \mu(P^{-1}(A)) = \mu(P^{-1}(A)\mathbf{R}) = \mu(P^{-1}(\mathbf{A}\mathbf{R})) = \nu(\mathbf{A}\mathbf{R}). \quad (72)$$

Therefore $P(\boldsymbol{\Theta})$ is uniformly distributed in $\mathbb{O}(n)$. \square

Lemma 9. *Let $\mathbf{U}\mathbf{D}\mathbf{V}^T$ be the singular value decomposition of matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$. Then*

$$\|\mathbf{W} - \mathbf{U}\mathbf{V}^T\|_F \leq \|\mathbf{W}\mathbf{W}^T - \mathbf{I}_n\|_F. \quad (73)$$

Proof.

$$\|\mathbf{W}\mathbf{W}^T - \mathbf{I}_n\|_F = \|(\mathbf{U}\mathbf{D}\mathbf{V}^T)(\mathbf{U}\mathbf{D}\mathbf{V}^T)^T - \mathbf{I}_n\|_F \quad (74)$$

$$= \|\mathbf{U}\mathbf{D}^2\mathbf{U}^T - \mathbf{I}_n\|_F \quad (75)$$

$$= \|\mathbf{D}^2 - \mathbf{I}_n\|_F. \quad (76)$$

Since $(x^2 - 1)^2 \geq (x - 1)^2$ for $x \in \mathbb{R}$, we have

$$\|\mathbf{D}^2 - \mathbf{I}_n\|_F \geq \|\mathbf{D} - \mathbf{I}_n\|_F \quad (77)$$

$$= \|\mathbf{U}\mathbf{D}\mathbf{V}^T - \mathbf{U}\mathbf{I}_n\mathbf{V}^T\|_F \quad (78)$$

$$= \|\mathbf{W} - \mathbf{U}\mathbf{V}^T\|_F. \quad (79)$$

\square

Lemma 10. *If random matrix Θ is uniformly distributed in $\mathcal{S}(n)$ and random matrix \mathbf{W} is uniformly distributed in $\mathbb{O}(n)$, then for $A \in \mathcal{B}$ where \mathcal{B} is the Borel algebra of $\mathcal{S}(n)$*

$$|\mathbb{P}\{\Theta \in A\} - \mathbb{P}\{\mathbf{W} \in A\}| \rightarrow 0 \quad (80)$$

as $n \rightarrow \infty$.

Proof. First, due to Lemma 1, for any $\epsilon > 0$ there exists an open set G such that $A \subseteq G$ and

$$|\mathbb{P}\{\Theta \in G\} - \mathbb{P}\{\Theta \in A\}| < \epsilon, \quad (81)$$

$$|\mathbb{P}\{\mathbf{W} \in G\} - \mathbb{P}\{\mathbf{W} \in A\}| < \epsilon. \quad (82)$$

Next, due to Lemma 2 and the Bounded Convergence Theorem, there exists a sequence of bounded and Lipschitz continuous function $f_k : \mathcal{S}(n) \rightarrow \mathbb{R}$ such that $0 \leq f_k \leq 1$ and

$$|\mathbb{E}[f_k(\Theta)] - \mathbb{P}\{\Theta \in G\}| = \left| \int_{\mathcal{S}(n)} f_k(\Theta) d\mu(\Theta) - \int_{\mathcal{S}(n)} \mathbf{1}_{\{\Theta \in G\}} d\mu(\Theta) \right| \quad (83)$$

$$\leq \int_{\mathcal{S}(n)} |f_k(\Theta) - \mathbf{1}_{\{\Theta \in G\}}| d\mu(\Theta) \rightarrow 0 \quad (84)$$

where μ denotes the uniform probability measure for Θ , as $k \rightarrow \infty$. Similarly, $|\mathbb{E}[f_k(\mathbf{W})] - \mathbb{P}\{\mathbf{W} \in G\}| \rightarrow 0$ as $k \rightarrow \infty$.

Thus, if for any bounded and Lipschitz function $f : \mathcal{S}(n) \rightarrow \mathbb{R}$

$$|\mathbb{E}[f(\Theta)] - \mathbb{E}[f(\mathbf{W})]| \rightarrow 0 \quad (85)$$

as $n \rightarrow \infty$, then the lemma is proved.

Let the metric function be $d(\mathbf{X}, \mathbf{Y}) = \frac{1}{n} \|\mathbf{X} - \mathbf{Y}\|_F$ for $\mathbf{X} \in \mathcal{S}(n)$ and $\mathbf{Y} \in \mathcal{S}(n)$. Let $\nu(\mathbf{W})$ the uniform probability measure on $\mathbb{O}(n)$. Let $\mathbf{U}\mathbf{D}\mathbf{V}^T = \Theta$ be the singular value decomposition of Θ . Then $P(\Theta) = \mathbf{U}\mathbf{V}^T$ is the projection of Θ onto $\mathbb{O}(n)$ under Frobenius norm. From Lemma 8, we have that $P(\Theta)$ is uniformly distributed in $\mathbb{O}(n)$ since Θ is uniformly distributed in $\mathcal{S}(n)$. Thus we have

$$\int_{\mathbb{O}(n)} f(\mathbf{W}) d\nu(\mathbf{W}) = \int_{\mathcal{S}(n)} f(P(\Theta)) d\mu(\Theta). \quad (86)$$

Define for $0 < \epsilon < 1$

$$\mathcal{U}_\epsilon = \{\Theta \in \mathcal{S}(n) : |\theta_i^T \theta_j| < \epsilon \text{ for } i \neq j\}. \quad (87)$$

We have, from Lemma 3,

$$\int_{\mathcal{U}_\epsilon} d\mu(\Theta) = \mathbb{P}\{\Theta \in \mathcal{U}_\epsilon\} = \mathbb{P}\{|\theta_i^T \theta_j| < \epsilon \text{ for } i \neq j\} \quad (88)$$

$$= 1 - \sum_{i < j} \mathbb{P}\{|\theta_i^T \theta_j| \geq \epsilon\} \geq 1 - n(n-1) \exp\left(-\frac{n\epsilon^2}{2}\right) \quad (89)$$

so that

$$\int_{\mathcal{S}(n) \setminus \mathcal{U}_\epsilon} d\mu(\Theta) = 1 - \int_{\mathcal{U}_\epsilon} d\mu(\Theta) \leq n(n-1) \exp\left(-\frac{n\epsilon^2}{2}\right). \quad (90)$$

Then

$$\left| \int_{\mathcal{S}(n)} f(\Theta) d\mu(\Theta) - \int_{\mathcal{S}(n)} f(P(\Theta)) d\mu(\Theta) \right| \quad (91)$$

$$\leq \int_{\mathcal{S}(n)} |f(\Theta) - f(P(\Theta))| d\mu(\Theta) \quad (92)$$

$$= \int_{\mathcal{S}(n) \setminus \mathcal{U}_\epsilon} |f(\Theta) - f(P(\Theta))| d\mu(\Theta) + \int_{\mathcal{U}_\epsilon} |f(\Theta) - f(P(\Theta))| d\mu(\Theta) \quad (93)$$

$$\leq 2M \int_{\mathcal{S}(n) \setminus \mathcal{U}_\epsilon} d\mu(\Theta) + \frac{K}{n} \int_{\mathcal{U}_\epsilon} \|\Theta - P(\Theta)\|_F d\mu(\Theta) \quad (94)$$

where $M = \sup |f|$ since f is bounded and K is the Lipschitz constant of f .

Let $\epsilon = n^{-1/4}$ and we have $\int_{\mathcal{S}(n) \setminus \mathcal{U}_\epsilon} d\mu(\Theta) \rightarrow 0$ as $n \rightarrow \infty$. For $\Theta \in \mathcal{U}_\epsilon$, $\|\Theta\Theta^T - \mathbf{I}_n\|_F < n\epsilon$. Applying Lemma 9, we have $\|\Theta - P(\Theta)\|_F < n\epsilon$. Then

$$\frac{1}{n} \int_{\mathcal{U}_\epsilon} \|\Theta - P(\Theta)\|_F d\mu(\Theta) < \epsilon \int_{\mathcal{U}_\epsilon} d\mu(\Theta) \leq \epsilon \int_{\mathcal{S}(n)} d\mu(\Theta) = \epsilon \rightarrow 0. \quad (95)$$

Therefore,

$$|\mathbb{E}[f(\Theta)] - \mathbb{E}[f(\mathbf{W})]| \rightarrow 0 \quad (96)$$

as $n \rightarrow \infty$. \square

Proof of Theorem 2. Let z_1, z_2, \dots, z_n be independent standard Gaussian random variables. Due to Lemma 5, random variable $\phi(z_i)$ is sub-gaussian. Then, applying Lemma 6 with each $y_i = 1$ and with the assumption $\mathbb{E}_{z_i \sim \mathcal{N}(0,1)}[\phi(z_i)^2] = 1$, we have for

$$\mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(z_i)^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (97)$$

as $n \rightarrow \infty$.

Next, let random matrix $\Theta = (\theta_1, \theta_2, \dots, \theta_n)^T$ be uniformly distributed in $\mathcal{S}(n)$ such that each θ_i is independent of others.

Due to Lemma 7, we have

$$\left|\mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(\theta_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\} - \mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(z_i)^2 - 1\right| \geq \epsilon\right\}\right| \rightarrow 0 \quad (98)$$

and therefore

$$\mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(\theta_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (99)$$

as $n \rightarrow \infty$.

Finally, let random matrix $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n)^T$ be uniformly distributed in $\mathbb{O}(n)$. Due to Lemma 10, we have

$$\left|\mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(\theta_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\} - \mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(\mathbf{w}_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\}\right| \rightarrow 0 \quad (100)$$

and therefore

$$\mathbb{P}\left\{\left|\frac{1}{n} \sum_i \phi(\mathbf{w}_i^T \mathbf{x})^2 - 1\right| \geq \epsilon\right\} \rightarrow 0 \quad (101)$$

as $n \rightarrow \infty$. \square

Proof of Theorem 3. Let z_1, z_2, \dots, z_n be independent standard Gaussian random variables. Due to Lemma 5, random variable $\phi'(z_i)$ is sub-gaussian. Then, applying Lemma 6 with the assumption $\mathbb{E}_{z_i \sim \mathcal{N}(0,1)}[\phi'(z_i)^2] = 1$, we have

$$\mathbb{P}\left\{\frac{1}{n} \left|\sum_i y_i^2 \phi'(z_i)^2 - y_i^2\right| \geq \epsilon\right\} \rightarrow 0 \quad (102)$$

as $n \rightarrow \infty$.

Next, let random matrix $\Theta = (\theta_1, \theta_2, \dots, \theta_n)^T$ be uniformly distributed in $\mathcal{S}(n)$ such that each θ_i is independent of others. Due to Lemma 7, we have

$$\left|\mathbb{P}\left\{\frac{1}{n} \left|\sum_i y_i^2 \phi'(\theta_i^T \mathbf{x})^2 - y_i^2\right| \geq \epsilon\right\} - \mathbb{P}\left\{\frac{1}{n} \left|\sum_i y_i^2 \phi'(z_i)^2 - y_i^2\right| \geq \epsilon\right\}\right| \rightarrow 0 \quad (103)$$

as $n \rightarrow \infty$.

Finally, let random matrix $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n)^T$ be uniformly distributed in $\mathbb{O}(n)$. Due to Lemma 10, we have

$$\left| \mathbb{P}\left\{\frac{1}{n} \left| \sum_i y_i^2 \phi'(\boldsymbol{\theta}_i^T \mathbf{x})^2 - y_i^2 \right| \geq \epsilon\right\} - \mathbb{P}\left\{\frac{1}{n} \left| \sum_i y_i^2 \phi'(\mathbf{w}_i^T \mathbf{x})^2 - y_i^2 \right| \geq \epsilon\right\} \right| \rightarrow 0 \quad (104)$$

and therefore

$$\mathbb{P}\left\{\frac{1}{n} \left| \sum_i \phi'(\mathbf{w}_i^T \mathbf{x})^2 y_i^2 - \sum_i y_i^2 \right| \geq \epsilon\right\} \rightarrow 0 \quad (105)$$

as $n \rightarrow \infty$. □

Appendix B Additional Experiments

Due to the space limitation, we only showed the experiments with Tanh and SELU activation functions in the main text. In this section, we show the experiments with ReLU, LeakyReLU, ELU and SELU. Additionally, we also measure the gradient exploding/vanishing during training on the real-world data.

B.1 Synthetic Data

In Figure 5 and 6, we show the experiments in addition to Figure 1. In Figure 7, we show the experiments in addition to Figure 2. In Figure 8, we show the experiments in addition to Figure 3. ELU shows similar behaviors as Tanh since $\text{ELU}(x) \approx x$ for $x \approx 0$.

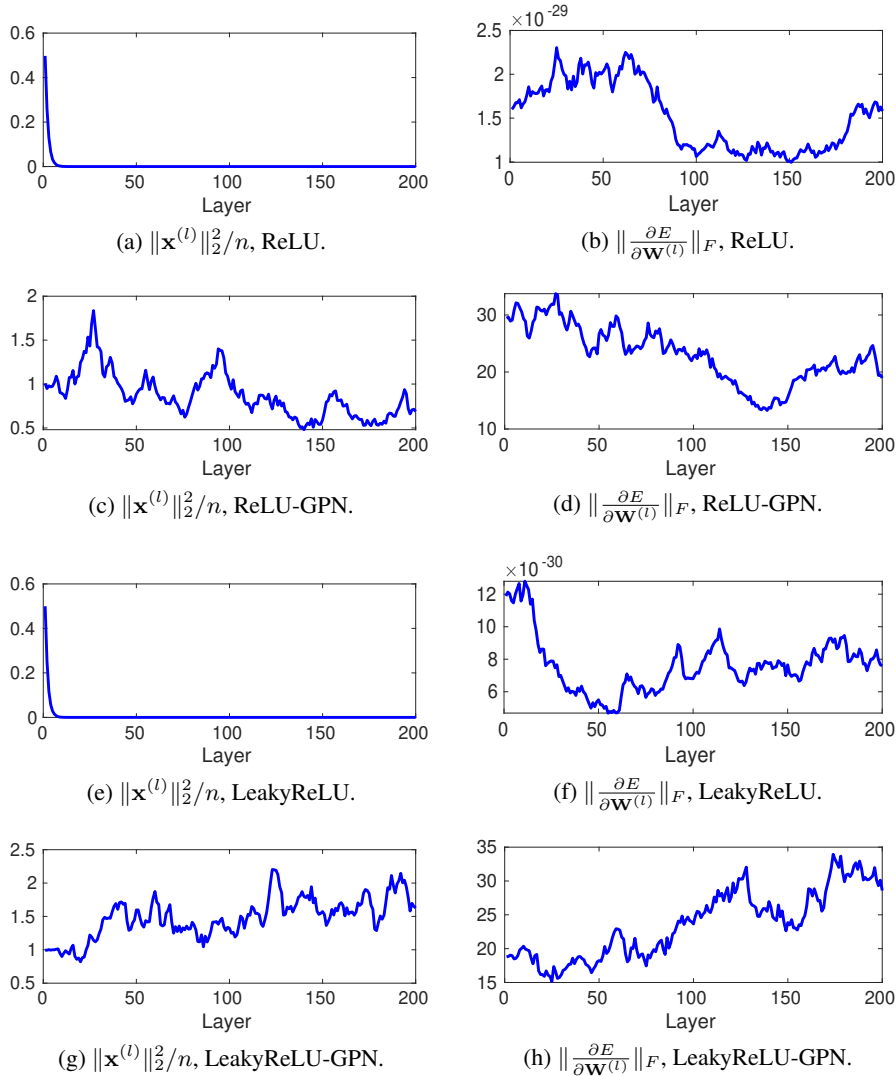


Figure 5: Results on synthetic data with different activation functions. “-GPN” denotes the function is Gaussian-Poincaré normalized. $\|\mathbf{x}^{(l)}\|_2$ denotes the l_2 norm of the outputs of the l -th layer. n denotes the width. $\|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F$ is the Frobenius norm of the gradient of the weight matrix in the l -th layer.

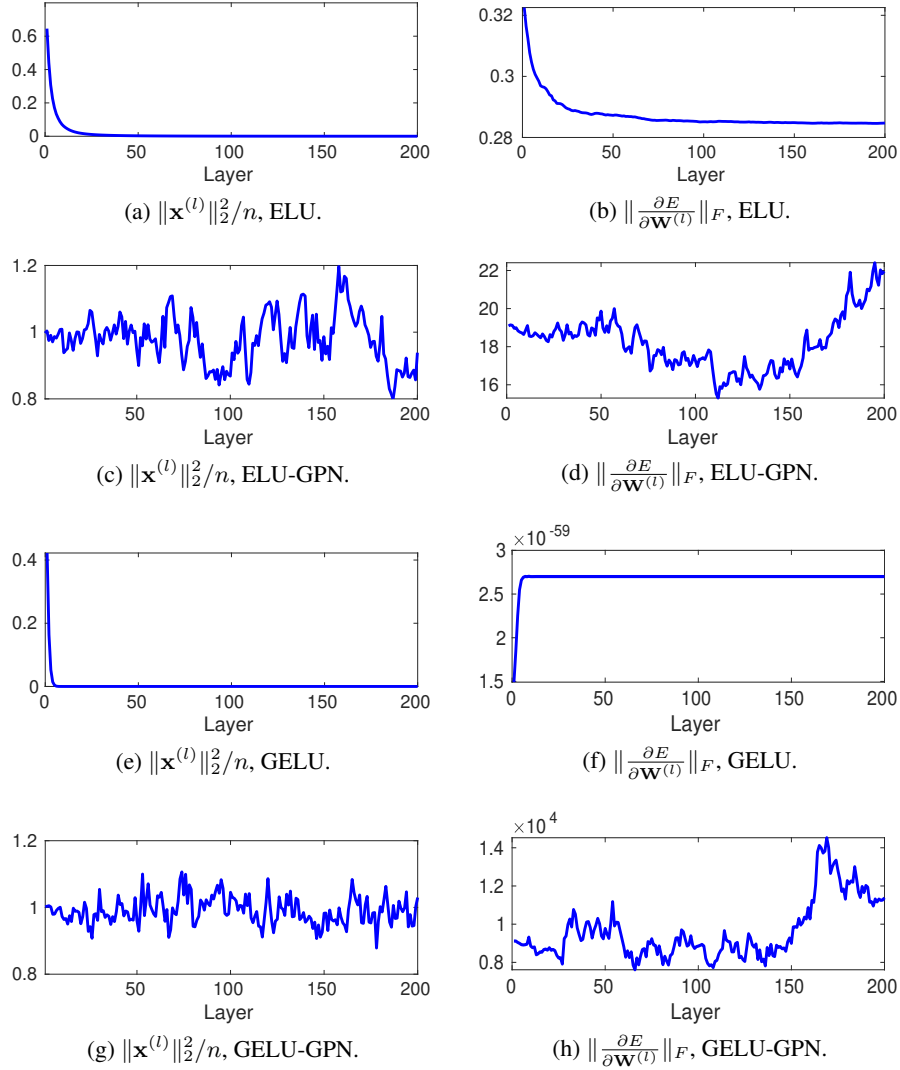


Figure 6: Results on synthetic data with different activation functions. “-GPN” denotes the function is Gaussian-Poincaré normalized. $\|\mathbf{x}^{(l)}\|_2$ denotes the l_2 norm of the outputs of the l -th layer. n denotes the width. $\|\frac{\partial E}{\partial \mathbf{W}^{(l)}}\|_F$ is the Frobenius norm of the gradient of the weight matrix in the l -th layer.




Figure 7: 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. Except for ELU, none of them has values concentrated around one.

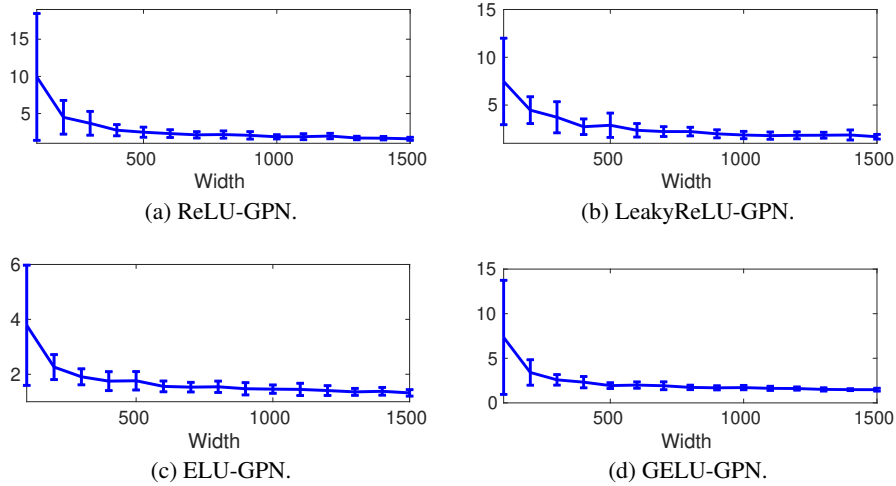


Figure 8: Gradient norm ratio for different layer width on synthetic data. The ratio is defined as $\max_l \left\| \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right\|_F / \min_l \left\| \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right\|_F$. The width ranges from 100 to 1500. The error bars show standard deviation.

B.2 Real-World Data

In Figure 9, we show the test accuracy during training on CIFAR-10 in addition to Figure 4 in the main text. In Figure 10, we show the experiments on MNIST.

In Figure 11, 12, 13 and 14, we show a measure of gradient exploding/vanishing during training for different activation functions. The measure is defined as the ratio of the maximum gradient norm and the minimum gradient norm across layers. Since we use the parametrization

$$\mathbf{W} = \left(\frac{\mathbf{v}_1}{\|\mathbf{v}_1\|_2}, \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|_2}, \dots, \frac{\mathbf{v}_n}{\|\mathbf{v}_n\|_2} \right)^T \quad (106)$$

with $\mathbf{V} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)^T$, the gradient norm ratio is defined on the unconstrained weights \mathbf{V} , that is,

$$\frac{\max_l \left\| \frac{\partial E}{\partial \mathbf{V}^{(l)}} \right\|_F}{\min_l \left\| \frac{\partial E}{\partial \mathbf{V}^{(l)}} \right\|_F}. \quad (107)$$

Note that for ReLU, LeakyReLU and GELU, the gradient vanishes during training in some experiments and therefore the plots are empty. From the figures, we can see that batch normalization leads to gradient explosion especially at the early stage of training. On the other hand, without batch normalization, the gradient is stable throughout training for GPN activation functions.




Figure 9: Test accuracy (percentage) during training on CIFAR-10.

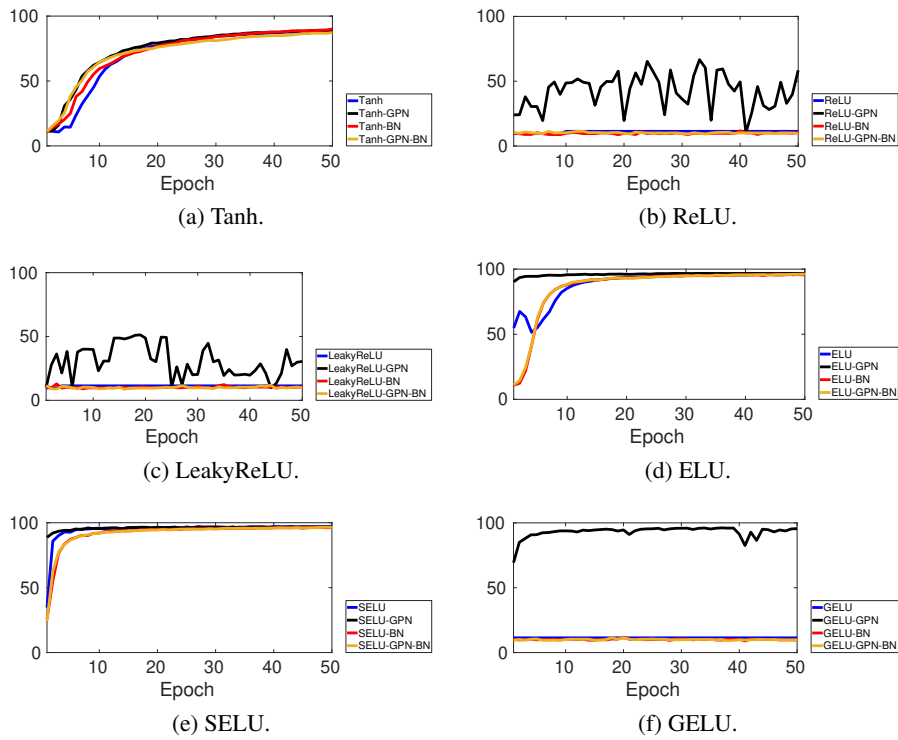


Figure 10: Test accuracy (percentage) during training on MNIST.

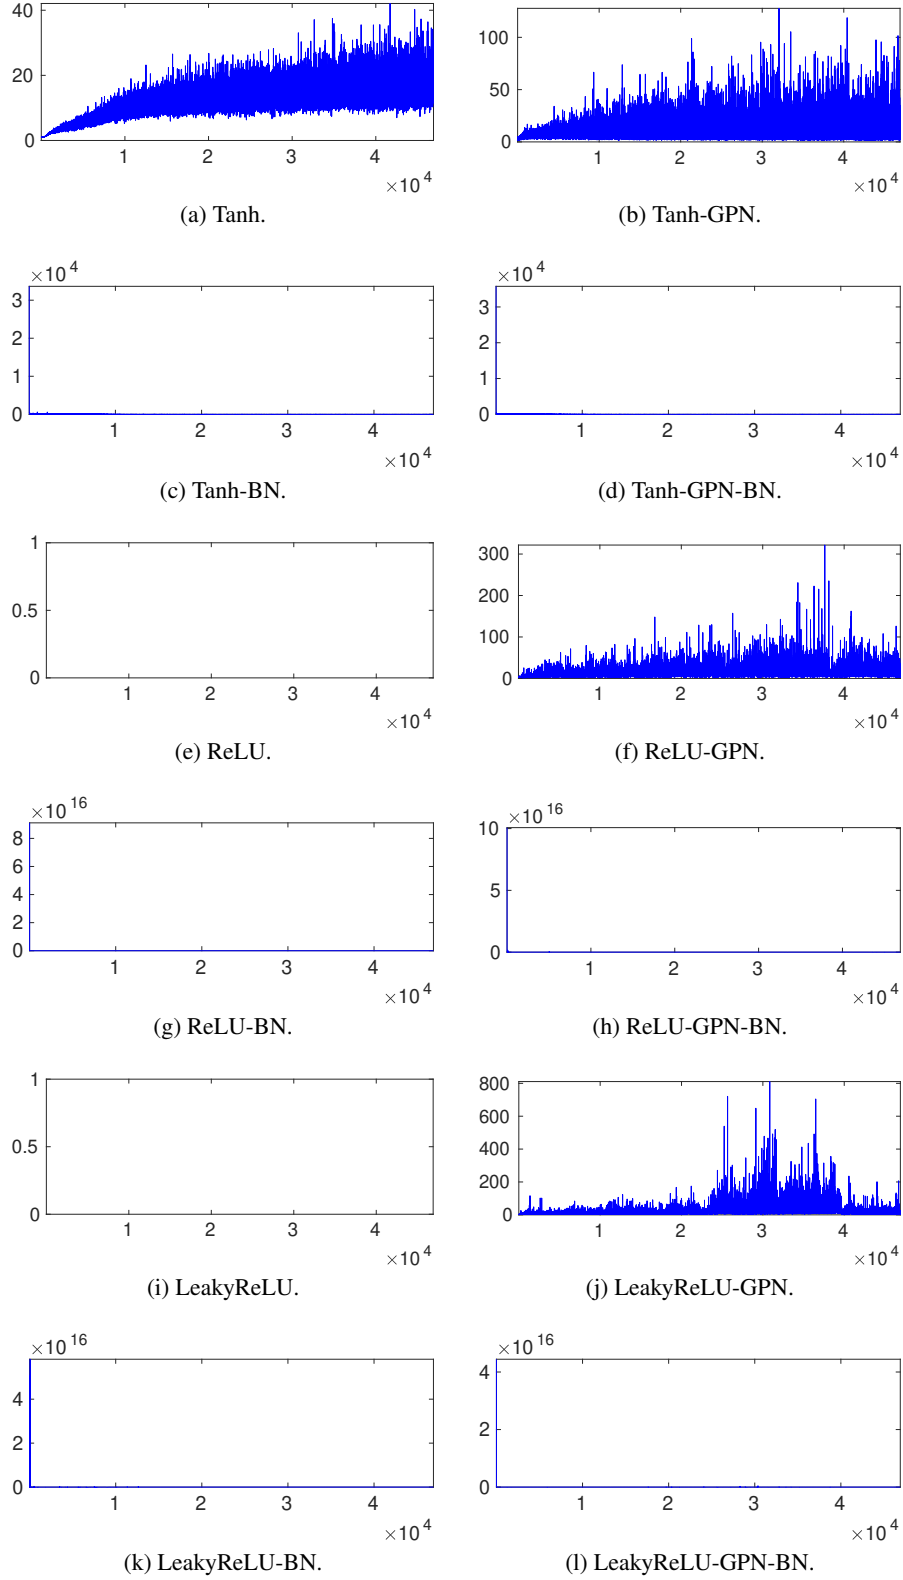


Figure 11: Gradient norm ratio during training on MNIST. Horizontal axis denotes the mini-batch updates. Vertical axis denotes the gradient norm ratio $\max_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F / \min_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F$. The gradient vanishes ($\|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F \approx 0$) for ReLU and LeakyReLU during training and hence the plots are empty.

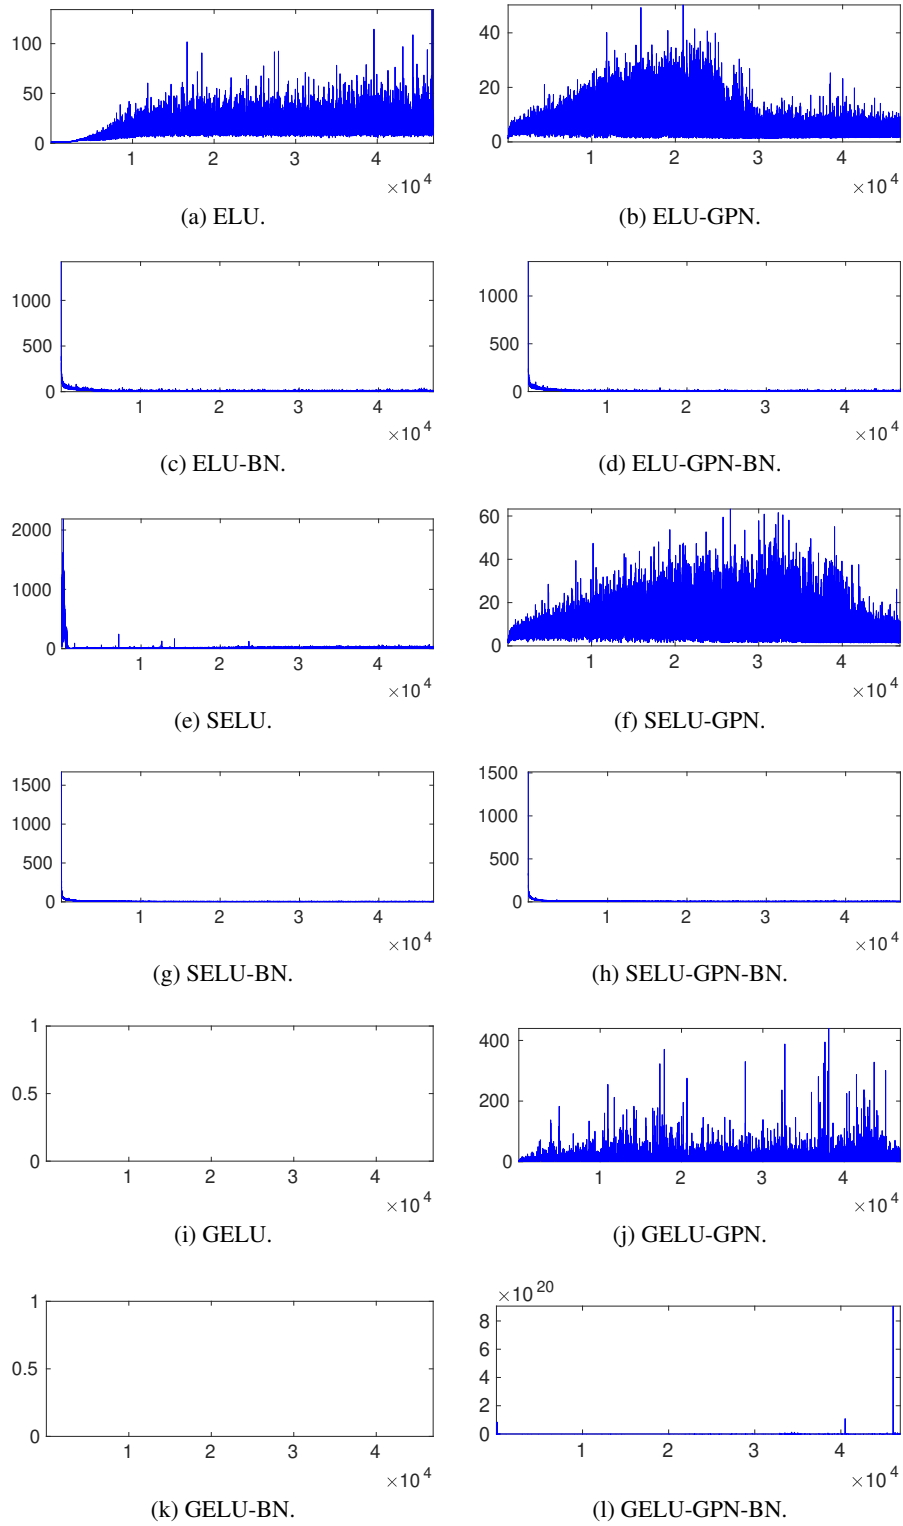


Figure 12: Gradient norm ratio during training on MNIST. Horizontal axis denotes the mini-batch updates. Vertical axis denotes the gradient norm ratio $\max_l \|\frac{\partial E}{\partial \mathbf{V}^{(l)}}\|_F / \min_l \|\frac{\partial E}{\partial \mathbf{V}^{(l)}}\|_F$. The gradient vanishes ($\|\frac{\partial E}{\partial \mathbf{V}^{(l)}}\|_F \approx 0$) for GELU and GELU-BN during training and hence the plots are empty.

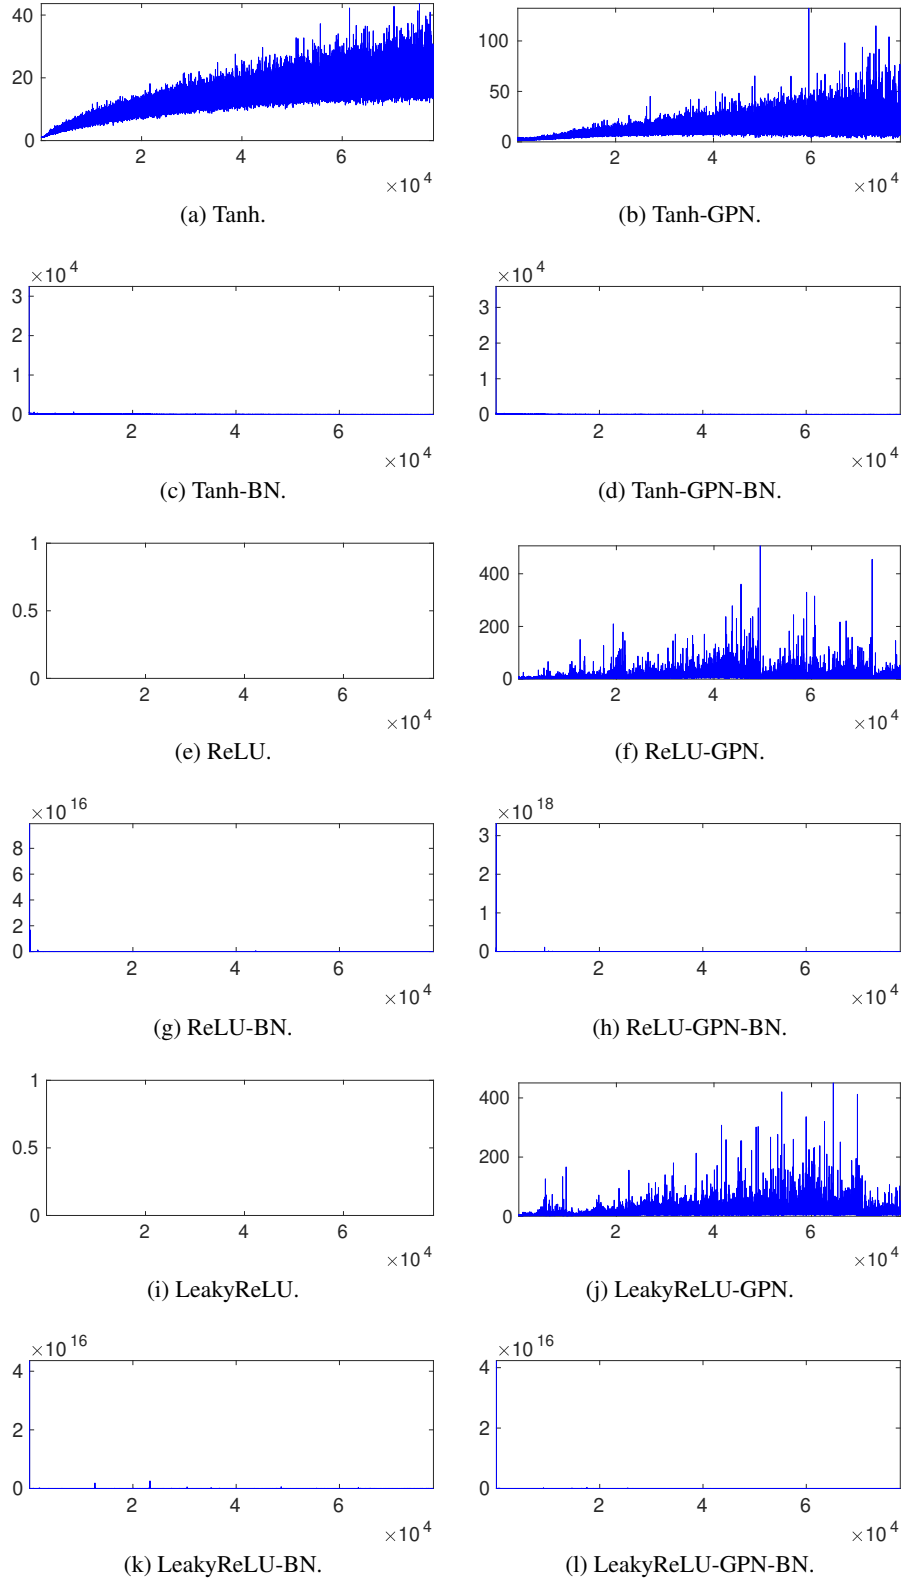


Figure 13: Gradient norm ratio during training on CIFAR-10. Horizontal axis denotes the mini-batch updates. Vertical axis denotes the gradient norm ratio $\max_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F / \min_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F$. The gradient vanishes ($\|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F \approx 0$) for ReLU and LeakyReLU during training and hence the plots are empty.

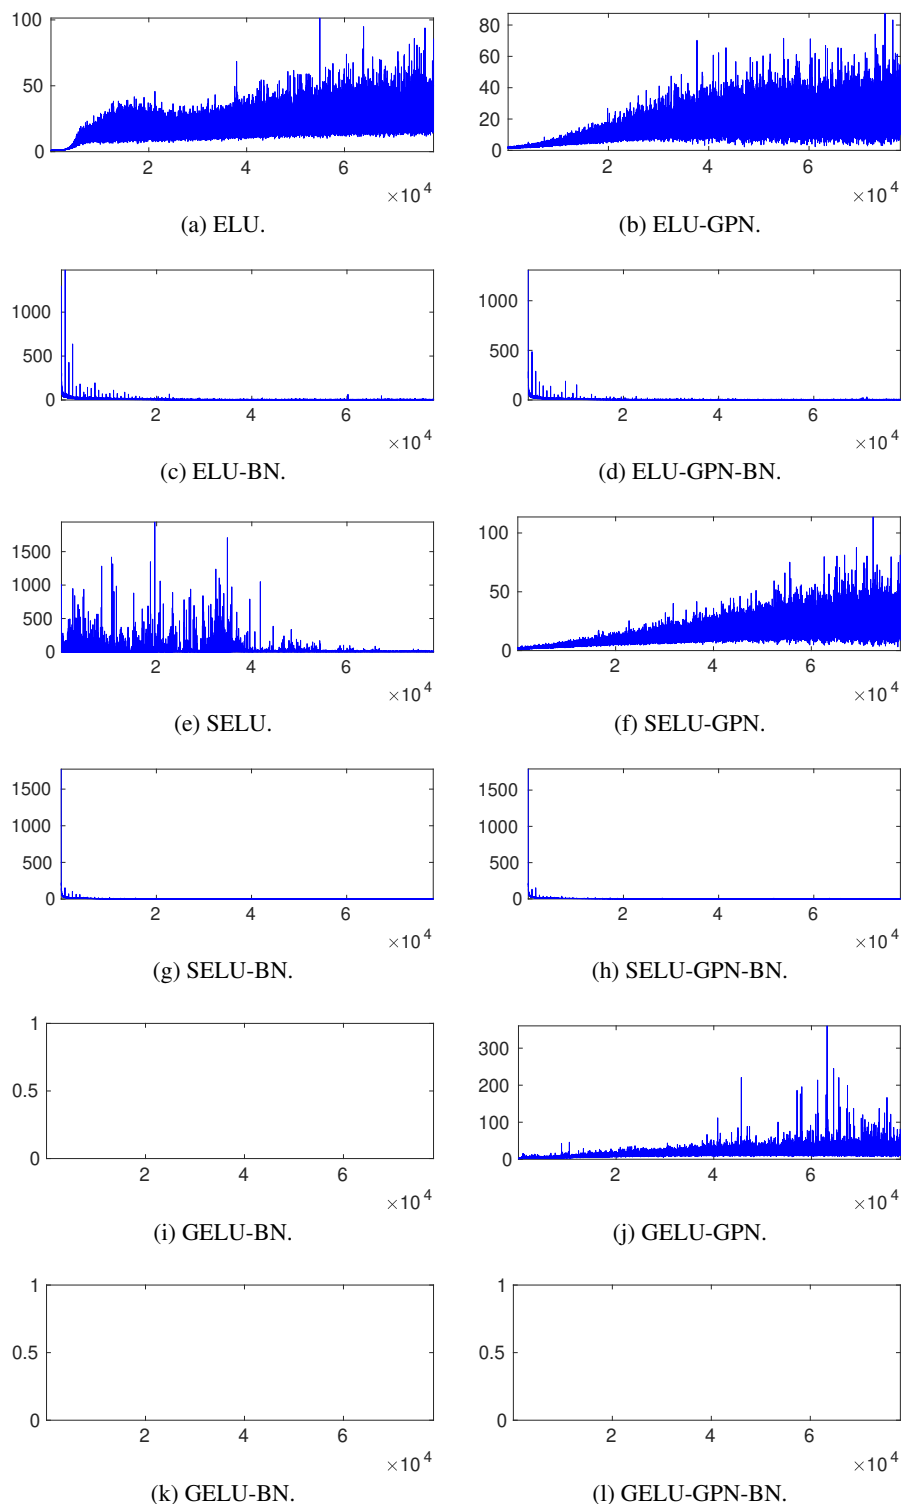


Figure 14: Gradient norm ratio during training on CIFAR-10. Horizontal axis denotes the mini-batch updates. Vertical axis denotes the gradient norm ratio $\max_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F / \min_l \|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F$. The gradient vanishes ($\|\frac{\partial E}{\partial \mathbf{V}(l)}\|_F \approx 0$) for GELU, GELU-BN and GELU-GPN-BN during training and hence the plots are empty. For GELU-GPN-BN, both gradient exploding and gradient vanishing are observed.