

ReLU Characteristic Activation Analysis

Wenlin Chen

University of Cambridge, Cambridge, UK
Max Planck Institute for Intelligent Systems, Tübingen, Germany

WC337@CAM.AC.UK

Hong Ge

University of Cambridge, Cambridge, UK

HG344@CAM.AC.UK

Abstract

We introduce a novel approach for analyzing the training dynamics of ReLU networks by examining the *characteristic activation boundaries* of individual ReLU neurons. Our proposed analysis reveals a critical instability in common neural network parameterizations and normalizations during stochastic optimization, which impedes fast convergence and hurts generalization performance. Addressing this, we propose *Geometric Parameterization (GmP)*, a novel neural network parameterization technique that effectively separates the radial and angular components of weights in the hyperspherical coordinate system. We show theoretically that GmP resolves the aforementioned instability issue. We report empirical results on various models and benchmarks to verify GmP’s theoretical advantages of optimization stability, convergence speed and generalization performance.

1. Introduction

In a standard neural network, each neuron performs an affine transformation on its input $\mathbf{x} \in \mathbb{R}^n$ followed by an element-wise nonlinear activation function g :

$$z = g(\mathbf{w}^T \mathbf{x} + b), \quad (1)$$

where the affine transformation is parameterized by a weight vector $\mathbf{w} \in \mathbb{R}^n$ and a bias scalar $b \in \mathbb{R}$. Rectified Linear Unit (ReLU) [11] is arguably the most popular activation function in modern deep learning architectures due to its simplicity and effectiveness:

$$g(s) = \text{ReLU}(s) = \max(0, s). \quad (2)$$

We introduce a novel concept for ReLU networks, called *characteristic activation boundary*, which is defined as the set of input locations with zero pre-activations. By definition, such boundaries separate the active and inactive regions of ReLU neurons in the input space, which play a critical role in ReLU networks since they serve as the fundamental building blocks for the decision boundaries.

In this work, we analyze the evolution dynamics of the characteristic activation boundaries in ReLU networks. Our novel analysis identifies a critical instability in many common parameterizations and normalizations that operate in the Cartesian coordinate, including Standard Parameterization, Weight Normalization and Batch Normalization. We show theoretically that this issue destabilizes the evolution of the characteristic boundaries in the presence of stochastic gradient noise, and empirically that it impedes fast convergence and hurts generalization performance. To address this, we introduce a novel neural network parameterization, named *Geometric Parameterization (GmP)*. As opposed to traditional parameterizations and normalizations which operate in the

Cartesian coordinate, GmP operates in the hyperspherical coordinate which automatically decouples the radial and angular components of the weights. Our theoretical results show that GmP stabilizes the evolution of characteristic activation boundaries during stochastic optimization. Our empirical studies confirm the efficacy of GmP on various models and benchmarks, including a ResNet trained on ImageNet. We report notable improvements in optimization stability, convergence speed, and generalization performance, which validate our theoretical hypotheses and results.

2. ReLU Characteristic Activation Analysis

This section defines characteristic activation boundary and its geometric connection to ReLU features, which serve as the basics of the proposed characteristic activation analysis for understanding neural network training dynamics.

2.1. Preliminary and Terminology

Standard Parameterization (SP) refers to the common weight-bias parameterization as defined in Equation (1). **Weight Normalization (WN)** [40] is a reparameterization technique that decouples the length l and the direction $\mathbf{v}/\|\mathbf{v}\|_2$ of \mathbf{w} in a standard ReLU unit (1):

$$z = \text{ReLU} \left(l \left(\frac{\mathbf{v}}{\|\mathbf{v}\|_2} \right)^T \mathbf{x} + b \right). \quad (3)$$

WN makes the length l and the direction $\mathbf{v}/\|\mathbf{v}\|_2$ of the weight vector independent of each other in the Cartesian coordinate system, which is effective in improving the conditioning of the gradients and thus speeding up optimization. **Batch Normalization (BN)** [18] is a widely-used normalization layer in modern deep learning architectures such as ResNet [16], which is effective in accelerating and stabilizing stochastic optimization of neural networks [21]. BN standardizes the pre-activation using the empirical mean and variance estimated using mini-batch statistics:

$$\text{BN}(\mathbf{w}^T \mathbf{x} + b) = \gamma \frac{\mathbf{w}^T \mathbf{x} - \hat{\mathbb{E}}_{\mathbf{x}}[\mathbf{w}^T \mathbf{x}]}{\sqrt{\hat{\text{Var}}_{\mathbf{x}}[\mathbf{w}^T \mathbf{x} + b]}} + \beta, \quad (4)$$

where $\gamma \in \mathbb{R}$ and $\beta \in \mathbb{R}$ are two free parameters to be learned from data, which adjusts the output of the BN layer as needed to increase its expressiveness.

2.2. ReLU Characteristic Activation Boundary

Noticing that the ReLU activation function (2) is active for positive arguments $s > 0$ and inactive for negative arguments $s < 0$, we introduce a novel concept called *characteristic activation boundary (CAB)* at the cut-off point $s = 0$ (i.e., with zero pre-activations), which will play a central role in our proposed characteristic activation analysis.

Definition 1 *The characteristic activation boundary (CAB) for a ReLU unit is defined by the set of input locations with zero pre-activations:*

$$\mathcal{B} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{w}^T \mathbf{x} + b = 0\}. \quad (5)$$

A CAB is an $(n - 1)$ -dimensional hyperplane that separates the active and inactive regions of a ReLU unit in the input space \mathbb{R}^n . Figure 1(a) visualizes a CAB in \mathbb{R}^2 .




Figure 1: (a) Characteristic activation boundary (CAB) \mathcal{B} (brown solid line) and spatial location $\phi = -\lambda \mathbf{u}(\theta)$ of a ReLU unit $z = \text{ReLU}(\mathbf{u}(\theta)^T \mathbf{x} + \lambda) = \text{ReLU}(\cos(\theta)x_1 + \sin(\theta)x_2 + \lambda)$ for inputs $\mathbf{x} \in \mathbb{R}^2$. The CAB forms a line in \mathbb{R}^2 , which acts as a boundary separating inputs into two regions. Green arrows denote the active region, and red arrows denote the inactive region. (b)-(e) Stability of the CAB of a ReLU unit in \mathbb{R}^2 under small perturbations $\varepsilon = \delta \mathbf{1}$ to the parameters. Solid lines denote characteristic activation boundaries \mathcal{B} , and colored dotted lines connect the origin and spatial locations ϕ of \mathcal{B} . Smaller changes between the perturbed and original boundaries imply higher stability. GmP is most stable against perturbations.

Definition 2 *The spatial location of a CAB is defined as*

$$\phi = -\frac{b \mathbf{w}}{\mathbf{w}^T \mathbf{w}} = -\frac{b}{\|\mathbf{w}\|_2} \frac{\mathbf{w}}{\|\mathbf{w}\|_2}, \quad (6)$$

which is a point on the corresponding CAB as $\mathbf{w}^T \phi + b = 0$. The vector that goes from the origin to ϕ specifies the shortest path between the origin and the CAB. Therefore, each spatial location uniquely determines a CAB. Figure 1(a) visualizes the spatial location of a CAB in \mathbb{R}^2 .

CABs play a critical role in ReLU networks, since they effectively specify the locations of ReLU features (i.e., non-linearities) which serve as the building blocks for the decision boundaries.

2.3. Instability of CABs During Stochastic Optimization

In the presence of stochastic gradient noise, we identify an instability in the evolution of CABs under common neural network parameterizations and normalizations.

Proposition 3 *A perturbation ε to the weight \mathbf{w} under SP (1) can result in an arbitrarily large change in the angular direction of the CAB if \mathbf{w} has a similar magnitude to ε ¹.*

The proof of Proposition 3 can be found in Appendix B.1. This indicates that CABs are vulnerable to small perturbations when \mathbf{w} has a small norm. This has the implication that even a small gradient noise could destabilize the evolution of ReLU features during stochastic optimization and thus destroy the learning signal for the decision boundaries of the network. Such instability prevents practitioners from using larger learning rates [12].

It might be tempting to think that WN does not suffer from this instability since it decouples the length l and the direction $\mathbf{v}/\|\mathbf{v}\|_2$ of \mathbf{w} as shown in Equation (3). However, we show that WN is also vulnerable to small perturbations.

1. $\|\mathbf{w}\|_2$ is supposed to be small during training as large weights would lead to overfitting, numerical instability and even divergence [12] (e.g., the popular weight decay method explicitly regularizes $\|\mathbf{w}\|_2$ to be close to zero).

Proposition 4 A perturbation $(\varepsilon, \varepsilon')$ to (\mathbf{v}, l) under WN (3) can result in an arbitrarily large change in the angular direction of the CAB if \mathbf{v} has a similar magnitude to ε .

The proof of Proposition 4 can be found in Appendix B.2. Furthermore, BN also suffers from this instability issue since it can be viewed as a kind of weight normalization.

Proposition 5 Without loss of generality, assume that the input \mathbf{x} has zero mean. A perturbation ε to the weight \mathbf{w} under BN (4) can result in an arbitrarily large change in the angular direction of the CAB if \mathbf{w} has a similar magnitude to ε .

The proof of Proposition 5 can be found in Appendix B.3. Figures 1(b)-1(e) simulate the evolution behaviors of the CABs under SP and WN in \mathbb{R}^2 , showing that even a small perturbation δ of magnitude 10^{-3} can drastically change the spatial locations of the CABs. More generally, this instability issue exists in every neural network parameterization and normalization technique that operates in the Cartesian coordinate, since the fundamental issue is that the change in angular direction of the CAB always has the same unstable form as in Equation (12). In the next section, we will address this issue with a new parameterization technique that operates in a different coordinate system.

3. Geometric Parameterization

This section introduces *Geometric Parameterization (GmP)* and demonstrates its nice theoretical property which addresses the instability issue in common parameterizations and normalizations as shown in the previous section.

3.1. CABs in the Hyperspherical Coordinate System

In a high dimensional input space, most data points live in a thin shell since the volume of a high dimensional space concentrates near its surface [4]. Intuitively, we want the spatial locations of CABs to be close to the thin shell where most data points live, since this spatial affinity between CABs and data points will introduce ReLU features (non-linearities) at suitable locations in the input space to separate different inputs \mathbf{x} by assigning them different activation values. The use of hyperspherical coordinate enables us to explicitly control the locations of such non-linearities.

Definition 6 The spatial location of a CAB in the hyperspherical coordinate system is given by

$$\phi(\lambda, \boldsymbol{\theta}) = -\lambda \mathbf{u}(\boldsymbol{\theta}), \quad (7)$$

where the radius λ corresponds to $b/\|\mathbf{w}\|_2$ in SP, and the unit directional vector $\mathbf{u}(\boldsymbol{\theta})$ corresponds to $\mathbf{w}/\|\mathbf{w}\|_2$ in SP and is determined by the angle $\boldsymbol{\theta} = [\theta_1, \dots, \theta_{n-1}]^T$:

$$\mathbf{u}(\boldsymbol{\theta}) = \begin{bmatrix} \cos(\theta_1) \\ \sin(\theta_1) \cos(\theta_2) \\ \sin(\theta_1) \sin(\theta_2) \cos(\theta_3) \\ \vdots \\ \sin(\theta_1) \sin(\theta_2) \cdots \sin(\theta_{n-2}) \cos(\theta_{n-1}) \\ \sin(\theta_1) \sin(\theta_2) \cdots \sin(\theta_{n-2}) \sin(\theta_{n-1}) \end{bmatrix}. \quad (8)$$

$\mathbf{u}(\boldsymbol{\theta}) \in S^{n-1}$ is a unit directional vector on the unit hypersphere $S^{n-1} := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$.

Definition 7 *The CAB in the hyperspherical coordinate system is given by*

$$\mathcal{B}(\lambda, \boldsymbol{\theta}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{u}(\boldsymbol{\theta})^T \mathbf{x} + \lambda = 0\}. \quad (9)$$

Geometrically speaking, the angle $\boldsymbol{\theta}$ controls the direction of a CAB, while the radius λ controls the distance between the origin and the CAB. Calculating the pre-activation of a ReLU unit for an input \mathbf{x} is equivalent to projecting \mathbf{x} onto the unit vector $\mathbf{u}(\boldsymbol{\theta})$ and then adding the radius λ to the signed norm of the projected vector. From this perspective, it is clear a CAB is a set of inputs whose projections over $\mathbf{u}(\boldsymbol{\theta})$ have signed norm $-\lambda$. For this reason, we refer to this radial-angular decomposition in the hyperspherical coordinate system as *Geometric Parameterization* (GmP).

3.2. Geometric Parameterization of ReLU Networks

Definition 8 *A ReLU unit under GmP is given by*

$$z = r \text{ReLU}(\mathbf{u}(\boldsymbol{\theta})^T \mathbf{x} + \lambda). \quad (10)$$

GmP has three learnable parameters $r, \lambda, \boldsymbol{\theta}$. The radial and angular parameters λ and $\boldsymbol{\theta} = [\theta_1, \dots, \theta_{n-1}]$ specify the spatial location of the CAB, while the scaling parameter r controls the scale of the activation. These parameters have $n + 1$ degrees of freedom in total (same as SP).

3.3. Theoretical Analysis

Section 2.3 identified an instability in the evolution of CABs under common parameterizations and normalizations due to a fundamental issue in the Cartesian coordinate system. In contrast, CABs under GmP should be much more stable under perturbation, since the radius λ and angle $\boldsymbol{\theta}$ of the spatial location ϕ are automatically disentangled in the hyperspherical coordinate system. This means that small perturbations to the parameters in GmP will only cause small changes in the spatial location of the CAB. Below, we show that under a small perturbation, the change in the angular direction of a CAB under GmP is bounded by the magnitude of the perturbation.

Theorem 9 *With an infinitesimal perturbation $\boldsymbol{\varepsilon} := [\varepsilon_1, \dots, \varepsilon_{n-1}]^T$ to the angular parameter $\boldsymbol{\theta}$, the change in the angular direction $\mathbf{u}(\boldsymbol{\theta}) \in S^{n-1}$ ($n \geq 2$) of a CAB under GmP is given by*

$$\langle \mathbf{u}(\boldsymbol{\theta}), \mathbf{u}(\boldsymbol{\theta} + \boldsymbol{\varepsilon}) \rangle = \sqrt{\varepsilon_1^2 + \sum_{i=2}^{n-1} \left(\prod_{j=1}^{i-1} \sin^2(\theta_j) \right) \varepsilon_i^2}. \quad (11)$$

The proof of Theorem 9 can be found in Appendix B.4, which is based on an elegant idea from differential geometry that $\langle \mathbf{u}(\boldsymbol{\theta}), \mathbf{u}(\boldsymbol{\theta} + \boldsymbol{\varepsilon}) \rangle$ is simply the norm $\|\boldsymbol{\varepsilon}\|_{\mathbf{M}}$ of the perturbation with respect to the metric tensor \mathbf{M} for the hyperspherical coordinate.

Corollary 10 *The change in $\boldsymbol{\theta}$ in Theorem 9 is upper bounded: $\langle \mathbf{u}(\boldsymbol{\theta}), \mathbf{u}(\boldsymbol{\theta} + \boldsymbol{\varepsilon}) \rangle \leq \|\boldsymbol{\varepsilon}\|_2$.*

Corollary 10 implies that optimizing the geometric parameters in GmP directly translates into a smooth evolution of the spatial locations of CABs even in the presence of stochastic gradient noise. Back to Figures 1(b)-1(e), CABs under GmP gradually moves away from its original spatial location as we increase δ , which is in sharp contrast to the unstable evaluation of CABs under SP and WN.

Appendix C contains additional remarks on GmP. Appendix D contains a thorough empirical evaluation to verify our theoretical analysis, showcasing GmP’s significant improvements in optimization stability, convergence speed and generalization performance.

Acknowledgements

We thank Isaac Reid for helpful feedback and discussions. WC acknowledges funding via a Cambridge Trust Scholarship (supported by the Cambridge Trust) and a Cambridge University Engineering Department Studentship (under grant G105682 NMZR/089 supported by Huawei R&D UK). HG acknowledges generous support from Huawei R&D UK.

Part of this work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk).

References

- [1] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2): 251–276, 1998.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] Xuchan Bao, James Lucas, Sushant Sachdeva, and Roger B Grosse. Regularized linear autoencoders recover the principal components, eventually. *Advances in Neural Information Processing Systems*, 33:6971–6981, 2020.
- [4] Avrim Blum, John Hopcroft, and Ravindran Kannan. *Foundations of Data Science*. Cambridge University Press, 2020. doi: 10.1017/9781108755528.
- [5] Vitaliy Chiley, Ilya Sharapov, Atli Kosson, Urs Koster, Ryan Reece, Sofia Samaniego de la Fuente, Vishal Subbiah, and Michael James. Online normalization for training neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [6] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [8] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- [9] Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In *International Conference on Machine Learning*, pages 2232–2241. PMLR, 2019.
- [10] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.

- [11] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.
- [12] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [13] Boris Hanin and David Rolnick. Complexity of linear regions in deep networks. In *International Conference on Machine Learning*, pages 2596–2604. PMLR, 2019.
- [14] Boris Hanin and David Rolnick. Deep relu networks have surprisingly few activation patterns. *Advances in neural information processing systems*, 32, 2019.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [17] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *Advances in neural information processing systems*, 30, 2017.
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015.
- [19] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Jonas Kohler, Hadi Daneshmand, Aurelien Lucchi, Thomas Hofmann, Ming Zhou, and Klaus Neymeyr. Exponential convergence rates for batch normalization: The power of length-direction decoupling in non-convex optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 806–815. PMLR, 2019.
- [22] Atli Kosson, Bettina Messmer, and Martin Jaggi. Rotational equilibrium: How weight decay balances learning across neural networks. In *NeurIPS 2023 Workshop on Mathematics of Modern Machine Learning*, 2023.
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [24] Frederik Kunstner, Philipp Hennig, and Lukas Balles. Limitations of the empirical fisher approximation for natural gradient descent. *Advances in neural information processing systems*, 32, 2019.

- [25] Antoine Labatie, Dominic Masters, Zach Eaton-Rosen, and Carlo Luschi. Proxy-normalizing activations to match batch normalization while removing batch dependence. *Advances in Neural Information Processing Systems*, 34:16990–17006, 2021.
- [26] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- [27] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.
- [28] Zhiyuan Li and Sanjeev Arora. An exponential learning rate schedule for deep learning. *arXiv preprint arXiv:1910.07454*, 2019.
- [29] Zhiyuan Li, Kaifeng Lyu, and Sanjeev Arora. Reconciling modern deep learning with traditional optimization analyses: The intrinsic learning rate. *Advances in Neural Information Processing Systems*, 33:14544–14555, 2020.
- [30] Senwei Liang, Zhongzhan Huang, Mingfu Liang, and Haizhao Yang. Instance enhancement batch normalization: An adaptive regulator of batch noise. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 4819–4827, 2020.
- [31] Ping Luo, Jiamin Ren, Zhanglin Peng, Ruimao Zhang, and Jingyu Li. Differentiable learning-to-normalize via switchable normalization. *arXiv preprint arXiv:1806.10779*, 2018.
- [32] Siyuan Ma, Raef Bassily, and Mikhail Belkin. The power of interpolation: Understanding the effectiveness of sgd in modern over-parametrized learning. In *International Conference on Machine Learning*, pages 3325–3334. PMLR, 2018.
- [33] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [34] Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. *Advances in neural information processing systems*, 27, 2014.
- [35] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [36] Vardan Papyan. Traces of class/cross-class structure pervade deep learning spectra. *The Journal of Machine Learning Research*, 21(1):10197–10260, 2020.
- [37] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *international conference on machine learning*, pages 2847–2854. PMLR, 2017.
- [38] Simon Roburin, Yann Dubois de Mont-Marin, Andrei Bursuc, Renaud Marlet, Patrick Perez, and Mathieu Aubry. A spherical analysis of adam with batch normalization, 2021. URL <https://openreview.net/forum?id=jHykXSik3ch>.

- [39] David Rolnick and Konrad Kording. Reverse-engineering deep relu networks. In *International conference on machine learning*, pages 8178–8187. PMLR, 2020.
- [40] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016.
- [41] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? *Advances in neural information processing systems*, 31, 2018.
- [42] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [43] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [44] Twan Van Laarhoven. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*, 2017.
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [46] Ruosi Wan, Zhanxing Zhu, Xiangyu Zhang, and Jian Sun. Spherical motion dynamics: Learning dynamics of normalized neural network using sgd and weight decay. *Advances in Neural Information Processing Systems*, 34:6380–6391, 2021.
- [47] Geoffrey S Watson. Distributions on the circle and sphere. *Journal of Applied Probability*, 19 (A):265–280, 1982.
- [48] Andrew TA Wood. Simulation of the von mises fisher distribution. *Communications in statistics-simulation and computation*, 23(1):157–164, 1994.
- [49] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [50] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [51] Shuangfei Zhai, Tatiana Likhomanenko, Etai Littwin, Jason Ramapuram, Dan Busbridge, Yizhe Zhang, Jiatao Gu, and Joshua M. Susskind. σ reparam: Stable transformer training with spectral reparameterization, 2023.
- [52] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. *arXiv preprint arXiv:1810.12281*, 2018.
- [53] Guodong Zhang, James Martens, and Roger B Grosse. Fast convergence of natural gradient descent for over-parameterized neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

- [54] Hongyi Zhang, Yann N Dauphin, and Tengyu Ma. Fixup initialization: Residual learning without normalization. *arXiv preprint arXiv:1901.09321*, 2019.




Figure 2: Visualization of characteristic activation boundaries (brown solid lines) and spatial locations $\phi = -\lambda u(\theta)$ of a ReLU unit $z = \text{ReLU}(u(\theta)x + \lambda)$ (blue solid lines) for inputs $x \in \mathbb{R}$. Green arrows denote active regions and red arrows denote inactive regions.




Figure 3: Visualization of characteristic activation boundaries (brown solid lines) and spatial locations $\phi = -\lambda \mathbf{u}(\theta)$ of a ReLU unit $z = \text{ReLU}(\mathbf{u}(\theta)^\top \mathbf{x} + \lambda) = \text{ReLU}(\cos(\theta)x_1 + \sin(\theta)x_2 + \lambda)$ for inputs $\mathbf{x} \in \mathbb{R}^2$. Green arrows denote active regions and red arrows denote inactive regions.

Appendix A. More Visualizations of the Characteristic Activation Boundaries

Figures 2 and 3 respectively visualize the characteristic activation boundaries in the input spaces \mathbb{R} and \mathbb{R}^2 under different conditions of the radius and angle parameters.

Appendix B. Technical Proofs

B.1. Proof of Proposition 3

Proof The change in the angular direction of the CAB under SP is given by

$$\langle \mathbf{w}, \mathbf{w} + \boldsymbol{\varepsilon} \rangle = \arccos \left(\frac{\mathbf{w}^\top (\mathbf{w} + \boldsymbol{\varepsilon})}{\|\mathbf{w}\|_2 \|\mathbf{w} + \boldsymbol{\varepsilon}\|_2} \right). \quad (12)$$

As $\|\mathbf{w}\|_2$ is small, we can construct a small perturbation $\boldsymbol{\varepsilon} = -(1 + \delta) \mathbf{w}$ with an infinitesimal δ . Then, $\langle \mathbf{w}, \mathbf{w} + \boldsymbol{\varepsilon} \rangle = \langle \mathbf{w}, -\delta \mathbf{w} \rangle = \pi$, which rotates the CAB by 180° . This completes the proof. In general, $\langle \mathbf{w}, \mathbf{w} + \boldsymbol{\varepsilon} \rangle$ can take arbitrary values in $[0, \pi]$ even for a small perturbation $\boldsymbol{\varepsilon}$ in this case. ■

B.2. Proof of Proposition 4

Proof The change in the angular direction of the CAB under WN is given by

$$\left\langle l \frac{\mathbf{v}}{\|\mathbf{v}\|_2}, (l + \varepsilon') \frac{\mathbf{v} + \boldsymbol{\varepsilon}}{\|\mathbf{v} + \boldsymbol{\varepsilon}\|_2} \right\rangle = \arccos \left(\frac{\mathbf{v}^\top (\mathbf{v} + \boldsymbol{\varepsilon})}{\|\mathbf{v}\|_2 \|\mathbf{v} + \boldsymbol{\varepsilon}\|_2} \right), \quad (13)$$

which has an identical form to Equation (12). Therefore, we can similarly construct a small perturbation $\varepsilon = -(1 + \delta) \mathbf{v}$ with an infinitesimal δ , which rotates the CAB by 180° . This completes the proof. \blacksquare

B.3. Proof of Proposition 5

Proof Since $\hat{\mathbb{E}}_{\mathbf{x}}[\mathbf{w}^T \mathbf{x}] = \mathbf{w}^T \hat{\mathbb{E}}_{\mathbf{x}}[\mathbf{x}] = \mathbf{0}$ by assumption, Equation (4) can be re-written as

$$\text{BN}(\mathbf{w}^T \mathbf{x} + b) = \gamma \frac{\mathbf{w}^T \mathbf{x}}{\sqrt{\mathbf{w}^T \hat{\Sigma} \mathbf{w}}} + \beta = \gamma \left(\frac{\mathbf{w}}{\|\mathbf{w}\|_{\hat{\Sigma}}} \right)^T \mathbf{x} + \beta, \quad (14)$$

where the norm $\|\mathbf{w}\|_{\hat{\Sigma}}$ is with respect to the data covariance matrix $\hat{\Sigma} = \hat{\text{Var}}[\mathbf{x}]$ estimated using mini-batch statistics. It can be seen that BN has the same form as WN (3) except that WN fixes $\hat{\Sigma} = \mathbf{I}$. Therefore, the instability argument for WN in the proof of Proposition 4 also holds for BN. This completes the proof. \blacksquare

B.4. Proof of Theorem 9

Proof The main idea of this proof is that the change in the angular direction under infinitesimal perturbation ε is the norm of ε with respect to the metric tensor \mathbf{M} for the hyperspherical coordinate system (since \mathbf{u} has unit length):

$$\langle \mathbf{u}(\boldsymbol{\theta}), \mathbf{u}(\boldsymbol{\theta} + \varepsilon) \rangle = \|\varepsilon\|_{\mathbf{M}} = \sqrt{\varepsilon^T \mathbf{M} \varepsilon}. \quad (15)$$

Therefore, we need to work out a formula for calculating \mathbf{M} .

Let us start with an input $\mathbf{x} = [x_1, \dots, x_n]^T \in \mathbb{R}^n$ ($n \geq 2$) in the Cartesian coordinate system, where the metric tensor is the Kronecker delta $m'_{ij} = \delta_{ij}$. For the geometric parameterization of the unit hypersphere S^{n-1} , we have

$$\begin{aligned} u_1(\boldsymbol{\theta}) &= \cos(\theta_1), \\ u_2(\boldsymbol{\theta}) &= \sin(\theta_1) \cos(\theta_2), \\ u_3(\boldsymbol{\theta}) &= \sin(\theta_1) \sin(\theta_2) \cos(\theta_3), \\ &\vdots \\ u_{n-2}(\boldsymbol{\theta}) &= \sin(\theta_1) \sin(\theta_2) \sin(\theta_3) \cdots \sin(\theta_{n-2}) \cos(\theta_{n-1}), \\ u_{n-1}(\boldsymbol{\theta}) &= \sin(\theta_1) \sin(\theta_2) \sin(\theta_3) \cdots \sin(\theta_{n-2}) \sin(\theta_{n-1}). \end{aligned} \quad (16)$$

The metric tensor \mathbf{M} for the geometric parameterization of S^{n-1} is the pullback of the Euclidean metric in \mathbb{R}^n :

$$m_{ab} = \sum_{i=1}^{n-1} \sum_{j=1}^{n-1} m'_{ij} \frac{\partial u_i}{\partial \theta_a} \frac{\partial u_j}{\partial \theta_b} = \sum_{i=1}^{n-1} \frac{\partial u_i}{\partial \theta_a} \frac{\partial u_i}{\partial \theta_b}. \quad (17)$$

- $a \neq b$: we have $m_{ab} = 0$, since it is a sum of terms that are either zero or with alternating signs which cancel out. Hence, \mathbf{M} is a diagonal matrix.

- $a = b$: Using the fact that $\sin^2(\theta_j) + \cos^2(\theta_j) = 1, \forall j$, we have $m_{11} = 1$ and

$$m_{aa} = \prod_{i=1}^{a-1} \sin^2(\theta_i), \quad 2 \leq a \leq n-1. \quad (18)$$

Therefore, the metric tensor for the hyperspherical coordinate is a diagonal matrix

$$\mathbf{M} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \sin^2(\theta_1) & 0 & \cdots & 0 & 0 \\ 0 & 0 & \sin^2(\theta_1) \sin^2(\theta_2) & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \prod_{i=1}^{n-2} \sin^2(\theta_i) & 0 \\ 0 & 0 & 0 & \cdots & 0 & \prod_{i=1}^{n-1} \sin^2(\theta_i) \end{bmatrix}. \quad (19)$$

Finally, the change in the angular direction of a unit vector $\mathbf{u}(\boldsymbol{\theta})$ under a perturbation $\boldsymbol{\varepsilon}$ to $\boldsymbol{\theta}$ is given by the norm of $\boldsymbol{\varepsilon}$ with respect to the tensor matrix \mathbf{M} :

$$\langle \mathbf{u}(\boldsymbol{\theta}), \mathbf{u}(\boldsymbol{\theta} + \boldsymbol{\varepsilon}) \rangle = \sqrt{\boldsymbol{\varepsilon}^T \mathbf{M} \boldsymbol{\varepsilon}} = \sqrt{\varepsilon_1^2 + \sum_{i=2}^{n-1} \left(\prod_{j=1}^{i-1} \sin^2(\theta_j) \right) \varepsilon_i^2}. \quad (20)$$

This completes the proof. ■

B.5. Proof of Corollary 10

Proof Since $-1 \leq \sin(\theta_j) \leq 1$, it follows that $0 \leq m_{i,i} = \prod_{j=1}^{i-1} \sin^2(\theta_j) \leq 1$ for all i . Therefore,

$$\sqrt{\varepsilon_1^2 + \sum_{i=2}^{n-1} \left(\prod_{j=1}^{i-1} \sin^2(\theta_j) \right) \varepsilon_i^2} \leq \sqrt{\varepsilon_1^2 + \sum_{i=2}^{n-1} \varepsilon_i^2} = \|\boldsymbol{\varepsilon}\|_2.$$

This completes the proof. ■

Appendix C. Remarks on Geometric Parameterization

C.1. Computational Cost

Let n =fan-in and m =fan-out for a layer. GmP needs to compute $2n - 2$ more scalars

$$\sin(\theta_1), \dots, \sin(\theta_{n-1}), \cos(\theta_1), \dots, \cos(\theta_{n-1}) \quad (21)$$

than SP for each of the m neurons, which incur an extra cost of $\mathcal{O}(mn)$ for all neurons in a layer. However, since the cost of computing the affine transformation for each layer is also $\mathcal{O}(mn)$, the total computational cost of GmP remains $\mathcal{O}(mn)$ for each layer, which is the same as SP.

C.2. Layer-Size Independent Parameter Initialization

Unlike existing neural network parameterizations which are sensitive to initialization, GmP can work with less carefully chosen initialization schemes independent of the width of the layer, thanks to an invariant property of the hyperspherical coordinate system. To see this, we first consider the distribution of the angular direction of the CAB under SP. Under popular initialization methods such as the Glorot initialization [10] and He initialization [15], each element in the initial weight vector \mathbf{w} under SP is independently and identically sampled from a zero mean Gaussian distribution with a layer-size dependent variance. Interestingly, this always induces a uniform distribution over the unit n -sphere for the direction of the CAB, no matter what variance value is used in that Gaussian distribution. This allows us to initialize the angular parameter θ uniformly at random by sampling from the von Mises–Fisher distribution [47, 48]. The parameter λ is initialized to zero due to its connection $\lambda = b/\|\mathbf{w}\|_2$ to SP and the common practice to set $b = 0$ at initialization. The scaling parameter r is initialized to one, based on the intuition that the scale roughly corresponds to the total variance of the weights \mathbf{w} in SP. We highlight that none of the parameters λ , θ , and r in GmP requires layer-size dependent initialization.

C.3. Internal Covariate Shift

One implicit assumption of our characteristic activation analysis is that the input distribution to a neuron always centers around the origin during training. This assumption automatically holds for one-hidden-layer networks since the training data can be centered during data pre-processing. However, this assumption is not necessarily satisfied for multi-hidden-layer networks, since the inputs to an intermediate layer are transformed by the weights and squashed by the activation function in the previous layer. We propose a simple technique called Input Mean Normalization (IMN), which is a *parameter-free* layer that centers the input to each intermediate layer using the empirical mean estimated by mini-batch statistics:

$$z = r \operatorname{ReLU}(\mathbf{u}(\theta)^T(\mathbf{x} - \hat{\mathbb{E}}[\mathbf{x}]) + \lambda). \quad (22)$$

It might be tempting to think that IMN is similar to Mean-only Batch Normalization (MBN) [40]. However, we emphasize that MBN is unable to address the internal covariate shift problem as it is applied to pre-activations rather than post-activations.

Appendix D. Empirical Evaluation

This section contains empirical evaluation of GmP on both illustrative demonstrations and challenging classification and regression benchmarks with neural network architectures of different sizes optimized by SGD and Adam. A detailed setup for each experiment can be found in Appendix E. Code for reproducing the experimental results is available at <https://github.com/Wenlin-Chen/geometric-parameterization>.

D.1. Illustrative Experiments

This section verifies the validity of the hypotheses of our proposed characteristic activation analysis on two illustrative experiments aided with visualization, and demonstrates that the improved stability under GmP is beneficial for neural network optimization and generalization.

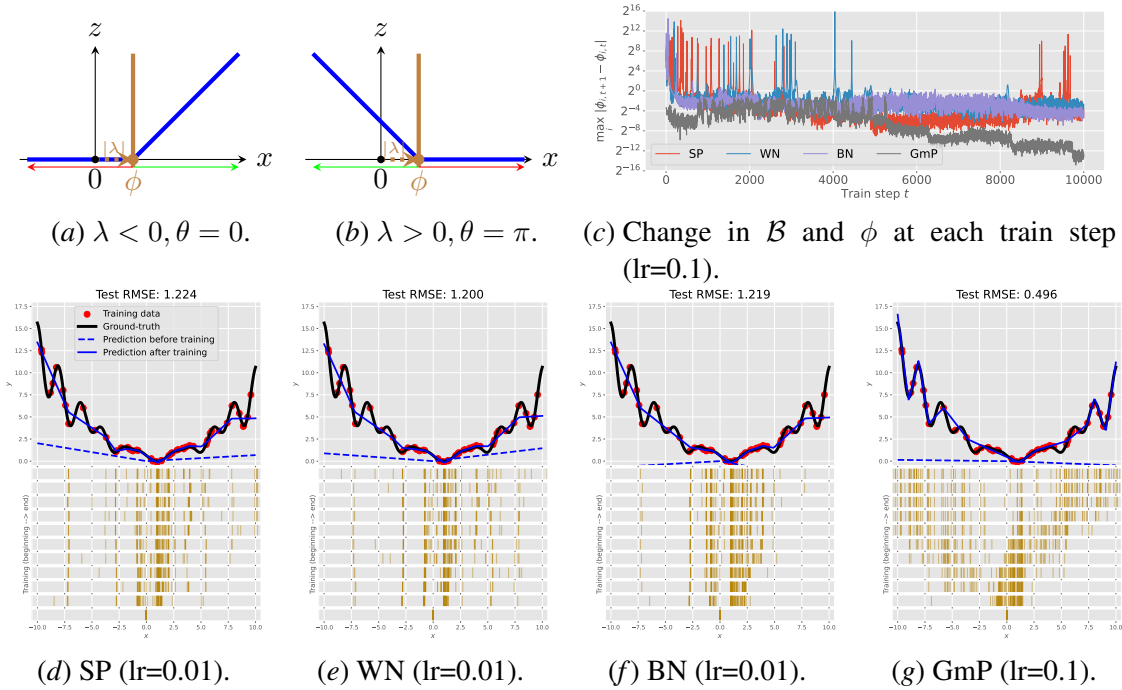


Figure 4: (a)-(b) Characteristic activation point \mathcal{B} (intersection of brown solid lines and the x -axis) and spatial location $\phi = -\lambda u(\theta)$ of a ReLU unit $z = \text{ReLU}(u(\theta)x + \lambda)$ (blue solid lines) for inputs $x \in \mathbb{R}$. Green arrows denote active regions, and red arrows denote inactive regions. (c) Evolution dynamics of the characteristic points \mathcal{B} in a one-hidden-layer network with 100 ReLU units for a 1D Levy regression problem under SP, WN, BN and GmP during training. Smaller values are better as they indicate higher stability of the evolution of the characteristic points during training. The y -axis is in \log_2 scale. (d)-(g): The top row illustrates the experimental setup, including the network’s predictions at initialization and after training, and the training data and the ground-truth function (Levy). Bottom row: the evolution of the characteristic activation point for the 100 ReLU units during training. Each horizontal bar shows the spatial location spectrum for a chosen optimization step, moving from the bottom (at initialization) to the top (after training with Adam). More spread of the spatial locations covers the data better and adds more useful non-linearities to the model, making prediction more accurate. Regression accuracy is measured by root mean squared error (RMSE) on a separate test set. Smaller RMSE values are better. We use cross-validation to select the learning rate for each method. The optimal learning rate for SP, WN, and BN is lower than that for GmP, since their training becomes unstable with higher learning rates, as shown in (c).

D.1.1. 1D LEVY REGRESSION

In Figure 4, we train a one-hidden-layer network with 100 ReLU units under SP, WN, BN and GmP on the 1D Levy regression dataset using Adam [20]. As shown in Figures 4(a)-4(b), Both the CAB \mathcal{B} and its spatial location ϕ reduce to the same point in \mathbb{R} , which will be referred to as the characteristic activation point. The angle θ of the characteristic activation point can only take two values 0 or π , corresponding to the two directions on the real line. Clearly, GmP significantly improves the stability of the evolution of the characteristic activation point and allows us to use

GEOMETRIC PARAMETERIZATION

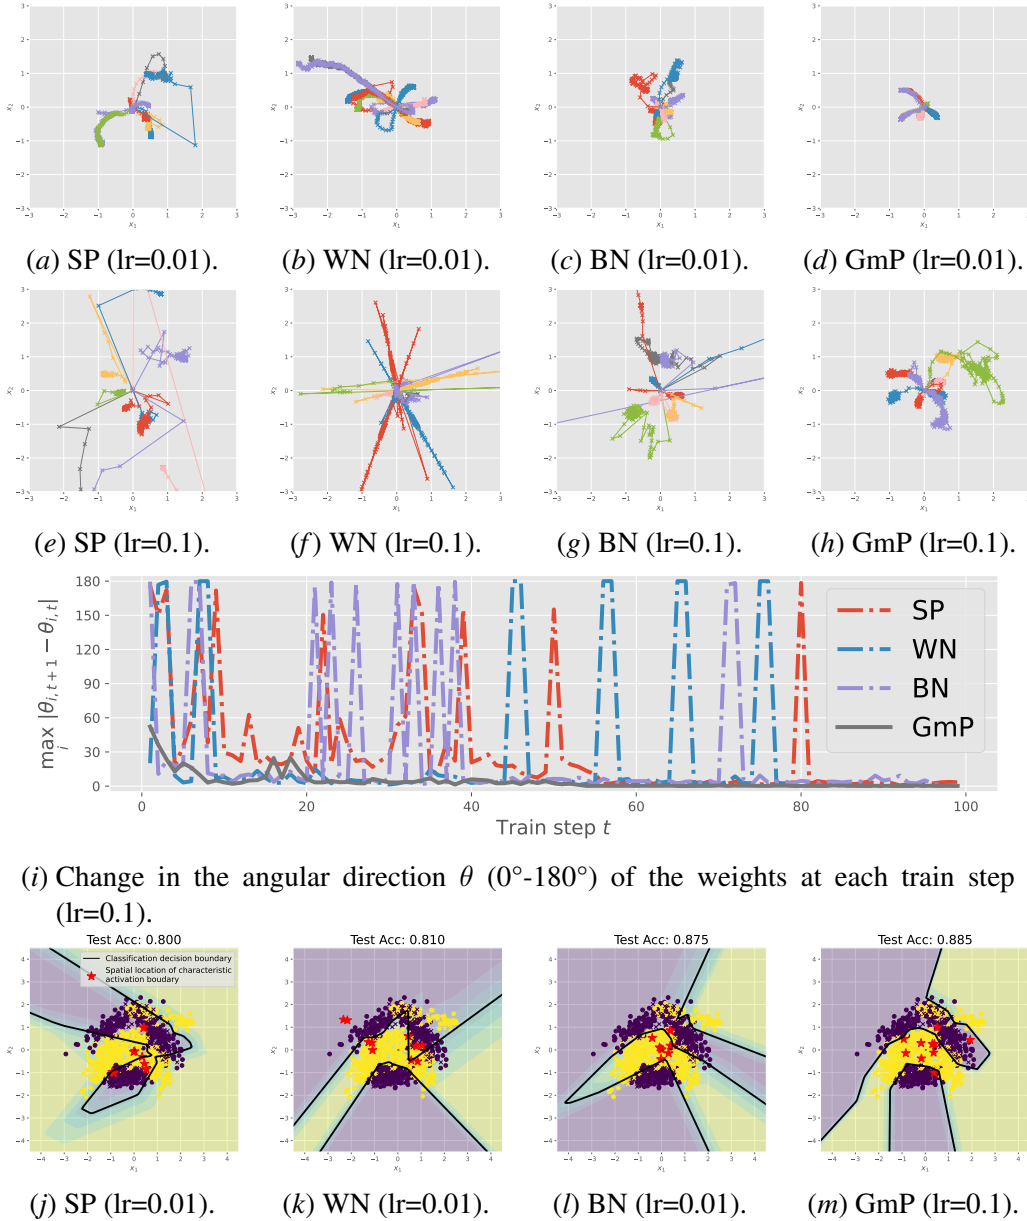


Figure 5: Performance of a single-hidden-layer neural network with 10 ReLU units on the 2D Banana classification dataset under SP, WN, BN and GmP trained using Adam. (a)-(h): Trajectories of the spatial locations of the 10 ReLU units during training. Each color depicts one ReLU unit. Smoother evolution means higher training stability. The evolution under GmP is stable, so we can use a $10\times$ larger learning rate. (i): Evolution dynamics of the angular direction θ of CABs. Smaller values are better as they indicate higher robustness against stochastic gradient noise. (j)-(m): Network predictions after training. Black bold lines depict the classification boundary between two classes. Classification accuracy is measured on a separate test set. Higher accuracy values are better. The red stars show the spatial locations of 10 ReLU units. Intuitively speaking, more evenly spread out red stars are better for classification accuracy, as they provide more useful non-linearity.

Table 1: Test RMSE for MLPs trained on six UCI benchmarks.

Benchmark	Boston	Concrete	Energy	Power	Wine	Yacht
SP	3.370 ± 0.145	5.472 ± 0.144	0.898 ± 0.274	4.065 ± 0.029	0.623 ± 0.008	0.639 ± 0.063
WN	3.459 ± 0.156	5.952 ± 0.148	2.093 ± 0.789	4.073 ± 0.026	0.632 ± 0.008	0.624 ± 0.076
BN	3.469 ± 0.153	5.695 ± 0.160	1.648 ± 0.302	4.164 ± 0.026	0.622 ± 0.011	0.777 ± 0.055
GmP	3.057 ± 0.144	5.153 ± 0.098	0.474 ± 0.013	4.022 ± 0.025	0.613 ± 0.006	0.584 ± 0.046

a $10\times$ large learning rate as selected by cross validation. Figure 4(c) shows that under GmP the maximum change $\max_i |\Delta\phi_{i,t}| = \max_i |\phi_{i,t+1} - \phi_{i,t}|$ at each train step t is always smaller than one throughout training, while under other parameterizations the changes can be up to 2^{16} at some steps. The stable evolution of the characteristic point under GmP leads to improved optimization stability and generalization performance (i.e., the best test RMSE) on this regression task, as shown in Figures 4(d)-4(g).

D.1.2. 2D BANANA CLASSIFICATION

In Figure 5, we train a one-hidden-layer network with 10 ReLU units under SP, WN, BN and GmP on the 2D Banana classification dataset using Adam. Figures 5(a)-5(h) show that GmP allows us to use a $10\times$ larger learning rate (as selected by cross validation) while maintaining a smooth evolution of the characteristic activation boundary. Figure 5(i) shows that GmP is the only method that guarantees stable updates for the angular directions of the CAB during training with a large learning rate: under GmP, the maximum change $\max_i |\Delta\theta_{i,t}| = \max_i |\theta_{i,t+1} - \theta_{i,t}|$ at each train step t remains low throughout training, while under other parameterizations the change can be up to 180° at some steps. This verifies the hypothesis in our proposed characteristic activation analysis. Figures 5(j)-5(m) show that under GmP, the spatial locations of CABs move towards different directions during training and spread over all training data points in different regions, which forms a classification decision boundary with a reasonable shape that achieves the best generalization performance (i.e., the highest test accuracy) among all compared methods.

D.2. Machine Learning Benchmarks

D.2.1. UCI REGRESSION WITH MLP

We further evaluate GmP on six UCI regression datasets [8]. We train an MLP with one hidden layer and 100 hidden units for 10 different random 80/20 train/test splits. We use the Adam optimizer [20]. We use cross-validation and find that the optimal learning rate is 0.1 for GmP and 0.01 for all the other methods. Table 1 shows that GmP consistently achieves the best test RMSE on all benchmarks, significantly outperforming other methods in most cases.

D.2.2. IMAGENET CLASSIFICATION WITH RESNET

Finally, we evaluate GmP with a gold-standard large residual neural network ResNet-34 [16] on the full ImageNet (ILSVRC 2012) dataset [7], which consists of 1,281,167 training images and 50,000 validation images that contain objects from 1,000 categories. The size of the images ranges from 75×56 to 4288×2848 . We follow exactly the same experimental setup for optimization and data augmentation as in He et al. [16]. Specifically, we use the SGD optimizer with momentum 0.9, which turns out to be better than Adam for image classification tasks [16]. We reduce the learning

Table 2: Validation acc. (%) for ResNet-34 trained on ImageNet.

Metric	Top-1 valid. acc.	Top-5 valid. acc.
WN+MBN	73.48 \pm 0.16	91.12 \pm 0.12
BN	74.76 \pm 0.09	91.61 \pm 0.05
GmP+IMN	75.57 \pm 0.12	92.68 \pm 0.11

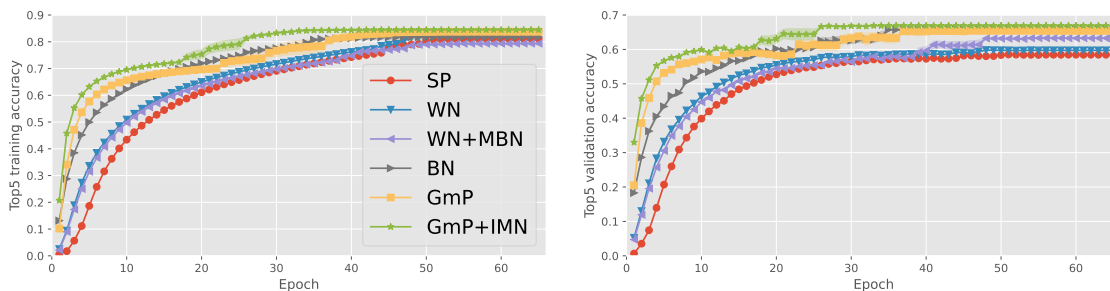
rate by 0.1 at epochs 30, 60 and 80. All models are trained for 90 epochs. We use a batch size of 256 for all methods. We use cross-validation and find that the optimal initial learning rate is 0.1 for all compared methods. We employ random horizontal flip, random resizing (256-480) with preserved aspect ratio, random crop (224), and color augmentation for data augmentation during training [23]. To address the internal covariate shift problem, we employ Input Mean Normalization (IMN) for GmP. Following Salimans and Kingma [40], Mean-only Batch Normalization (MBN) is used for WN. Table 2 reports the single-center-crop top-1 and top-5 validation accuracy for all compared methods, which shows that GmP+IMN significantly outperforms BN and WN+MBN in terms of both top-1 and top-5 validation accuracy. This demonstrates that our method is useful for improving large-scale residual network training.

D.2.3. ABLATION STUDY: IMAGENET32 CLASSIFICATION WITH VGG

We perform ablation study to provide further insights into how the batch size and intermediate normalization layer affect the convergence speed and generalization performance of different parameterizations. To maintain a manageable computational cost, we conduct these experiments with a medium-sized convolutional neural network VGG-6 [42] on ImageNet32 [6], which contains all 1.3M images and 1,000 categories from ImageNet (ILSVRC 2012) [7], but with the images resized to 32×32 . We follow exactly the same experimental setup for optimization and data augmentation as in Chrabaszcz et al. [6]. We use cross-validation and find that the optimal initial learning rate is 0.1 for GmP and 0.01 for all the other methods. Table 3 shows that GmP+IMN consistently achieves the best top-1 and top-5 validation accuracy for all batch sizes considered. Furthermore, the improvement of GmP+IMN over other methods gets larger as the batch size increases, highlighting the robustness and scalability of GmP with large batch sizes. In addition to achieving the best performance, Figure 6 shows that GmP+IMN (the green curve) also converges significantly faster than other compared methods: its top-5 validation accuracy converges within 25 epochs, which is 10 epochs earlier than the second best method BN. The ablation study GmP vs GmP+IMN shows that IMN significantly improves the performance of GmP, which is expected since it addresses the problem of covariate shifts between hidden layers. Notably, Wide ResNet (WRN 28-2) [50] trained with BN and a batch size of 500 only achieved 43.08% top-1 validation accuracy as reported in Chrabaszcz et al. [6], underperforming VGG-6 trained with GmP+IMN (43.62% as shown in Table 3). This reveals the significance of better parameterizations: *even a small non-residual network like VGG-6 with GmP+IMN can outperform large, wide residual networks like WRN 28-2.*

Table 3: Top-1 and top-5 validation accuracy (%) for VGG-6 trained on ImageNet32.

Metric	Top-1 validation accuracy			Top-5 validation accuracy		
	256	512	1024	256	512	1024
SP	38.31 ± 0.13	36.99 ± 0.11	35.02 ± 0.03	62.48 ± 0.14	60.71 ± 0.18	58.14 ± 0.39
WN	39.13 ± 0.10	37.92 ± 0.12	36.17 ± 0.03	63.28 ± 0.02	61.93 ± 0.09	60.16 ± 0.18
WN+MBN	42.22 ± 0.01	40.96 ± 0.02	39.33 ± 0.07	66.04 ± 0.07	65.08 ± 0.03	63.32 ± 0.08
BN	42.79 ± 0.03	41.90 ± 0.19	41.39 ± 0.02	67.17 ± 0.08	66.50 ± 0.25	65.89 ± 0.06
GmP	40.76 ± 0.09	41.65 ± 0.09	41.29 ± 0.08	65.08 ± 0.08	65.76 ± 0.05	65.49 ± 0.06
GmP+IMN	43.14 ± 0.05	43.62 ± 0.08	42.70 ± 0.15	67.36 ± 0.05	67.76 ± 0.09	66.98 ± 0.18



(a) Top-5 training accuracy.

(b) Top-5 validation accuracy.

Figure 6: Comparison of convergence speed for VGG-6 trained on the ImageNet32 dataset with a batch size of 1024.

Appendix E. Detailed Experimental Setups

E.1. UCI Regression with MLP

We train an MLP with one hidden layer and 100 hidden units for 10 different random 80/20 train/test splits. We use the Adam optimizer [20] with full-batch training. We use cross-validation to select the learning rate for each compared method from the set $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$. We find that the optimal initial learning rate is 0.1 for GmP and 0.01 for all the other compared methods. We report test root mean squared error (RMSE).

E.2. ImageNet Classification with ResNet

We train a ResNet-34 [16] on the ImageNet (ILSVRC 2012) dataset [7], which consists of 1,281,167 training images and 50,000 validation images that contain objects from 1,000 categories. The size of the images ranges from 75×56 to 4288×2848 . We follow exactly the same experimental setup for optimization and data augmentation as in He et al. [16]. We use the SGD optimizer with momentum 0.9, which turns out to be better than Adam for image classification tasks [16]. We reduce the learning rate by 0.1 at epochs 30, 60 and 80. All models are trained for 90 epochs. We use a batch size of 256 for all methods. We use cross-validation to select the learning rate for each compared method from the set $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$. We find that the optimal initial learning rate is 0.1 for all compared methods. We employ random horizontal flip, random resizing (256-480) with preserved aspect ratio, random crop (224), and color augmentation for data augmentation during training [23]. To address the internal covariate shift problem, we employ Input

Mean Normalization (IMN) for GmP. Following [40], Mean-only Batch Normalization (MBN) is used for WN. We report single-center-crop top-1 and top-5 validation accuracy.

E.3. Ablation Study: ImageNet32 Classification with VGG

To maintain a manageable computational cost for the ablation study, we train a VGG-6 [42] on ImageNet32 [6], which contains all 1.3M images and 1,000 categories from ImageNet (ILSVRC 2012) [7], but with the images resized to 32×32 . We follow exactly the same experimental setup for optimization and data augmentation as in Chrabaszcz et al. [6]. We use the SGD optimizer with momentum 0.9, which turns out to be better than Adam for image classification tasks [16]. We reduce the learning rate by 0.1 at epochs 30, 60 and 80. All models are trained for 90 epochs. We train the model using three common batch sizes $\{256, 512, 1024\}$ for all methods. We use cross-validation to select the learning rate for each compared method from the set $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$. We find that the optimal initial learning rate is 0.1 for GmP and 0.01 all the other methods. We employ random horizontal flips for data augmentation during training. We conduct an ablation study to explore the effects of input mean normalization (IMN) for GmP and mean batch normalization (MBN) for WN in deep networks. We report top-1 and top-5 validation accuracy.

Appendix F. Related Work

F.1. Neural Network Training Dynamics.

Neural Tangent Kernels [19, 27] show that wide networks evolve like linear models during training, while Neural Network Gaussian Processes [26, 35] provide insights into how wide neural networks generalize. [9] studies the evolution of the Hessian spectrum of neural networks during training. [3] investigates the curvatures of different principle components around the optimum of a regularized linear autoencoder. [1, 24, 53] analyze the training dynamics with natural gradient descent (NGD); see Appendix F.3 for a discussion of the connection between GmP and NGD. [32] investigates the effectiveness of stochastic gradient descents for neural network training. [5, 17, 22, 25, 28, 29, 36, 38, 41, 44, 46, 52, 54] study the effects of Adam, BN and weight decay on training dynamics. [13, 14, 34, 37, 39] pay special attention to analyzing the ReLU activations in neural networks. In contrast, our proposed characteristic activation analysis studies the evolution of the characteristic activation boundaries during stochastic optimization.

F.2. Parameterization and Normalization.

In addition to SP, WN [40] and BN [18], there are many other neural network normalization and parameterization techniques. Instead of normalizing the batch dimension as in BN, LayerNorm [2] operates on the feature dimension, which is preferred for small batches or variable-length inputs such as text [45]. Other variants of BN include SwitchNorm [31] and IEBN [30]. There are also normalization techniques designed for specific applications. For instance, InstanceNorm [43] and GroupNorm [49] are special cases of LayerNorm designed for CNNs, while spectral normalization [33, 51] is specifically designed for GANs and transformers. All these techniques operate in the Cartesian coordinate and thus suffers from the instability issue, whereas GmP operates in the hyperspherical coordinate to overcome this instability issue.

F.3. Connections to Natural Gradient Descent

The parameter space of a neural network can be thought of as a Riemannian manifold M , for which the neural network parameterization specifies the coordinate for M . In this paper, we reveal that standard parameterization is vulnerable to small perturbations of the parameters (e.g., SGD noise) whereas our proposed geometric parameterization is much more robust against perturbations.

One thing to note is that the full version of natural gradient descent is invariant to neural network parameterizations since it is defined in an abstract form without any specific parameterization/coordinate. In the abstract form, we let the loss function be $l(\theta) = -\log p(Y|X, \theta)$ where θ are the abstract parameters and (X, Y) is an abstract data point. The fisher matrix $F(\theta) = \mathbb{E}_{p(X, Y)}[\nabla l(\theta)\nabla l(\theta)^T]$ is the metric tensor for M , where $p(X, Y)$ is the abstract data distribution. A single step of the full version of natural gradient descent is given by

$$\theta_{t+1} \leftarrow \theta_t - \eta_t \cdot \text{Exp}_{\theta_t}[F(\theta_t)^{-1}\nabla h(\theta_t)], \quad (23)$$

where $h : M \rightarrow \mathbb{R}$ is any differentiable function and η_t is the learning rate. Note that the exponential map Exp_{θ_t} maps the update $F(\theta_t)^{-1}\nabla h(\theta_t)$ from the tangent space $T_{\theta_t}M$ back to the manifold M , which results in an exact update invariant to the parameterization/coordinate. However, it is often not practical to calculate the exponential map since it requires solving a linear system of size $\dim(M)$, which is the total number of parameters in the neural network.

The most commonly used version of natural gradient descent in practice is its first-order approximation given by

$$\theta_{t+1} \leftarrow \theta_t - \eta_t \cdot F(\theta_t)^{-1}\nabla h(\theta_t), \quad (24)$$

which is a second-order optimization method that is invariant up to first-order transformation of the parameterization/coordinate. Since our proposed geometric parameterization is a nonlinear transformation of standard parameterization, it will still make a difference under the first order approximation of the full version of natural gradient descent.

Appendix G. Limitations and Future Work

Theoretically, this work analyzed ReLU networks due to their wider adoption. However, the general characteristic activation analysis technique can be extended to other activation functions. Also, we only performed the characteristic activation analysis for single-hidden-layer ReLU networks and proposed a practical workaround to address the problem of covariate shift between hidden layers by using input mean normalization. For future work, this analysis needs to be generalized to examine training dynamics in multiple-hidden-layer networks to understand the theoretical behavior of deep networks. One potential difficulty with multiple-hidden-layer networks is that the characteristic activation boundary becomes a piecewise linear partition of the input space, which is less straightforward to analyze. A possible solution would be to consider how the assignment of each data point to the partition evolves during training, similar to how we track the characteristic activation boundaries. Empirically, we believe that our experiments are sufficient to verify our theory and support our claims. For future work, it would be interesting to investigate the theoretical properties and empirical performance of a combination of GmP and existing normalization techniques (e.g., BN). We also leave the experiments of examining the performance of GmP with larger neural network architectures (e.g., transformers) on large datasets from different domains (e.g., NLP) for future work.