

---

# A Method on Searching Better Activation Functions

---

Haoyuan Sun<sup>\*,1</sup>, Zihao Wu<sup>\*,2</sup>, Bo Xia<sup>1</sup>, Pu Chang<sup>3</sup>, Zibin Dong<sup>2</sup>, Yifu Yuan<sup>2</sup>,  
Yongzhe Chang<sup>†,1</sup>, Xueqian Wang<sup>†,1</sup>

<sup>\*</sup>Equal Contribution   <sup>†</sup>Corresponding Authors

<sup>1</sup>Tsinghua University   <sup>2</sup>Tianjin University   <sup>3</sup>Anhui Polytechnic University

## Abstract

The success of artificial neural networks (ANNs) hinges greatly on the judicious selection of an activation function, introducing non-linearity into network and enabling them to model sophisticated relationships in data. However, the search of activation functions has largely relied on empirical knowledge in the past, lacking theoretical guidance, which has hindered the identification of more effective activation functions. In this work, we offer a proper solution to such issue. Firstly, we theoretically demonstrate the existence of the worst activation function with boundary conditions (WAFBC) from the perspective of information entropy. Furthermore, inspired by the Taylor expansion form of information entropy functional, we propose the Entropy-based Activation Function Optimization (EAFO) methodology. EAFO methodology presents a novel perspective for designing static activation functions in deep neural networks and the potential of dynamically optimizing activation during iterative training. Utilizing EAFO methodology, we derive a novel activation function from ReLU, known as Correction Regularized ReLU (CRReLU). Experiments conducted with vision transformer and its variants on CIFAR-10, CIFAR-100 and ImageNet-1K datasets demonstrate the superiority of CRReLU over existing corrections of ReLU. Extensive empirical studies on task of large language model (LLM) fine-tuning, CRReLU exhibits superior performance compared to GELU, suggesting its broader potential for practical applications.

## 1 Introduction

Flourishing development of artificial intelligence is predominantly attributable to rapid advancements in artificial neural networks (ANNs) observed in recent years. Activation functions (AFs) play a critical role in the performance of ANNs due to their fundamental role in enabling nonlinear representations. Despite continuous development of novel activation functions and their empirical success in improving network performance, theoretical analysis towards these activation functions remain scarce in the research literature. In other words, proposal of improved activation functions is often based on empirical evidence without theoretical validations, which greatly hinders the search for better activation functions. Hence, a theoretically reliable methodology on searching better activation functions holds significant value for the machine learning community and future research.

In our work, we initiate our exploration from the correlation between information entropy and Bayesian error rate. Subsequently, we establish the connection between activation function and information entropy, ultimately deriving the specific form that the worst activation function does exist under boundary conditions. Based on the derivation, we propose a novel method for optimizing activation functions, namely **the Entropy-based Activation Function Optimization (EAFO)** methodology. Utilizing EAFO methodology, we derive a novel activation function known as **Correction Regularized ReLU (CRReLU)** with the beginning of conventional ReLU [1–3]. The derived CRReLU activation function possesses several desirable properties, including the avoidance of neuron death,

the preservation of neuron sparsity, and so on. Experiments involving the vision transformer [4] and its variants [5, 6], conducted on CIFAR-10, CIFAR-100 [7] and ImageNet-1K [8] datasets, have consistently demonstrated the superior performance of CRReLU compared to other activation function baselines. Extensive experimental studies on the task of large language model (LLM) fine-tuning with direct preference optimization (DPO) method [9] also demonstrate that CRReLU surpasses GELU in performance, suggesting the wider applicability of CRReLU in practical scenarios. Moreover, the EAFO methodology also shows potential to further optimize activation functions during the iterative training of ANNs, although the specific optimization techniques remain a topic of ongoing research.

In summary, our main contributions are as follows:

- We theoretically prove the existence of the worst activation function with boundary conditions from the perspective of information entropy; and starting from the worst activation function, performance of activation functions always improves.
- We propose the Entropy-based Activation Function Optimization (EAFO) methodology, which provides a novel perspective for designing static activation functions in deep neural networks and the potential of dynamically optimizing activation during iterative training.
- We derive a novel activation function known as Correction Regularized ReLU (CRReLU) starting from ReLU utilizing the EAFO methodology. Experiments across several mainstream architectures, datasets and tasks demonstrate that the proposed CRReLU exceeds existing activation functions, exhibiting exceptional performance.

## 2 Methodology

### 2.1 Methodology Problem Setup

#### 2.1.1 Bayesian Error Rate and Information Entropy

A deep neural network can be simplified as comprising a feature extraction layer, which is subsequently followed by a fully connected layer for final classification. From a probabilistic perspective, in binary classification, the feature extraction layer can be conceptualized as transforming the shape of mixture distribution, thereby enabling the final fully connected layer to separate two distributions with a hyperplane. Hence, the more overlapping two distributions are, the higher Bayesian error rate and the worse classification performance. Furthermore, a lower information entropy corresponds to a higher likelihood of forming two distinct peaks (i.e. the smaller classification uncertainty, the easier to classify); and an increase in the overlap between two distributions also leads to the increase of information entropy (i.e. the greater classification uncertainty, the harder to classify). In addition, the above statements can be extended to multi-class classification, and further elaboration is omitted here.

#### 2.1.2 Activation Function and Information Entropy

Assuming the inverse function of the activation function is  $y(x)$ , and the activation function is monotonically increasing. Many previous activation functions, such as Sigmoid and Tanh [10], satisfy the assumption that the function has an inverse function in entire definition domain. Furthermore, when an activation function fails to meet the assumption, we can transform the part of such function satisfying this assumption, as is the case with the positive part of ReLU.

Then we set data distribution before passing through the activation function obeys the distribution  $p(x)$ . Thus, data distribution after passing through activation function is :  $q(x) = p(y(x))y'(x)$ , where  $y'(x)$  represents the derivative of  $y(x)$ . Hence, we can express the information entropy as:

$$\mathbb{H}(y(x)) = - \int q(x) \log q(x) dx = - \int p(y(x))y'(x) \log(p(y(x))y'(x)) dx = \int \mathbb{G}(y'(x), y(x)) dx$$

Therefore, the information entropy can be deemed as a **functional**, which takes a function  $y(x)$  as input and produces a real number as output.

### 2.2 Worst Activation Function with Boundary Condition (WAFBC)

Firstly, we would like to determine the extremum (whether it is a maximum or minimum) of the functional  $\mathbb{H}(y(x))$ . For further deductions, taking the simplest functional into consideration, e.g. setting  $\mathbb{H}(y(x)) = \int \mathbb{G}(y'(x), y(x), x) dx$ .

In order to research the influence brought by variations of function  $y(x)$ , we apply a small perturbation  $\varepsilon\eta(x)$  to function  $y(x)$ , and then the functional  $\mathbb{H}(y(x) + \varepsilon\eta(x))$  takes the form as:

$$\mathbb{H}(y(x) + \varepsilon\eta(x)) = \int \mathbb{G}(y'(x) + \varepsilon\eta'(x), y(x) + \varepsilon\eta(x), x) dx$$

We apply **Taylor expansion** to functional  $\mathbb{H}(y(x) + \varepsilon\eta(x))$ , we can obtain the following equation:

$$\begin{aligned} \mathbb{H}(y(x) + \varepsilon\eta(x)) &= \int \left[ \mathbb{G}(y'(x), y(x), x) + \varepsilon \frac{\partial \mathbb{G}}{\partial y'} \eta'(x) + \varepsilon \frac{\partial \mathbb{G}}{\partial y} \eta(x) + \mathcal{O}(\varepsilon) \right] dx \\ &= \mathbb{H}(y(x)) + \varepsilon \int \left[ \frac{\partial \mathbb{G}}{\partial y} \eta(x) + \frac{\partial \mathbb{G}}{\partial y'} \eta'(x) \right] dx + \mathcal{O}(\varepsilon) \end{aligned} \quad (1)$$

As illustrated in Section 2.1.2,  $q(x) = p(y(x))y'(x)$  is the data distribution after passing through activation function. We can easily get that for the inverse function  $y(x)$  of activation function, when  $x$  approaches the lower bound (e.g. the initial activation function value approaches lower bound),  $y(x)$  should approaches negative infinity; and when  $x$  approaches the upper bound (e.g. the initial activation function value approaches upper bound),  $y(x)$  should approaches positive infinity. And for the sake that  $\varepsilon\eta(x)$  is a small perturbation applied to  $y(x)$ , we can draw the conclusion that  $\eta(x)$  must be 0 at the boundaries.

Utilizing the method of integration by parts and boundary condition towards Equation 1, we can derive the following results:

$$\int \frac{\partial \mathbb{G}}{\partial y'} \eta'(x) dx = \int \frac{\partial \mathbb{G}}{\partial y'} d\eta(x) = \eta(x) \frac{\partial \mathbb{G}}{\partial y'} \Big|_x - \int \eta(x) \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) dx = - \int \eta(x) \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) dx$$

Thus,  $\mathbb{H}(y(x) + \varepsilon\eta(x))$  has the following expression:

$$\mathbb{H}(y(x) + \varepsilon\eta(x)) = \mathbb{H}(y(x)) + \varepsilon \int \left[ \frac{\partial \mathbb{G}}{\partial y} - \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \right] \eta(x) dx + \mathcal{O}(\varepsilon)$$

In analogy to the extremum of ordinary functions, it is expected that the first-order term should be 0 at the extremum point. Such requirement for arbitrary  $\eta(x)$  leads to the Euler-Lagrange equation:

$$\frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) - \frac{\partial \mathbb{G}}{\partial y} = 0 \quad (2)$$

**Proposition 1.** *If  $\mathbb{G}$  is independent of  $x$ , i.e.  $\mathbb{G} = \mathbb{G}(y, y')$ , based on the Euler-Lagrange equation expressed in Equation 2, then we have:*

$$\mathbb{G} - y' \frac{\partial \mathbb{G}}{\partial y'} = C \quad (3)$$

Detailed proof of Proposition 1 can be seen in Appendix C and further discussion can be seen in Appendix D.

### 2.3 Entropy-based Activation Function Optimization (EAFO)

In Section 2.2, we have derived the extremum of the functional, showing the analytical form in Equation 10. However, the solution obtained is the global maximum, rather than the minimum. The minimum of the functional is needed if we would like to obtain the best activation function. Nonetheless, based on calculation, the actual situation is that this functional only has a global maximum but no global minimum exists. **Hence, there is no best activation function, but only better activation functions.** In this scenario, WAFBC represents a global maximum of the functional, indicating that the performance of activation functions consistently improves from WAFBC to any alternative activation functions. Therefore, we propose the following question: Is there a methodology to begin with an existing, high-performing activation function, and subsequently develop an activation function with superior performance?

Let's reconsider the Taylor expansion of the functional

$$\mathbb{H}(y(x) + \varepsilon\eta(x)) = \mathbb{H}(y(x)) + \varepsilon \int \left[ \frac{\partial \mathbb{G}}{\partial y} - \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \right] \eta(x) dx + \mathcal{O}(\varepsilon)$$

To minimize the information entropy of novel activation function, it is advisable to reduce the first-order term of Taylor expansion. In order to ensure that the information entropy of novel activation function has been indeed reduced, we would like to set  $\eta(x)$  as the opposite sign to  $\frac{\partial \mathbb{G}}{\partial y} - \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right)$ , which means we set:

$$\eta(x) = - \left( \frac{\partial \mathbb{G}}{\partial y} - \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \right) \quad (4)$$

Substitute the analytical form of functional  $\mathbb{G}(y'(x), y(x))$  into Equation 4, perform the calculation, we can derive the following equation:

$$\eta(x) = - \left( p(y(x)) \frac{y''(x)}{y'(x)} + p'(y(x)) y'(x) \right) \quad (5)$$

, where  $p(x)$  is the probability density function (PDF) of data distribution before passing through the activation function;  $p'(x)$  is the first order derivative of PDF;  $y(x)$  is inverse function of the activation function;  $y'(x)$  is the first order derivative of  $y(x)$ ;  $y''(x)$  is the second order derivative of  $y(x)$ .

As a result, we have derived a correction term that is capable of decreasing information entropy, expressing its general form in Equation 5. Subsequently, we can obtain the inverse function of the optimized activation function, denoted as  $g(x) = y(x) + \eta(x)$ . Finally, the optimized activation function can be obtained by deriving the inverse function of  $g(x)$ .

We provide further discussion on EAFO methodology in Appendix E.

## 2.4 Correction Regularized ReLU (CRReLU) : From ReLU to Better

Now, let's consider the derivation from ReLU to CRReLU. For the sake of concise representation, we rewrite the data distribution and the derivative of data distribution as:

$$p(y) = C \cdot e^{-\frac{y^2}{2}} \quad , \quad p'(y) = -C \cdot ye^{-\frac{y^2}{2}}$$

Furthermore, ReLU has a mathematical function defined as  $y = x$  when  $x$  is positive, meaning we have  $y(x) = x$ ,  $y'(x) = 1$  and  $y''(x) = 0$ . Therefore,

$$p'(y(x)) = p'(x) = -C \cdot ye^{-\frac{y^2}{2}} = -C \cdot xe^{-\frac{x^2}{2}}$$

Ultimately, by incorporating  $p'(y) = -C \cdot xe^{-\frac{x^2}{2}}$ ,  $y'(x) = 1$  and  $y''(x) = 0$  into Equation 5, we can obtain:

$$\eta(x) = -C \cdot xe^{-\frac{x^2}{2}}$$

Furthermore, we make constant  $C$  as a learnable parameter  $\varepsilon$  with the purpose of enabling self-optimization in networks. According to EAFO methodology, we can get the inverse function of revised activation function as follows:

$$g(x) = x - \varepsilon xe^{-\frac{x^2}{2}} \quad x \geq 0 \quad (6)$$

Finally, the optimized activation function CRReLU can be obtained by deriving the inverse function of  $g(x)$ . However, obtaining the inverse function of Equation 6 presents a challenge using conventional methods; as a consequence, we use the following function as a form of practical approximation.

$$f(x) = x + \varepsilon xe^{-\frac{x^2}{2}} \quad x \geq 0 \quad (7)$$

We show the rationalization and reliability of utilizing Equation 7 as the approximate inverse function of Equation 6 in Proposition 2

**Proposition 2.** *Known  $g(x) = x - \varepsilon xe^{-\frac{x^2}{2}}$ ,  $f(x) = x + \varepsilon xe^{-\frac{x^2}{2}}$ , for  $x \geq 0$ , the absolute value of error between  $g(f(x))$  and  $x$  is bounded with  $\left| e^{-1} \varepsilon^2 + 0.5e^{-\frac{3}{2}} \varepsilon^3 \right|$ .*

Detailed proof of Proposition 2 can be seen in Appendix G.

As illustrated in Section 2.2,  $\varepsilon\eta(x)$  is the small perturbation; hence, from a theoretical perspective, we can set  $\varepsilon\eta(x)$  as an infinitesimal. Furthermore, in this case, given the knowledge that  $\eta(x)$  is a bounded function, we can easily deduce that  $\varepsilon$  is also an infinitesimal. Therefore, the absolute value of error between  $g(f(x))$  and  $x$  is an infinitesimal of higher order. In practice, we typically initialize  $\varepsilon$  to a small value, such as 0.01 (as described in Section 3), implying that the absolute value of error is a small value.

Finally, let’s consider the part where  $x$  is negative. When  $x$  is negative, the inverse function of ReLU can be visualized as a ray emanating from the origin and extending to infinity, possessing an infinite slope; and when  $x$  is positive, it constitutes a ray with the slope of 1. Hence, the correction term solution for both positive and negative values of  $x$  can be considered identical, differing only by constant  $C$ . In Equation 7 and Proposition 2, it is shown that incorporating the correction term into a linear activation function can have beneficial effects by reducing the information entropy. Therefore, we can obtain the full form of Correction Regularized ReLU as:

$$f(x) = \max(0, x) + \varepsilon x e^{-\frac{x^2}{2}} \tag{8}$$

### 3 Experiments on Image Classification Task

**Datasets.** In experiments of image classification task, we adopt three datasets, ordered as CIFAR-10 [7], CIFAR-100 [7] and ImageNet-1K [8] in terms of the number of classification categories.

**Baselines.** We conduct experiments comparing the performance of CRReLU with several typical existing corrections of ReLU as illustrated in Appendix A and Appendix B : PReLU [11], ELU [12], CELU [13], GELU [14], Swish (SiLU) [15] and Mish [16].

Table 1: Test accuracy of experiments conducted on CIFAR-10 for 100 epochs.

Top-one Accuracy		GELU	ELU	PReLU	CELU	SiLU	Mish	CRReLU (ours)
CIFAR-10	ViT-Tiny	0.706	0.669	0.786	0.669	0.683	0.687	<b>0.802</b>
CIFAR-10	DeiT-Tiny	0.716	0.671	0.753	0.671	0.694	0.695	<b>0.768</b>
CIFAR-10	TNT-Small	0.743	0.689	0.761	0.689	0.719	0.725	<b>0.775</b>

Table 2: Test accuracy of experiments conducted on CIFAR-100 for 100 epochs.

Top-one Accuracy		GELU	ELU	PReLU	CELU	SiLU	Mish	CRReLU (ours)
CIFAR-100	ViT-Tiny	0.322	0.287	0.421	0.287	0.306	0.297	<b>0.459</b>
CIFAR-100	DeiT-Tiny	0.460	0.400	0.493	0.400	0.429	0.429	<b>0.508</b>
CIFAR-100	TNT-Small	0.484	0.435	0.498	0.435	0.459	0.464	<b>0.508</b>

Table 3: Test accuracy of experiments conducted on ImageNet-1K for 100 epochs.

Top-one Accuracy		GELU	ELU	PReLU	CELU	SiLU	Mish	CRReLU (ours)
ImageNet-1K	ViT-Tiny	0.542	0.384	0.572	0.384	0.469	0.479	<b>0.579</b>
ImageNet-1K	DeiT-Tiny	<b>0.619</b>	0.497	0.612	0.497	0.584	0.592	<b>0.615</b>

We further elaborate on the aforementioned experimental results in Appendix I and delineate experimental results of the LLM fine-tuning task in Appendix H. Finally, we provide a concise summary and discussion of the methodology, the study’s limitations and future work in Appendix J and the details of experimental settings in Appendix K.

## Acknowledgment

This work is partly supported by the National Natural Science Foundation of China (No.62103225), Natural Science Foundation of Shenzhen (No.JCYJ20230807111604008), Natural Science Foundation of Guangdong Province (No.2024A1515010003) and National Key Research and Development Program (No.2022YFB4701402).

## References

- [1] Richard Hahnloser, Rahul Sarpeshkar, Misha Mahowald, Rodney Douglas, and H. Seung. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405:947–51, 07 2000.
- [2] Kevin Jarrett, Koray Kavukcuoglu, Marc’Aurelio Ranzato, and Yann LeCun. What is the best multi-stage architecture for object recognition? In *2009 IEEE 12th International Conference on Computer Vision*, pages 2146–2153, 2009.
- [3] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *International Conference on Machine Learning*, 2010.
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [5] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021.
- [6] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer, 2021.
- [7] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [8] 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, 2009.
- [9] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.
- [10] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4:251–257, 1991.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, 2015.
- [12] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus), 2016.
- [13] Jonathan T. Barron. Continuously differentiable exponential linear units, 2017.
- [14] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023.
- [15] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions, 2017.
- [16] Diganta Misra. Mish: A self regularized non-monotonic activation function. In *British Machine Vision Conference*, 2020.
- [17] OpenAI. Gpt-4 technical report, 2023.

- [18] Thibaut Lavril et al. Hugo Touvron. Llama: Open and efficient foundation language models, 2023.
- [19] Gemma Team. Gemma: Open models based on gemini research and technology, 2024.
- [20] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.
- [21] Chenbin Pan, Burhaneddin Yaman, Tommaso Nesti, Abhirup Mallik, Alessandro G Allievi, Senem Velipasalar, and Liu Ren. Vlp: Vision language planning for autonomous driving, 2024.
- [22] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
- [23] H. Larochelle, Yoshua Bengio, Jérôme Louradour, and Pascal Lamblin. Exploring strategies for training deep neural networks. *J. Mach. Learn. Res.*, 10:1–40, 2009.
- [24] Andrew L. Maas. Rectifier nonlinearities improve neural network acoustic models. 2013.
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019.
- [26] Minhyeok Lee. Gelu activation function in deep learning: A comprehensive mathematical analysis and performance, 2023.
- [27] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y. Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks, 2024.
- [28] Christopher K. I. Williams. Computing with infinite networks. In *Neural Information Processing Systems*, 1996.
- [29] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes, 2018.
- [30] Daniel S. Park, Jaehoon Lee, Daiyi Peng, Yuan Cao, and Jascha Sohl-Dickstein. Towards nngp-guided neural architecture search, 2020.
- [31] Tianxiang Gao, Xiaokai Huo, Hailiang Liu, and Hongyang Gao. Wide neural networks as gaussian processes: Lessons from deep equilibrium models, 2023.
- [32] Zhongzhan Huang, Wenqi Shao, Xinjiang Wang, Liang Lin, and Ping Luo. Convolution-weight-distribution assumption: Rethinking the criteria of channel pruning, 2021.
- [33] Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. Understanding dataset difficulty with  $\mathcal{V}$ -usable information, 2022.
- [34] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022.
- [35] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [36] Samira Abnar, Mostafa Dehghani, and Willem Zuidema. Transferring inductive biases through knowledge distillation, 2020.

## A Related Work

With the development of deep learning, deep neural networks (DNNs) have gained significant prominence and achieved notable success across various domains. Recent advancements in the field of natural language processing, exemplified by large language models such as GPT-4 [17], LLama-3 [18], and Gemini [19], have propelled machine understanding and generation of natural language to unprecedented levels of accuracy. Additionally, deep neural networks have also achieved important applications in computer vision [4–6], deep reinforcement learning [20], autonomous driving[21], and many other areas.

The nonlinearity of activation functions in neural networks is crucial for both enabling the efficient learning of complex patterns, and facilitating the extraction of intricate and hierarchical representations from input data, thus allowing them to capture more complex relationships between input and output variables. In contrast, however, the nonlinear activation functions of deep neural networks also presents challenges during training, encompassing challenges like gradient vanishing [22], gradient exploding [23], and so on.

To address these challenges, researchers have explored alternative approaches for improvement, including the enhancement of activation functions. In the nascent stages of activation function development, scholars predominantly focused on rudimentary thresholding functions, initially directing their attention towards squashing functions such as the Sigmoid function and the Tanh function [10]. In order to mitigate the issues of vanishing and exploding gradients, various non-squashing functions have been proposed. Notably, ReLU [1–3] has played a pivotal role in the remarkable success of deep learning. The derivative of ReLU for positive inputs is one, thereby preventing the gradient from vanishing; however, negative values are mapped to zero, leading to two main issues: (1) The absence of information flow for negative values, known as dying ReLU ; (2) The shift in subsequent layers due to positive bias maintained by activation.

Given the aforementioned challenges, researchers have dedicated significant efforts to improving the effectiveness of activation functions. The Leaky ReLU [24] activation function permits a small negative slope, ensuring some gradient can still be propagated even when input is less than zero. The Parametric ReLU (PReLU) [11] is an extension of the Leaky ReLU, where  $\alpha$  is considered a learnable parameter that is learned from data rather than being predetermined. The Exponential Linear Unit (ELU) [12] outputs a negative value when  $x$  is less than 0, leading to the advantageous property of the average output approaching 0. The Continuously Differentiable Exponential Linear Unit (CELU) [13] proposes an alternative parameterization that simplifies analysis of the rectifier function and facilitates the tuning process of parameters in ELU. The Swish (also known as SiLU) [15] has been shown to enhance training stability and performance in deep learning models due to its smooth nature and improved gradient propagation. In Mish [16] activation function, unboundedness of positive values avoids the saturation led by a plateau, slight allowance for negative values enables better gradient flow, and the smoother activation function allows better information to flow deep into neural networks, thus resulting in better accuracy and generalization in performance.

## B Motivation

In Appendix A, it is apparent that researchers have dedicated substantial efforts to the exploration of improved activation functions, which are widely acknowledged to hold considerable significance for the advancement of deep learning. However, it has also come to our attention that proposals for these activation functions lack a theoretical framework, indicating such searches are, to some extent, inefficient and aimless.

GELU(Gaussian Error Linear Unit)[14] was first proposed in 2016 and has since gained significant success in a variety of fields, especially with the emergence of large language models in recent years. It has been successfully incorporated into several cutting-edge neural network architectures, such as BERT[25] , ViT [4] , GPT-4[17] , and so on, demonstrating its versatility and effectiveness. In the work conducted by Lee [26] (2023), insightful mathematical properties of the GELU are finally unveiled, including its differentiability, boundedness, stationarity, and smoothness. Hence, it is often the case that superior performance exhibited by novel activation functions frequently lacks mathematical explanations for their observed enhancements. Understanding may merely limited



to the fact that it exhibits improved performance, which hampers exploration for better activation functions and interpretability of neural networks.

In light of the aforementioned challenges, our work endeavors to propose a methodology for searching better activation functions, not only enabling the discovery of improved activation functions but also elucidating the reasons behind their superior performance at the same time.

## C Proof of Proposition 1

*Proof.* From Equation 2, we know that:

$$\frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) - \frac{\partial \mathbb{G}}{\partial y} = 0$$

Considering the total differential of  $\mathbb{G}$ :

$$\frac{d\mathbb{G}}{dx}(y', y, x) = \frac{\partial \mathbb{G}}{\partial x} \cdot \frac{dx}{dx} + \frac{\partial \mathbb{G}}{\partial y} \cdot \frac{dy}{dx} + \frac{\partial \mathbb{G}}{\partial y'} \cdot \frac{dy'}{dx} = \frac{\partial \mathbb{G}}{\partial x} + \frac{\partial \mathbb{G}}{\partial y} \cdot y' + \frac{\partial \mathbb{G}}{\partial y'} \cdot y''$$

Thus, we have:

$$\begin{aligned} \frac{d}{dx} \left( y' \frac{\partial \mathbb{G}}{\partial y'} \right) &= y'' \frac{\partial \mathbb{G}}{\partial y'} + y' \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \\ &= \frac{d\mathbb{G}}{dx}(y', y, x) - \frac{\partial \mathbb{G}}{\partial y} \cdot y' - \frac{\partial \mathbb{G}}{\partial x} + y' \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \\ &= \frac{d}{dx} \mathbb{G}(y', y, x) - \frac{\partial \mathbb{G}}{\partial x} - y' \cdot \left( \frac{\partial \mathbb{G}}{\partial y} - \frac{d}{dx} \left( \frac{\partial \mathbb{G}}{\partial y'} \right) \right) \\ &= \frac{d}{dx} \mathbb{G}(y', y, x) - \frac{\partial \mathbb{G}}{\partial x} \end{aligned}$$

Therefore, we know that

$$\frac{\partial \mathbb{G}}{\partial x} - \frac{d}{dx} \left( \mathbb{G} - y' \frac{\partial \mathbb{G}}{\partial y'} \right) = 0$$

For the sake that  $\mathbb{G}$  is independent of  $x$ , then we have that  $\frac{\partial \mathbb{G}}{\partial x} = 0$ . Hence,

$$\frac{d}{dx} \left( \mathbb{G} - y' \frac{\partial \mathbb{G}}{\partial y'} \right) = 0$$

Finally, we can draw the conclusion that:

$$\mathbb{G} - y' \frac{\partial \mathbb{G}}{\partial y'} = C$$

, which completes the proof. □

## D Further Discussion on Worst Activation Function with Boundary Condition (WAFBC)

Substitute  $\mathbb{G} = p(y(x))y'(x) \log(p(y(x))y'(x))$  into Equation 3 and perform the calculation, the final result is:

$$\frac{dy}{dx} p(y(x)) = C$$

Integrating both sides of the equation simultaneously, the final solution is:

$$x = c_1 \int p(y) dy + c_2 \quad (9)$$

Based on the solution we get in Equation 9, for the sake that  $y(x)$  is the inverse function of the activation function, the first integral equation can finally be solved to obtain the form of the activation function as:

$$f(x) = C_1 \int_{-\infty}^x p(t) dt + C_2 \quad (10)$$

, where  $C_1$  and  $C_2$  are two constants based on the upper bound and lower bound of activation function. Equation 10 shows the analytical form of the worst activation function with boundary condition. Through the above derivation, extremum of the functional is determined. Furthermore, we would like to deduce whether it is a maximum value or a minimum one. Applying Legendre condition to the functional extremum, then we have:

$$\mathbb{G}_{y'y'} = -\frac{p(y(x))}{y'} \leq 0$$

Therefore, the derived extremum is a maximum extremum, and is a global maximum extremum actually, meaning the deduced activation function has the worst performance. Actually, the WAFBC possesses some intriguing properties, for example, it inherently has upper and lower bounds, which can explain why bounded activation functions like Sigmoid and Tanh do not perform as well as unbounded functions like ReLU.

Let's take several typical boundary conditions into consideration. Firstly, setting  $f(x)$  approaches 1, when  $x$  tends to positive infinity; and  $f(x)$  approaches 0, when  $x$  tends to negative infinity. Therefore, the solution takes the form of cumulative distribution function (CDF), which can be expressed as:

$$f(x) = \int_{-\infty}^x p(t)dt$$

Similarly, if fixing the difference between the upper and lower bounds of the activation function to be  $e$ , and making the activation function symmetric about the origin, the form can be written as:

$$f(x) = e \int_0^x p(t)dt$$

Furthermore, in the event that the input data distribution is assumed to be approximately uniformly distributed, the worst activation function can be approximated as a linear function. Were it to approximate the input data distribution as a normal distribution, then the form of the worst activation function would be closer to Sigmoid and Tanh. We show the comparison of function curves in Figure 1 and Figure 2.

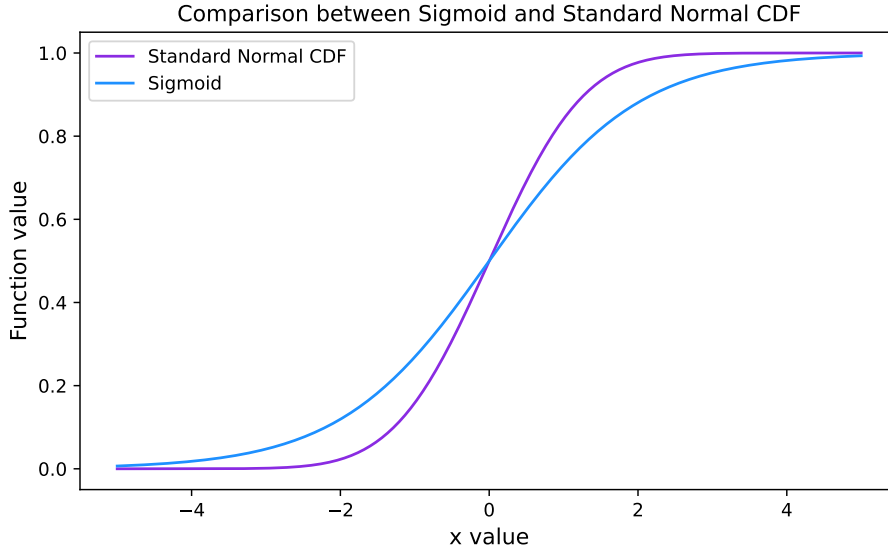


Figure 1: Comparison between Sigmoid and standard normal CDF

## E Further Discussion on EAFO methodology

**EAFO methodology outline** . In summary, we express the theoretical EAFO methodology as follows:  
 1) Utilize Equation 5 and derive correction term  $\eta(x)$  given data distribution  $p(y)$  and inverse function




Figure 2: Comparison between Tanh and Standard Normal CDF multiplied by  $e$  (has been transformed to achieve symmetry about origin)

of activation function  $y(x)$ . 2) Sum the correction term with the inverse function to obtain the inverse function of the optimized function, i.e.  $g(x) = y(x) + \eta(x)$ . 3) Derive the rigorous or approximate inverse function of  $g(x)$ , yielding the optimized activation function.

Furthermore, EAFO methodology has also shown the potential of dynamically optimizing activation during iterative training. We are acknowledged that activation of neural networks with Multi-Layer Perceptrons (MLPs) architecture is typically fixed. Recent studies, such as work done by Liu et al. [27], have suggested the optimization of activation in innovative network architectures (Kolmogorov-Arnold Networks). Furthermore, across true data distributions  $p(y)$ , utilizing EAFO methodology, we may continuously optimize activation  $y(x)$  practically under Multi-Layer Perceptrons (MLPs) architecture with numerical methods. Moreover, in theory, it is feasible to optimize activation functions using methods such as gradient descent optimization of the information entropy functional through numerical methods; however, we are also aware that this would result in an explosion of computational complexity in large neural networks, which calls for practically efficient algorithms. Hence, the EAFO methodology is still in the theoretical stage presently, providing guidance for calculating the analytical form of better activation functions.

## F Further Discussion on Correction Regularized ReLU (CRRReLU)

As illustrated in Section 2.2, it is theoretically true that the worst activation function exists, and we can determine its exact form. Actually, beginning with the worst activation function, the value of the functional  $\mathbb{G}$  consistently decreases, indicating an improvement in the performance of activation function. This reveals the feasibility of searching an improved activation function, which constitutes the crux of "optimization". In Section 2.3, EAFO is proposed as the optimization methodology. Hence, we can easily think of optimizing from WAFBC to get a better-performing activation function. While it is true that such an idea is feasible, we also observe that WAFBC itself takes the form of a variable upper bound integral, which yields a complex form of  $\eta(x)$  and renders the deduced result not practically significant. Moreover, commencing optimization from WAFBC also leads to sluggish advancement. Therefore, in practical applications, we are inclined to start from an activation function that already demonstrates relatively good performance.

Hence, we would like to take ReLU [1–3] as the beginning, and show the process of finding a better activation function. Before the deduction, we also notice that ReLU is lack of an inverse function over the entire domain. In this section, we would like to utilize following strategies for mitigating the aforementioned dilemma: the initial activation function only necessitates an inverse function

in specific regions where it is required; and when encountering parts without an inverse function, we may employ practical approximations. Therefore, we initially examine the region where  $x$  is positive in the case of ReLU. As shown in Equation 5, the derivation of correction term  $\eta(x)$  only requires original distribution  $p(y)$  and inverse function of the activation function  $y(x)$ . Knowledge of activation function is easily available, whereas original distribution remains unexplored. However, in real experiments, original distribution of experimental data would surely exhibit a substantial degree of morphological variability, thus lacking a perfect analytical form. Hence, we assume the situation is that networks are large enough, according to the Central Limit Theorem, the data processed by them can be approximated as a Gaussian distribution [28–31][32]. Certainly, such assumption may not always hold in networks of real experiments; nevertheless, approximation of the exact solution for inverse function and existence of the learnable parameter  $\varepsilon$  have significantly mitigated the impact of such assumption, which can also be demonstrated by the insensitivity of CRReLU to data distribution shown in Section 3.

**Discussion on introduced learnable parameter  $\varepsilon$ .** In Section 2.2, we have successfully demonstrated existence of the worst activation function, and from the worst as a starting point, it always moves towards improvement, regardless of the direction taken. However, commencing from a specific activation function, like ReLU here, does not invariably result in improvement across all directions, i.e. certain optimization paths may lead to deteriorated outcomes. Therefore, from the practical perspective, we introduce learnable parameter  $\varepsilon$  with the aim of enabling self-optimization of networks. From another perspective, in the derivation from ReLU to CRReLU, we assume that data follows Gaussian distribution, which might not be true in real experiments. Existence of the learnable parameter  $\varepsilon$  also weakens this assumption to some extent.

Finally, we provide further details of python-like pseudocode of CRReLU, and further discussion on properties of CRReLU.

---

**Algorithm 1:** Correction Regularized ReLU (CRReLU) Pseudocode

---

```
import torch
import torch.nn as nn
import torch.nn.functional as F

class CRReLU(nn.Module):
    def __init__(self, lr=0.01):
        super(CRReLU, self).__init__()
        self.lr = nn.Parameter(torch.tensor(lr))

    def forward(self, x):
        return F.relu(x) + self.lr * x * torch.exp(-x**2/2)
```

---

We show the function curves with different  $\varepsilon$  values for CRReLU in Figure 3. As depicted in the figure, existence of the correction term in CRReLU brings several good properties. It allows propagation of gradient when input is less than zero, serving to alleviate the dying ReLU phenomenon to a certain degree; simultaneously, as  $x$  approaches negative infinity, CRReLU also converges to 0, thereby guaranteeing sparsity of models in the negative part.

## G Proof of Proposition 2

Before the proof of Proposition 2, we would like to show four facts without proof.

**Fact 1.**  $f(x) = xe^{-\frac{x^2}{2}}$  is a bounded function, and range of the function is  $[-e^{-\frac{1}{2}}, e^{-\frac{1}{2}}]$ .

**Fact 2.**  $f(x) = x^2e^{-x^2}$  is a bounded function, and range of the function is  $[0, e^{-1}]$ .

**Fact 3.**  $f(x) = x^3e^{-\frac{3}{2}x^2}$  is a bounded function, and range of the function is  $[-e^{-\frac{3}{2}}, e^{-\frac{3}{2}}]$ .

**Fact 4.**  $\forall x \in \mathcal{R}, 1 - e^{-x} - x \leq 0$ .

We now commence the proof of Proposition 2.




Figure 3: CRRReLU with different  $\varepsilon$  value

*Proof.* Substituting the analytic expression into the formula and performing algebraic simplifications, we can obtain:

$$\begin{aligned}
 g(f(x)) &= g\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right) = x + \varepsilon x e^{-\frac{x^2}{2}} - \varepsilon \left(x + \varepsilon x e^{-\frac{x^2}{2}}\right) e^{-\frac{1}{2}\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right)^2} \\
 &= x + \varepsilon x \left(e^{-\frac{x^2}{2}} - e^{-\frac{1}{2}\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right)^2}\right) - \varepsilon^2 x e^{-\frac{x^2}{2}} e^{-\frac{1}{2}\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right)^2} \\
 &= x + \varepsilon x e^{-\frac{x^2}{2}} \left[1 - e^{-\frac{1}{2}\left(2\varepsilon x e^{-\frac{x^2}{2}} + \varepsilon^2 x^2 e^{-x^2}\right)}\right] - \varepsilon^2 x e^{-\frac{x^2}{2}} e^{-\frac{1}{2}\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right)^2}
 \end{aligned}$$

Thus,

$$\begin{aligned}
 |g(f(x)) - x| &= \left| \varepsilon x e^{-\frac{x^2}{2}} \left[1 - e^{-\frac{1}{2}\left(2\varepsilon x e^{-\frac{x^2}{2}} + \varepsilon^2 x^2 e^{-x^2}\right)}\right] - \varepsilon^2 x e^{-\frac{x^2}{2}} e^{-\frac{1}{2}\left(x + \varepsilon x e^{-\frac{x^2}{2}}\right)^2} \right| \\
 &\leq \left| \varepsilon x e^{-\frac{x^2}{2}} \left[1 - e^{-\frac{1}{2}\left(2\varepsilon x e^{-\frac{x^2}{2}} + \varepsilon^2 x^2 e^{-x^2}\right)}\right] \right| \\
 &\leq \left| \varepsilon x e^{-\frac{x^2}{2}} \left[ \frac{2\varepsilon x e^{-\frac{x^2}{2}} + \varepsilon^2 x^2 e^{-x^2}}{2} \right] \right| \\
 &= \left| \varepsilon x e^{-\frac{x^2}{2}} \left(-\varepsilon x e^{-\frac{x^2}{2}} - \frac{1}{2}\varepsilon^2 x^2 e^{-x^2}\right) \right| = \left| \varepsilon^2 x^2 e^{-x^2} + \frac{1}{2}\varepsilon^3 x^3 e^{-\frac{3}{2}x^2} \right| \\
 &\leq \left| e^{-1}\varepsilon^2 + 0.5e^{-\frac{3}{2}}\varepsilon^3 \right|
 \end{aligned}$$

The first inequality is established owing to Fact 1 and the fact that when  $x$  is positive, the second term of absolute value must be positive. The second inequality is established owing to Fact 4. The third

inequality is established owing to Fact 2 and Fact 3. Hence, we can draw the conclusion that the absolute value of error between  $g(f(x))$  and  $x$  is bounded with  $\left|e^{-1}\varepsilon^2 + 0.5e^{-\frac{3}{2}}\varepsilon^3\right|$ , which completes the proof.  $\square$

## H Experiment on the Task of Large Language Model (LLM) Fine-Tuning

In experiments of large language model (LLM) fine-tuning task, we employ two human preference datasets: SHP [33] and HH [34].

Table 4: Metrics comparison between CRReLU and GELU in the task of LLM fine-tuning.

Evaluation Metrics		Evaluation Margin Reward $\uparrow$	Evaluation Accuracy $\uparrow$	Evaluation Loss $\downarrow$
$\beta = 0.1$	CRReLU	<b>0.1428</b>	<b>0.6210</b>	<b>0.6476</b>
	GELU	0.1419	0.6196	0.6480
$\beta = 1$	CRReLU	<b>0.4626</b>	<b>0.5756</b>	<b>0.9201</b>
	GELU	0.4556	0.5731	0.9298
$\beta = 2$	CRReLU	<b>0.7736</b>	<b>0.5628</b>	<b>1.462</b>
	GELU	0.7176	0.5606	1.481
$\beta = 5$	CRReLU	<b>1.846</b>	<b>0.5635</b>	<b>3.268</b>
	GELU	1.651	0.5566	3.305

In order to further validate the effectiveness of CRReLU on larger networks and generalization to a richer range of applications, we further perform supplementary experiments on LLM fine-tuning task. We employ the Direct Preference Optimization (DPO) [9] method to fine-tune GPT-2 [35] on Stanford Human Preferences (SHP) dataset [33] and Anthropic HH dataset [34]. The parameter number of GPT-2 is 137 M, a relatively modest magnitude, hence we conduct full fine-tuning instead of LoRA-based one on  $2\times$ RTX3090. Firstly, we carry out supervised fine-tuning (SFT) with the purpose of mitigating distribution shift between the true reference distribution which is unavailable, and the reference policy utilized by DPO. Subsequently, we separately set the penalty coefficient  $\beta$  as **0.1, 1, 2, and 5**, in order to compare the performance of CRReLU and GELU under different penalty coefficients, and then execute DPO. We report evaluation metrics of fine-tuning process in Table 4, demonstrating CRReLU generally outperforms GELU in LLM fine-tuning task.

## I Experimental Result Analysis on Image Classification Task

**Experiments of ViTs on CIFAR-10 and CIFAR-100.** Vision Transformer and its variants possess sufficiently complex structure and representational capability, garnering widespread attention from the community. Moreover, the assumption of Gaussian distribution has been theoretically proved as reasonable for sufficiently large MLPs [28–31] and CNNs [32]; however, the distribution of data under attention mechanism of transformers remains unexplored. Hence, we select vision transformer and its variants as our test model in order to further investigate the insensitivity of CRReLU to data distribution. Phase of experiments on CIFAR-10 and CIFAR-100 involves the selection of Vision Transformer (ViT) [4], Data-Efficient Image Transformer (DeiT) [5] and Transformer in Transformer (TNT) [6]. We report the top-one accuracy on CIFAR-10 in Table 1 and CIFAR-100 in Table 2, demonstrating CRReLU outperforms other existing corrections of ReLU on CIFAR dataset.

**Experiments of ViTs on ImageNet-1K.** ImageNet-1K dataset poses a significant challenge to information processing capability of neural networks due to its large image size and extensive range of classification categories. Hence, phase of experiments on ImageNet-1K involves the selection of Vision Transformer (ViT) [4] and Data-Efficient Image Transformer (DeiT) [5]. We report the top-one accuracy on ImageNet-1K in Table 3. Experiments on ViT clearly demonstrate superiority of CRReLU over other activation functions, and those on DeiT, GELU shows 0.4% higher accuracy compared to CRReLU. Such result is attributed to the teacher-student strategy structure of DeiT

model. We utilize the fine-tuned "deit-tiny-patch16-224" model as teacher model, which is trained with GELU. As explained in the work [36], through distillation, transformers will inherit inductive bias. Hence, training a student model with GELU on ImageNet-1K with the help of teacher model, which has already been pre-trained on ImageNet-1K with GELU, is certain to achieve better results than other activation functions.

## J Discussion, Limitations and Future Work

Pursuit of better activation functions has been a longstanding and fundamental topic in the realm of machine learning. However, prior research has consistently concentrated on empirical search, without an emphasis on understanding the underlying mathematical mechanisms. This work aims to offer a proper solution to such issue. Our investigation into the relationship between activation functions and information theory concepts reveals that information entropy can be represented as a functional. Existence of the worst activation function with boundary condition (WAFBC) furnishes a solid theoretical basis for exploring better activation functions. In the process of solving WAFBC, we draw inspiration from the Taylor expansion form, leading us to propose Entropy-based Activation Function Optimization (EAFO) methodology. EAFO methodology presents a novel perspective for designing static activation functions in deep neural networks and shows the potential of dynamically optimizing activation during iterative training. Utilizing EAFO methodology, we derive a novel activation function from ReLU, called Correction Regularized ReLU (CRReLU). Experiments involving image classification task and large language model (LLM) fine-tuning task demonstrate that CRReLU is comparable to or surpasses existing corrections of ReLU. Overall, the EAFO methodology provides numerous promising avenues for future research on activation functions, and the CRReLU introduces a novel addition to the set of high-performing activation functions.

**Limitations and Future Work.** Our findings raise several important questions for future work. *Firstly*, how can EAFO framework be systematically generalized to non-invertible activation functions? In the initial setting of EAFO methodology, the choice of activation function is restricted to those with invertible counterparts. Despite ReLU being a prominent example of activation function without an inverse, we derive CRReLU utilizing EAFO; however, the derivation also partly benefits from the simplicity of ReLU's form and several heuristic approaches. *Secondly*, how to effectively implement activation function iteration optimization during neural network training? Notwithstanding the demonstrated feasibility of iterative activation function optimization during neural network training, it is currently hindered by the high computational complexity, particularly in large-scale neural networks. Applicability of the EAFO methodology to optimize activation in alternative network structures, such as Kolmogorov-Arnold Networks (KANs), also deserves further in-depth research. Therefore, the development of practical and efficient algorithms is an exciting direction for future work. *Finally*, while we have empirically validated the exceptional performance of CRReLU on image classification task and large language model fine-tuning task, its performance on other tasks remains to be explored, thereby warranting further investigation.

## K Details of Experimental Settings

**Experimental hyperparameters.** For all transformer-based architectures, we directly set  $\epsilon$  to 0.01 without further optimization. Detailed experimental hyperparameters are provided in this part.

We conduct all experiments on 4×RTX3090 for 100 epochs using the AdamW optimizer with weight decay of 0.05, gradient clipping norm of 1.0, cross entropy loss function, and cosine annealing learning rate scheduler with linear warm-up.

### K.1 Task of Image Classification

Table 5: Experimental settings of ViT, DeiT and TNT on CIFAR-10 and CIFAR-100 datasets

Image Size	$32 \times 32$
Patch Size	4
Embedding Dim	192 for ViT-Tiny and DeiT-Tiny ; 384 for TNT-small
Optimizer	AdamW with weight decay = 0.05
Learning Rate	Cosine Annealing Learning Rate Scheduler Initial lr = $2.5 \times 10^{-4}$ ; lr drop = -1 ; min lr = $1 \times 10^{-5}$
Warm up	warmup epochs = 20 ; warmup learning rate = $1 \times 10^{-6}$
Gradient Clipping	1.0
Training Epochs	100
Batch Size	256
Loss Function	CrossEntropy Loss
Normalization	Layer Norm
Data Augmentation	True (provided by timm)
Drop Out and Drop Path	False

Table 6: Experimental settings of ViT and DeiT on ImageNet-1K dataset

Image Size	$224 \times 224$
Patch Size	16
Embedding Dim	192
Optimizer	AdamW with weight decay = 0.05
Learning Rate	Cosine Annealing Learning Rate Scheduler Initial lr = $2.5 \times 10^{-4}$ ; lr drop = -1 ; min lr = $1 \times 10^{-5}$
Warm up	warmup epochs = 20 ; warmup learning rate = $1 \times 10^{-6}$
Gradient Clipping	1.0
Training Epochs	100
Batch Size	256
Loss Function	CrossEntropy Loss
Normalization	Layer Norm
Data Augmentation	True (provided by timm)
Drop Out and Drop Path	False



Table 7: We record changes in parameter number when employing various activation functions. GELU, ELU, CELU, SiLU (Swish), and Mish are considered activation functions without learnable parameter (**AFs without LP**), while PReLU and CRRReLU are considered activation functions with learnable parameter (**AFs with LP**). The results demonstrate that increase in parameter number introduced by the learnable parameter is negligible.

Parameter Number		CIFAR-10	CIFAR-100	ImageNet-1K
ViT-Tiny	AFs without LP	5399818	5417188	5754472
	AFs with LP	5399830	5417200	5754484
DeiT-Tiny	AFs without LP	5365076	5399816	5910800
	AFs with LP	5365088	5399828	5910812
TNT-Small	AFs without LP	21525298	21559948	/
	AFs with LP	21525322	21559972	/

## K.2 Task of Large Language Model (LLM) Fine-tuning

Table 8: Experimental settings of GPT2 fine-tuning task

Batch Size	32
Optimizer	RMSprop (More Memory-Efficient)
Learning Rate	$5 \times 10^{-7}$ with linear warmup steps of 150
Trainer	FSDPTrainer (2 GPUs)
Max Gradient Norm	10.0
Max Length for an Input (Prompt + Response)	512
Max Length for Prompt	256