

Research Article

Mathematical Analysis and Performance Evaluation of the GELU Activation Function in Deep Learning

Minhyeok Lee 

School of Electrical and Electronics Engineering, Chung-Ang University, Seoul 06974, Republic of Korea

Correspondence should be addressed to Minhyeok Lee; mlee@cau.ac.kr

Received 20 May 2023; Revised 1 July 2023; Accepted 12 July 2023; Published 10 August 2023

Academic Editor: Qiang Wu

Copyright © 2023 Minhyeok Lee. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Selecting the most suitable activation function is a critical factor in the effectiveness of deep learning models, as it influences their learning capacity, stability, and computational efficiency. In recent years, the Gaussian error linear unit (GELU) activation function has emerged as a dominant method, surpassing traditional functions such as the rectified linear unit (ReLU) in various applications. This study presents a rigorous mathematical investigation of the GELU activation function, exploring its differentiability, boundedness, stationarity, and smoothness properties in detail. In addition, we conduct an extensive experimental comparison of the GELU function against a broad range of alternative activation functions, utilizing a residual convolutional network trained on the CIFAR-10, CIFAR-100, and STL-10 datasets as the empirical testbed. Our results demonstrate the superior performance of GELU compared to other activation functions, establishing its suitability for a wide range of deep learning applications. This comprehensive study contributes to a more profound understanding of the underlying mathematical properties of GELU and provides valuable insights for practitioners aiming to select activation functions that optimally align with their specific objectives and constraints in deep learning.

1. Introduction

Deep learning has gained significant attention in recent years [1–5], leading to substantial progress in various fields such as computer vision [6–8], healthcare [9–11], and finance [12, 13]. However, the effectiveness and robustness of deep learning models are highly dependent on the choice of an appropriate activation function. The activation function plays a crucial role in introducing nonlinearities to the neural network, allowing it to capture complex patterns and relationships in the input data. Consequently, the selection of an activation function that aligns optimally with the specific task and data characteristics is a crucial consideration for practitioners.

Although several activation functions have been proposed in the literature [14, 15], each possessing unique properties and advantages, the rectified linear unit (ReLU) has emerged as the most widely used activation function due to its simplicity, efficiency, and effectiveness in various applications. However, recent studies have shown that the

ReLU function may suffer from the dying ReLU problem [16], where a large fraction of the neurons can become inactive and unresponsive, hindering the learning process. Therefore, researchers have proposed and investigated alternative activation functions that address this limitation and offer improved performance.

Amidst the plethora of activation functions that have been proposed, certain variants have attained widespread popularity due to their compelling theoretical properties and empirical success. The Gaussian error linear unit (GELU) activation function [17] is one such instance that has rapidly gained traction as a popular choice for a broad spectrum of deep learning applications. The burgeoning interest in GELU can be attributed to its desirable attributes, including its smoothness, differentiability, and ability to approximate the widely used ReLU function. The GELU activation function has been successfully integrated into several state-of-the-art neural network architectures, such as BERT [18], ViT [19], and GPT [20], demonstrating its versatility and effectiveness.

Despite the widespread adoption of GELU activation and normalization methods in deep learning, a comprehensive mathematical understanding of their combined effects on the training dynamics of deep neural networks remains an area of open investigation. In this paper, we address this gap by providing a rigorous mathematical analysis of the properties of GELU activation and normalization methods in deep learning, with a focus on their impact on the optimization process and generalization performance of deep neural networks.

In this research endeavor, we aim to unravel the intricate interactions between GELU activation and normalization techniques, investigating their impact on the optimization landscape of deep neural networks. To achieve this goal, we undertake a rigorous mathematical analysis of their combined effects, drawing on advanced mathematical formulations to elucidate their influence on the convergence and generalization performance of neural network models. Through this endeavor, we hope to offer valuable insights that empower practitioners to make informed decisions when selecting activation functions, ultimately driving more efficient and effective deep learning models.

Our analysis delves into the nuances of GELU activation's mathematical properties, including its differentiability, boundedness, stationarity, and smoothness. In addition, we undertake a comprehensive empirical comparison of the GELU function against a diverse array of alternative activation functions, employing a residual convolutional network and the CIFAR-10, CIFAR-100, and STL-10 datasets as our testbed. By evaluating the efficacy of each function on this benchmark dataset, we gain a deeper understanding of their relative strengths and weaknesses, thereby informing our insights into the broader implications of activation function selection.

2. Background

2.1. Deep Learning Models. This section presents a formal mathematical description of deep learning, focusing on the key components and operations involved in training deep neural networks. We use precise notations and rigorous mathematical expressions to represent the neural network architecture, activation functions, and learning mechanisms.

2.1.1. Neural Network Architecture. A deep neural network can be modeled as a composition of functions, representing a sequence of interconnected layers. Let \mathcal{L} denote the total number of layers in the network, with $L_1, L_2, \dots, L_{\mathcal{L}}$ representing the individual layers. Each layer L_i consists of n_i neurons, where $i \in 1, \dots, \mathcal{L}$. The weights and biases associated with layer L_i are denoted by $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$ and $b_i \in \mathbb{R}^{n_i}$, respectively.

Given an input vector $x \in \mathbb{R}^{n_0}$, the output of the network can be represented as a composition of functions:

$$f(x) = f_{\mathcal{L}}(f_{\mathcal{L}-1}(\dots f_2(f_1(x)) \dots)), \quad (1)$$

where $f_i: \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ denotes the transformation function associated with layer L_i . The transformation function f_i can be expressed as follows:

$$f_i(z) = \phi_i(W_i z + b_i), \quad (2)$$

where $\phi_i: \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ denotes the activation function at layer L_i and $z \in \mathbb{R}^{n_{i-1}}$ represents the input to layer L_i .

In essence, these equations capture how the input is transformed across the layers of a deep neural network. This is achieved by first transforming the input with a linear function and then applying a nonlinear activation function to obtain the output of the layer. The composition of these functions across multiple layers results in a complex mapping that enables the network to learn intricate data representations.

2.1.2. Activation Functions. Activation functions are pivotal in instilling nonlinearities into the network, thereby facilitating the learning of intricate patterns. Common activation functions encompass the ReLU, hyperbolic tangent (\tanh), and GELU. It is important to note that the GELU function defined below is an approximation, given by the following equation:

$$\begin{aligned} \text{ReLU}(x) &= \max(0, x), \\ \tanh(x) &= \frac{e^x - e^{-x}}{e^x + e^{-x}}, \\ \text{GELU}(x) &= 0.5x \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right). \end{aligned} \quad (3)$$

Nonlinearity enables neural networks to learn complex, hierarchical representations from the input data, enabling the network to model more sophisticated relationships between the input and output. Without nonlinearity, neural networks would simply be limited to linear transformations, severely constraining their modeling capabilities.

The introduction of nonlinearity in neural networks has enabled significant progress in a wide range of applications, including computer vision, natural language processing, and speech recognition. The ability to learn nonlinear relationships has allowed deep neural networks to achieve state-of-the-art performance on complex tasks such as image classification, object detection, and language translation.

However, nonlinearity can also introduce challenges in the training of deep neural networks, such as the vanishing gradient problem and the exploding gradient problem. Nonlinear activation functions can lead to the amplification or attenuation of gradients during backpropagation, making it difficult to update the weights and biases of the neural network. Consequently, a careful selection of activation functions is critical to ensure stable and efficient training of deep neural networks.

2.1.3. Loss Function and Optimization. To optimize the neural network, a loss function $\mathcal{L}(y, \hat{y})$ that measures the discrepancy between the predicted output $\hat{y} = f(x)$ and the true output y is required. The choice of loss function is dependent on the task at hand. For regression tasks,

common loss functions include the mean squared error (MSE), mean absolute error (MAE), and Huber loss:

$$\begin{aligned} \text{MSE}(y, \hat{y}) &= \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, \\ \text{MAE}(y, \hat{y}) &= \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \\ \text{Huber}(y, \hat{y}) &= \frac{1}{n} \sum_{i=1}^n \begin{cases} \frac{1}{2}(y_i - \hat{y}_i)^2, & \text{if } |y_i - \hat{y}_i| \leq \delta, \\ \delta \left(|y_i - \hat{y}_i| - \frac{1}{2}\delta \right), & \text{otherwise,} \end{cases} \end{aligned} \quad (4)$$

where y_i is the true output for the i -th sample and \hat{y}_i is the predicted output.

For classification tasks, the cross-entropy loss is commonly used. Other loss functions include the hinge loss for support vector machines (SVMs) and the triplet loss for metric learning:

$$\begin{aligned} \text{CrossEntropy}(y, \hat{y}) &= - \sum_{i=1}^n y_i \log(\hat{y}_i), \\ \text{Hinge}(y, \hat{y}) &= \max(0, 1 - y\hat{y}), \\ \text{Triplet}(a, p, n) &= \max(0, d(a, p) - d(a, n) + \alpha), \end{aligned} \quad (5)$$

where y_i is the true class label for the i -th sample, \hat{y}_i is the predicted probability of the i -th sample belonging to the true class, a , p , and n represent the anchor, positive, and negative samples, respectively, and d denotes the distance metric used for embedding the samples.

The choice of the loss function is a critical factor in deep learning, as it acts as the objective function that the optimizer endeavors to minimize during the training process. The selection of an appropriate loss function depends on various factors such as the problem's nature, the output type, and the performance metric of interest. The optimization process strives to reduce the loss function by modifying the neural network parameters via optimization algorithms such as SGD and Adam.

It is crucial to choose a suitable loss function that is tailored to the problem, as selecting an unsuitable one can negatively impact the neural network's learning dynamics. Ineffective loss functions can lead to inadequate convergence, underfitting, or overfitting. On the other hand, selecting an appropriate loss function can expedite convergence rates, enhance generalization performance, and mitigate the risk of overfitting.

Researchers have proposed various modifications and extensions of the standard loss functions to address specific scenarios [21]. For instance, focal loss [22] bestows higher weights to hard-to-classify instances and has shown to be effective in imbalanced classification problems. On the other hand, adversarial loss [23] seeks to enhance the neural network's resilience to adversarial attacks and has been employed in security-critical applications such as image and

text classification. The selection of an appropriate loss function hinges on the problem's characteristics and the task objectives. To ensure the neural network's optimal performance, it is crucial to carefully evaluate and consider different loss functions.

To minimize the loss function, we utilize optimization algorithms [24] that update the weights and biases of the network iteratively. The most common optimization algorithm is gradient descent, which updates the parameters $\theta = \{W_i, b_i\}_{i=1}^{\mathcal{L}}$ by following the negative gradient of the loss function with respect to the parameters:

$$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}(y, \hat{y}), \quad (6)$$

where $\eta > 0$ is the learning rate and $\nabla_{\theta} \mathcal{L}(y, \hat{y})$ represents the gradient of the loss function with respect to the parameters.

The gradient of the loss function can be computed using the backpropagation algorithm, which applies the chain rule of calculus to compute the gradients in a layer-wise manner, starting from the output layer and propagating backward through the network. The gradients with respect to the weights and biases of layer L_i are given by the following equation:

$$\begin{aligned} \frac{\partial \mathcal{L}(y, \hat{y})}{\partial W_i} &= \delta_i z_{i-1}^{\top}, \\ \frac{\partial \mathcal{L}(y, \hat{y})}{\partial b_i} &= \delta_i, \end{aligned} \quad (7)$$

where $\delta_i \in \mathbb{R}^{n_i}$ denotes the error at layer L_i and $z_{i-1} \in \mathbb{R}^{n_{i-1}}$ represents the input to layer L_i . The error term δ_i can be computed recursively using the error term of the subsequent layer δ_{i+1} :

$$\delta_i = (\delta_{i+1} W_{i+1}) \odot \phi'_i(z_i), \quad (8)$$

where \odot denotes the element-wise product and $\phi'_i(z_i)$ represents the element-wise derivative of the activation function at layer L_i with respect to its input z_i .

The Adam optimizer [25] is a sophisticated and widely used optimization algorithm in deep learning that combines the advantages of adaptive learning rates with the momentum method. It has been demonstrated to be effective in training deep neural networks due to its ability to adapt the learning rate for each parameter individually, leading to faster convergence and improved generalization performance.

The Adam optimizer operates by maintaining an exponential moving average of the first and second moments of the gradients. Let g_t denote the gradient of the loss function $\mathcal{L}(y, \hat{y})$ with respect to the parameters θ at iteration t . The first moment, m_t and the second moment, v_t , are updated as follows:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t, \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) g_t^2, \end{aligned} \quad (9)$$

where β_1 and β_2 are the exponential decay rates for the first and second moments, respectively, and g_t^2 denotes the element-wise square of the gradient. Typically, the values of β_1 and β_2 are set to 0.9 and 0.999, respectively.

The first and second moments are initialized to zero, which can result in biased estimates during the initial iterations. To mitigate this, Adam employs bias correction to obtain unbiased estimates of the first and second moments, denoted as \widehat{m}_t and \widehat{v}_t :

$$\begin{aligned}\widehat{m}_t &= \frac{m_t}{1 - \beta_1^t}, \\ \widehat{v}_t &= \frac{v_t}{1 - \beta_2^t},\end{aligned}\quad (10)$$

where t represents the current iteration.

With the unbiased estimates of the first and second moments, the Adam optimizer updates the parameters θ as follows:

$$\theta \leftarrow \theta - \eta \frac{\widehat{m}_t}{\sqrt{\widehat{v}_t + \epsilon}}, \quad (11)$$

where η is the learning rate and ϵ is a small constant added to prevent division by zero, typically set to 10^{-8} .

2.2. GELU Activation Function. The GELU activation function, introduced by [17], is a smooth and differentiable approximation of the rectifier function. It has gained popularity in deep learning due to its desirable properties, such as nonlinearity, differentiability, and smoothness. As a result, GELU has been employed in various state-of-the-art architectures [26–29], including BERT [18], ViT [19], and GPT [20].

2.2.1. Motivation. The impetus for the development of the GELU activation function is to offer a smooth and differentiable alternative to the widely used ReLU activation function, without compromising its inherent benefits. The ReLU function, denoted as $\text{ReLU}(x) = \max(0, x)$, imparts nonlinearities to the network; however, it is non-differentiable at $x = 0$. This nondifferentiability can result in complications during gradient-based optimization, such as dead neurons or erratic training dynamics.

In order to mitigate these concerns, the GELU activation function is devised as a smooth approximation to the ReLU function, ensuring differentiability at every point while preserving the requisite nonlinear properties for deep learning applications. The GELU function draws inspiration from the Gaussian cumulative distribution function (CDF), which is characterized by its inherent smoothness and differentiability properties.

2.2.2. Derivation of the GELU Function. The GELU activation function is primarily derived from the Gaussian CDF, which can be defined as follows:

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt, \quad (12)$$

where $\Phi(x)$ signifies the likelihood of a random variable with a standard normal distribution taking a value less than

or equal to x . The GELU function can be expressed as a product of the input x and the Gaussian CDF:

$$\text{GELU}(x) = x \cdot \Phi(\alpha x), \quad (13)$$

where $\alpha > 0$ acts as a scaling factor, modulating the smoothness of the GELU function. Typically, $\alpha = 1$.

To further simplify the GELU function and enhance computational efficiency, an approximation of the Gaussian CDF is commonly used in practice:

$$\Phi(\alpha x) \approx \frac{1}{2} \left(1 + \tanh(\beta(\alpha x + \gamma(\alpha x)^3)) \right), \quad (14)$$

where $\beta > 0$ and $\gamma \in \mathbb{R}$ are constants, selected to minimize approximation error. Substituting this approximation into the GELU function, we arrive at the final approximate form of the GELU activation function (Figure 1):

$$\text{GELU}(x) = 0.5x \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right). \quad (15)$$

This version of the GELU function is smooth, differentiable, and computationally efficient, rendering it suitable for deployment in deep learning architectures.

2.3. Normalization Methods. Normalization methods aim to mitigate the internal covariate shift in deep neural networks by normalizing the inputs at each layer. These methods result in more stable training dynamics and allow for faster convergence by reducing the dependence of gradients on the input distribution. Normalization methods have become an essential component of modern deep learning architectures, as they enable training deeper networks with larger learning rates.

2.3.1. Batch Normalization. Batch normalization (BN) [30] is a widely used normalization technique that reduces internal covariate shift by normalizing activations across a minibatch during training. Given a minibatch $\mathcal{B} = x_1, x_2, \dots, x_m$ of m input activations at a particular layer, BN computes the mean $\mu_{\mathcal{B}}$ and variance $\sigma_{\mathcal{B}}^2$ of the minibatch as follows:

$$\mu_{\mathcal{B}} = \frac{1}{m} \sum_{i=1}^m x_i, \quad (16)$$

$$\sigma_{\mathcal{B}}^2 = \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2.$$

The input activations are then normalized using the computed mean and variance:

$$\widehat{x}_i = \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}}, \quad (17)$$

where $\epsilon > 0$ is a small constant added for numerical stability. Finally, BN applies a learned affine transformation to the normalized activations:

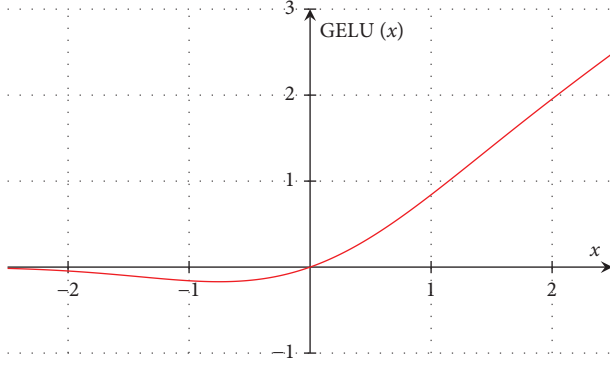


FIGURE 1: The GELU function defined in equation (15).

$$y_i = \gamma \hat{x}_i + \beta, \quad (18)$$

where γ and β are learnable parameters of the same shape as the input activations, allowing the model to learn the appropriate scale and shift for the normalized activations.

2.3.2. Layer Normalization. Layer normalization (LN) [31] is another normalization technique that addresses some of the limitations of BN, such as the dependence on minibatch size and reduced performance in recurrent networks. Unlike BN, which normalizes activations across a minibatch, LN normalizes activations across the feature dimension at each layer.

Given an input activation $x \in \mathbb{R}^d$ at a particular layer, LN computes the mean μ_x and variance σ_x^2 of the input activation as follows:

$$\begin{aligned} \mu_x &= \frac{1}{d} \sum_{i=1}^d x_i, \\ \sigma_x^2 &= \frac{1}{d} \sum_{i=1}^d (x_i - \mu_x)^2. \end{aligned} \quad (19)$$

Similar to BN, LN normalizes the input activations using the computed mean and variance and applies a learned affine transformation:

$$\begin{aligned} \hat{x}_i &= \frac{x_i - \mu_x}{\sqrt{\sigma_x^2 + \epsilon}}, \\ y_i &= \gamma \hat{x}_i + \beta, \end{aligned} \quad (20)$$

where $\epsilon > 0$ is a small constant added for numerical stability, and γ and β are learnable parameters of the same shape as the input activations.

2.3.3. Group Normalization. Group normalization (GN) [32] is a normalization technique that generalizes BN and LN by dividing the feature channels into groups and normalizing within each group. GN addresses some of the limitations of BN and LN, such as reduced performance in small minibatches and the need to choose between normalizing across the batch or feature dimensions.

Given an input activation $x \in \mathbb{R}^{C \times H \times W}$, where C is the number of channels and H and W are the spatial dimensions, GN divides the channels into G groups, with each group containing C/G channels. For each group $g \in 1, \dots, G$, GN computes the mean μ_g and variance σ_g^2 of the input activations within the group as follows:

$$\begin{aligned} \mu_g &= \frac{1}{C/G \cdot H \cdot W} \sum_{i=1}^{C/G} \sum_{j=1}^H \sum_{k=1}^W x_{g,i,j,k}, \\ \sigma_g^2 &= \frac{1}{C/G \cdot H \cdot W} \sum_{i=1}^{C/G} \sum_{j=1}^H \sum_{k=1}^W (x_{g,i,j,k} - \mu_g)^2, \end{aligned} \quad (21)$$

where $x_{g,i,j,k}$ denotes the activation value of the i -th channel in the g -th group at spatial location (j, k) .

GN then normalizes the input activations within each group using the computed mean and variance and applies a learned affine transformation:

$$\hat{x}_{g,i,j,k} = \frac{x_{g,i,j,k} - \mu_g}{\sqrt{\sigma_g^2 + \epsilon}}, \quad (22)$$

$$y_{g,i,j,k} = \gamma_g \hat{x}_{g,i,j,k} + \beta_g,$$

where $\epsilon > 0$ is a small constant added for numerical stability and γ_g and β_g are learnable parameters of the same shape as the input activations within the group.

3. Comprehensive Mathematical Analysis

We delve into a thorough mathematical examination of the GELU activation function and normalization methods, concentrating on their differentiability, boundness, stationarity, and smoothness properties.

3.1. Differentiability. Here, we offer a mathematical exploration of the differentiability of the GELU activation function. The differentiability of an activation function holds paramount importance for gradient-based optimization algorithms, as it guarantees the existence and computability of the gradients essential for backpropagation.

3.1.1. Derivative of the GELU Function. Now, we determine the derivative of the GELU activation function concerning its input x . The differentiability of the GELU function plays a crucial role in gradient-based optimization algorithms, as it ensures the existence and computability of the gradients necessary for backpropagation.

As shown in the previous sections, the GELU function can be represented in terms of the Gaussian CDF as given in equation (13). To compute the derivative of the GELU function with respect to its input x , we apply the chain rule of calculus:

$$\frac{d\text{GELU}(x)}{dx} = \frac{d(x \cdot \Phi(\alpha x))}{dx} = x \cdot \frac{d\Phi(\alpha x)}{dx} + \Phi(\alpha x). \quad (23)$$

Now, we need to compute the derivative of the Gaussian CDF with respect to its argument, $d\Phi(\alpha x)/dx$. Since $\Phi(x)$ is the integral of the Gaussian probability density function (PDF), we can differentiate the Gaussian CDF to obtain the Gaussian PDF scaled by the factor α :

$$\frac{d\Phi(\alpha x)}{dx} = \frac{\alpha}{\sqrt{2\pi}} e^{-(\alpha x)^2/2}. \quad (24)$$

Substituting this result back into the expression for the derivative of the GELU function, we obtain the following equation:

$$\frac{d\text{GELU}(x)}{dx} = \frac{\alpha x}{\sqrt{2\pi}} e^{-(\alpha x)^2/2} + \Phi(\alpha x). \quad (25)$$

As we utilize an approximation of the Gaussian CDF in equation (15), this approximation enables us to compute an approximate derivative of the GELU function. By employing the chain rule and the product rule, we can calculate the derivative of the GELU function with respect to x in the following manner:

$$\begin{aligned} \frac{d\text{GELU}(x)}{dx} &= 0.5 \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right) \\ &\quad + 0.5x \cdot \frac{d \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right)}{dx} \\ &= 0.5 \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right) \\ &\quad + 0.5x \cdot \text{sech}^2 \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \\ &\quad \cdot \frac{d \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right)}{dx}. \end{aligned} \quad (26)$$

Now, we compute the derivative of $\sqrt{\frac{2}{\pi}}(x + 0.044715x^3)$ with respect to x :

$$\frac{d \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right)}{dx} = \sqrt{\frac{2}{\pi}} (1 + 3 \cdot 0.044715x^2). \quad (27)$$

Substituting this expression back into the derivative of the GELU function, we obtain

$$\begin{aligned} \frac{d\text{GELU}(x)}{dx} &= 0.5 \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right) \\ &\quad + 0.5x \cdot \text{sech}^2 \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \\ &\quad \cdot \sqrt{\frac{2}{\pi}} (1 + 3 \cdot 0.044715x^2). \end{aligned} \quad (28)$$

Figure 2 shows the GELU function and its derivative with respect to x . As can be seen, the GELU function is differentiable at all points in its domain, ensuring the existence and computability of the gradients required for gradient-based optimization. Also, note that derivatives can be negative in some points. In deep learning, the optimization process seeks to minimize the loss function $\mathcal{L}(y, \hat{y})$

by updating the parameters of the model θ . This is carried out by iteratively adjusting the parameters in the negative direction of the gradient, as shown in equation (6). Now, let us analyze the implications of negative derivative values in the context of deep learning training. If we denote f_l as the output of the l -th layer before the activation, then the output after applying GELU activation is $g_l = \text{GELU}(f_l)$. During

backpropagation, we calculate the gradient of the loss function with respect to the output of each layer:

$$\delta_l = \frac{\partial \mathcal{L}(y, \hat{y})}{\partial g_l}. \quad (29)$$

Using the chain rule, we can calculate the gradient of the loss function with respect to the input of the activation function:

$$\frac{\partial \mathcal{L}(y, \hat{y})}{\partial f_l} = \delta_l \frac{d \text{GELU}(f_l)}{df_l}. \quad (30)$$

Negative derivative values of the GELU activation function indicate that the gradient of the activation function is negative, i.e., $d \text{GELU}(f_l)/df_l < 0$. This implies that a small positive increment in f_l would lead to a decrease in the value of g_l . In this case, the gradient update for the parameters would be:

$$\theta \leftarrow \theta - \eta \delta_l \frac{\partial f_l}{\partial \theta} \frac{d \text{GELU}(f_l)}{df_l}. \quad (31)$$

Since $d \text{GELU}(f_l)/df_l < 0$, the update step becomes:

$$\theta \leftarrow \theta + \eta \delta_l \frac{\partial f_l}{\partial \theta} \left(-\frac{d \text{GELU}(f_l)}{df_l} \right). \quad (32)$$

This implies that the model parameters will be updated in the direction opposite to the gradient of f_l concerning θ . This update will have the effect of augmenting the value of the loss function, as θ will be modified to attain a higher value of the loss function. Notably, the negative derivative values of the GELU activation function are minimal, which can be advantageous for deep learning training, particularly during the initial stages. At the onset of the training process, the values of x are generally zero-centered due to the initialization of weights and biases. In this scenario, negative derivative values can aid the model in escaping local minima in the earlier stage, resulting in more effective optimization.

For minuscule values of x , the negative derivative values remain small, and the update step in the optimization




FIGURE 2: The GELU function and its derivative with respect to x .

process assists the model in evading local minima. As the training progresses, the variance of the values of x enlarges, and negative values of x produce gradients close to zero. This leads to stable training, as the update step in the optimization process diminishes, preventing substantial alterations in the model parameters.

3.2. Boundness. In the present investigation, an analysis of the boundness property of the GELU activation function is conducted. Activation functions that exhibit boundedness are known to aid in circumventing the issue of vanishing or exploding gradients, which may arise during the training process by constraining the activations within a predetermined range.

3.2.1. Boundness of the GELU Function. To analyze the boundness of the GELU activation function, we examine the limits of the function as the input x approaches positive or negative infinity:

$$\begin{aligned} \lim_{x \rightarrow -\infty} \text{GELU}(x) &= \lim_{x \rightarrow -\infty} 0.5x \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right) = 0, \\ \lim_{x \rightarrow \infty} \text{GELU}(x) &= \lim_{x \rightarrow \infty} 0.5x \left(1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.044715x^3) \right) \right) = \infty. \end{aligned} \quad (33)$$

To ascertain the minimum value of the GELU activation function, we can study the first derivative of $\text{GELU}(x)$ with respect to x , identifying critical points where the derivative equates to zero. Solving $d \text{GELU}(x)/dx = 0$ delivers the minimum value of $\text{GELU}(x)$, approximately -0.17 , occurring at $x \approx -0.75$.

By taking into account these limits and conducting an analysis of the critical points, it can be inferred that the

GELU activation function has a lower bound of approximately -0.17 and is unbounded in the positive direction. This property, combined with the insights discussed in the following section regarding the practical upper bound, ensures that the activations are confined within a specific range during training. Consequently, this characteristic assists in mitigating the challenge of the vanishing or exploding gradient problem.

3.2.2. Upper-Boundness of the GELU Function. The present work furnishes a mathematical proof that establishes the upper-bound property of the combination of normalization methods and the GELU activation function. In particular, the study focuses on a layer constituted by a linear transformation operation followed by a normalization method, and finally, a GELU activation function.

Let L_i denote a layer in the neural network with input $z_i \in \mathbb{R}^{n_{i-1}}$, weights $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$, and biases $b_i \in \mathbb{R}^{n_i}$. We first apply the linear transformation to the input:

$$z'_i = W_i z_i + b_i. \quad (34)$$

Next, we apply a normalization method to the transformed input z'_i . For simplicity, we will use generic normalization denoted by the function $\mathcal{N}(z'_i)$. The normalized input is then

$$z''_i = \mathcal{N}(z'_i). \quad (35)$$

We apply the GELU activation function to the normalized input:

$$z_{i+1} = \text{GELU}(z''_i). \quad (36)$$

To show that the combination of normalization methods and GELU is upper-bounded, we need to find an upper-bound $M > 0$ such that for any input z_i , we have $z_{i+1} \leq M$.

Since the normalization method is applied before the GELU activation function, we know that z''_i has a fixed range, typically with mean 0 and variance 1. As a result, there exists a constant $K > 0$ such that $|z''_i|_\infty \leq K$.

The GELU function is upper-bounded by x :

$$\text{GELU}(x) \leq x. \quad (37)$$

Thus, we have

$$\begin{aligned} z_{i+1} = \text{GELU}(z''_i) &\leq z''_i \\ &\leq |z''_i|_\infty \\ &\leq K. \end{aligned} \quad (38)$$

Hence, we have shown that the combination of normalization methods and the GELU activation function is upper-bounded, with the upper bound being $M = K$. This property ensures that the activations remain within a fixed range during training, further helping to mitigate the vanishing or exploding gradient problem.

However, without normalization, the transformed input z'_i in equation (34) can become larger as the learning progresses, which can lead to larger values of x in $\text{GELU}(x) \leq x$.

As learning progresses, the weights W_i and biases b_i are updated, and their magnitudes may increase. Consequently, the magnitudes of z'_i can also increase, resulting in larger values of x :

$$x = z'_i \implies |x| \leq |W_i| |z_i| + |b_i|, \quad (39)$$

where $|\cdot|$ denotes an appropriate norm, e.g., the Euclidean norm. As the magnitudes of W_i and b_i grow, the bound on $|x|$ can also grow.

When there are multiple layers, this effect can deepen. Consider a deep neural network with N layers and let z'_j denote the transformed input at layer j , $j \in 1, 2, \dots, N$. Without normalization, the transformed input at layer j can be expressed as follows:

$$z'_j = W_j z_{j-1} + b_j. \quad (40)$$

For $j = 1, 2, \dots, N$, the bound on the magnitudes of z'_j can be recursively computed as follows:

$$|z'_j| \leq |W_j| |z'_{j-1}| + |b_j|. \quad (41)$$

As learning progresses and the magnitudes of W_j and b_j increase, the bound on $|z'_j|$ can grow, leading to larger values of x in $\text{GELU}(x) \leq x$ at each layer. This growth in the bound of $|z'_j|$ can compound across multiple layers, potentially leading to undesirable effects such as the vanishing or exploding gradient problem.

3.3. Stationarity. The present investigation concerns an analysis of the stationarity of the GELU activation function, with particular emphasis on its continuity, differentiability, and Lipschitz continuity properties. The stationarity of an activation function is of utmost significance as it aids in maintaining a well-behaved optimization landscape, which, in turn, facilitates more efficient convergence during the training process.

3.3.1. Continuity and Differentiability. The GELU activation function, as defined in equation (15), is a continuous function for all $x \in \mathbb{R}$. Since the composition of continuous functions is also continuous, we observe that the GELU function is continuous, given that both the scalar multiplication, addition, and the hyperbolic tangent function are continuous. Furthermore, the GELU function is differentiable everywhere, as shown in Section 3.1.

3.3.2. Lipschitz Continuity. Lipschitz continuity is a stronger form of continuity that provides an upper bound on the rate of change of a function [33]. A function is said to be Lipschitz continuous if there exists a constant $L > 0$ such that for all $x, y \in \mathbb{R}$, the following inequality holds:

$$|\text{GELU}(x) - \text{GELU}(y)| \leq L|x - y|. \quad (42)$$

For the GELU activation function to exhibit Lipschitz continuity, we need to show that its derivative is bounded. As per equation (25), the derivative of the GELU function with $\alpha = 1$ is given by the following equation:

$$\frac{d\text{GELU}(x)}{dx} = x \cdot \frac{1}{\sqrt{2\pi}} e^{-x^2/2} + \Phi(x). \quad (43)$$

We aim to find a constant $L > 0$ such that for all $x \in \mathbb{R}$, we have $|d\text{GELU}(x)/dx| \leq L$.

We use the second derivative to establish a tight bound on the derivative. The second derivative of the GELU function is given by the following equation:

$$\begin{aligned} \frac{d^2 \text{GELU}(x)}{dx^2} &= \frac{d}{dx} \left(x \cdot \frac{1}{\sqrt{2\pi}} e^{-x^2/2} + \Phi(x) \right), \\ &= \frac{1}{\sqrt{2\pi}} e^{-x^2/2} (1 - x^2) + \frac{1}{\sqrt{2\pi}} e^{-x^2/2}, \quad (44) \\ &= \frac{1}{\sqrt{2\pi}} e^{-x^2/2} (2 - x^2). \end{aligned}$$

Setting the second derivative equal to zero and solving for x , we obtain two critical points at $x = -\sqrt{2}$ and $x = \sqrt{2}$. Evaluating the first derivative at $x = \sqrt{2}$, we find:

$$\left. \frac{d \text{GELU}(x)}{dx} \right|_{x=\sqrt{2}} = \frac{e}{\sqrt{2\pi}} \approx 1.084. \quad (45)$$

Thus, we have found that the absolute value of the derivative of the GELU function is bounded by $L = e/\sqrt{2\pi} \approx 1.084$, thereby proving its Lipschitz continuity.

3.4. Smoothness of Feature Space. The present study undertakes a thorough and rigorous investigation of the smoothness of the feature space that is induced by the GELU activation function. The property of smoothness is highly desirable for activation functions, as it plays a crucial role in achieving well-conditioned optimization landscapes, thereby facilitating more efficient convergence during the training phase. The smoothness of the GELU function is examined in this work by scrutinizing its higher-order derivatives and employing the concept of Holder continuity.

3.4.1. Higher-Order Derivatives. To analyze the smoothness of the GELU activation function, we first examine its higher-order derivatives. Higher-order derivatives provide insights into the local curvature and smoothness of a function. The first derivative of the GELU function with respect to its input x is given by equation (43). To compute the second derivative, we differentiate the first derivative with respect to x , as in equation (44).

The second derivative gives information about the concavity of the GELU function. Since the second derivative is continuous, it implies that the GELU function is twice differentiable, and therefore, smooth.

3.4.2. Holder Continuity. Another measure of smoothness is the Holder continuity of a function [34]. A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is said to be Holder continuous with exponent $\alpha \in [0, 1]$ if there exists a constant $C > 0$ such that for all $x, y \in \mathbb{R}$, the following inequality holds:

$$|f(x) - f(y)| \leq C|x - y|^\alpha. \quad (46)$$

The larger the Holder exponent α , the smoother the function. If $\alpha = 1$, the function is Lipschitz continuous, and if $1 < \alpha \leq 2$, the function is twice continuously differentiable.

We have already shown in Section 3.3 that the GELU activation function is Lipschitz continuous, which implies that it is also Holder continuous with exponent $\alpha = 1$. Furthermore, the existence of the second derivative, as shown in the previous section, implies that the GELU function is Holder continuous with an exponent $\alpha \in [1, 2]$. This result demonstrates the smoothness of the feature space induced by the GELU activation function.

4. Experimental Comparison

In this section, we present a comprehensive experimental comparison of various activation functions within the context of residual convolutional networks trained on the CIFAR-10, CIFAR-100, and STL-10 datasets. The objective is to investigate the impact of diverse activation functions and compare the activation functions empirically.

The activation functions under scrutiny encompass a wide range of popular and effective choices, including ELU, Hardshrink, Hardsigmoid, Hardtanh, Hardswish, LeakyReLU, LogSigmoid, PReLU, ReLU, ReLU6, RReLU, SELU, CELU, GELU, Sigmoid, Softplus, Softshrink, Softsign, Tanh, and Tanhshrink [14, 15]. Several activation functions are displayed in Figure 3. For each activation function, the same training procedure was followed, employing cross-entropy loss as the criterion and the Adam optimizer for parameter updates. We trained the residual network for 20 epochs with a batch size of 128 and a learning rate of 0.001.

The residual network operates by ingesting a tensor of 3-channel images and feeding it through a sequence of convolutional layers. Each of these layers incorporates both BN and nonlinear activation. The residual blocks, which form the building blocks of the network, are composed of two convolutional layers and a skip connection. We used a preactivated residual network, where each block consists of two layers of BN, nonlinear activation, and a convolutional operation in order.

The architecture of the network includes an initial convolutional layer that enlarges the dimension, followed by six residual blocks, an adaptive pooling layer, and a fully connected layer for classification. Consequently, the network consists of 14 layers in total. The stride in the third and fifth residual blocks is set to 2, which effectively reduces the spatial dimensions of the feature maps.

An in-depth analysis of the results on the CIFAR-10 dataset presented in Table 1 and Figure 4 reveals intriguing patterns and trends among the activation functions. The test loss and test accuracy, which serve as the primary evaluation metrics, provide valuable insights into the efficacy of each activation function in the context of the residual convolutional network.

Several activation functions exhibit commendable performance, with GELU standing out as the top-performing function, achieving the lowest test loss of 0.3685 and the highest test accuracy of 89.52%. Hardswish and ReLU6 follow closely behind, registering test accuracies of 88.77% and 88.70%, respectively. These results suggest that GELU, Hardswish, and ReLU6 may be more suitable for this




FIGURE 3: Comparison of various activation functions.

particular network architecture and dataset, delivering superior performance in comparison to other activation functions.

Conversely, Sigmoid emerges as the least effective activation function, with a test loss of 3.2102 and a markedly low test accuracy of 33.90%. This result underlines the limitations of the Sigmoid function, which may suffer from issues such as vanishing gradients, particularly in deeper networks. The relatively poor performance of Sigmoid highlights the importance of selecting appropriate activation functions for the task at hand.

Other activation functions, such as ELU, LeakyReLU, and PReLU, exhibit satisfactory performance, with test accuracies ranging between 85% and 87%. These functions demonstrate their potential utility in deep learning applications, though they may not be the optimal choices for this specific network and dataset.

The disparities in performance among the activation functions can be attributed to various factors, including the nature of the dataset, the architecture of the network, and the inherent properties of the activation functions themselves. These results emphasize the significance of conducting empirical comparisons to identify the most suitable activation functions for a given deep learning problem.

In order to further substantiate the superior performance of the GELU activation function, we conducted additional experiments on two benchmark datasets, CIFAR-100 and STL-10, which are known for their complexity and diversity. In these additional experiments, we

TABLE 1: Test loss and test accuracy for different activation functions on CIFAR-10 dataset.

| Activations | Test loss | Test accuracy (%) |
|-------------|---------------|-------------------|
| ELU | 0.4232 | 86.22 |
| Hardshrink | 1.1266 | 60.81 |
| Hardsigmoid | 1.4296 | 54.00 |
| Hardtanh | 0.5573 | 82.01 |
| Hardswish | 0.3921 | 88.77 |
| LeakyReLU | 0.4036 | 87.93 |
| LogSigmoid | 0.5755 | 81.42 |
| PReLU | 0.5552 | 86.33 |
| ReLU | 0.4478 | 87.19 |
| ReLU6 | 0.4145 | 88.70 |
| RReLU | 0.4308 | 85.91 |
| SELU | 0.4983 | 83.37 |
| CELU | 0.4260 | 86.21 |
| Sigmoid | 3.2102 | 33.90 |
| Softplus | 0.5762 | 80.82 |
| Softshrink | 0.5626 | 81.93 |
| Softsign | 0.6819 | 78.33 |
| Tanh | 0.5318 | 82.91 |
| Tanhshrink | 0.5776 | 80.78 |
| GELU | 0.3685 | 89.52 |

*Bold indicates the best performance; italic indicates the second-best.

selected several activation functions that have shown promising results on the CIFAR-10 dataset. Table 2 presents the test loss and test accuracy for different activation functions, including GELU, on these two datasets. The results demonstrate that GELU consistently outperforms its

TABLE 2: Test loss and test accuracy for selected activation functions on CIFAR-100 and STL-10 datasets.

| Datasets | Activation | Test loss | Test accuracy (%) |
|-----------|------------|---------------|-------------------|
| CIFAR-100 | ELU | 1.5609 | 57.26 |
| | Hardswish | 1.3122 | 64.12 |
| | LeakyReLU | 1.4248 | 61.71 |
| | ReLU | 1.4223 | 61.84 |
| | ReLU6 | 1.4185 | 61.58 |
| | RReLU | 1.4509 | 59.81 |
| | SELU | 1.8315 | 51.09 |
| | GELU | <i>1.3351</i> | 64.71 |
| STL-10 | ELU | 1.5533 | 41.78 |
| | Hardswish | 1.2457 | 54.40 |
| | LeakyReLU | 1.1650 | 56.26 |
| | ReLU | 1.2105 | 54.86 |
| | ReLU6 | 1.5044 | 47.01 |
| | RReLU | 1.2814 | 51.25 |
| | SELU | 1.5221 | 41.18 |
| | GELU | <i>1.1853</i> | 58.48 |

*Bold indicates the best performance; italic indicates the second-best.

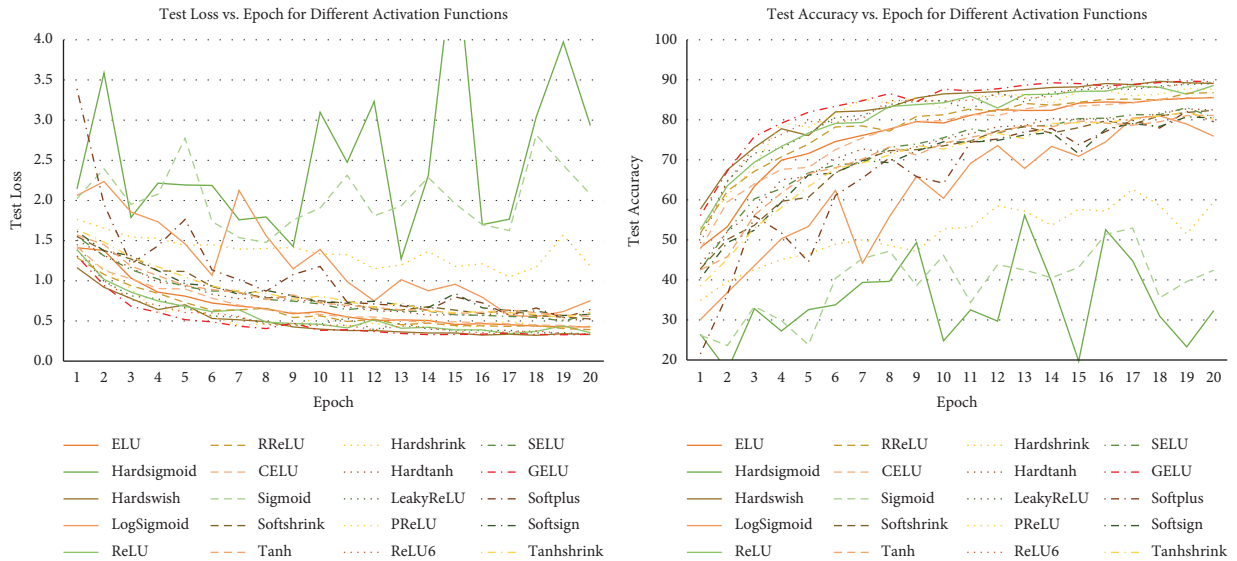


FIGURE 4: Experimental comparison of activation functions with respect to training epoch.

counterparts in terms of test accuracy, thereby reinforcing the assertion that it is a highly effective activation function for deep learning applications.

For the CIFAR-100 dataset, the GELU activation function achieved the highest test accuracy of 64.71%, surpassing the second-best performance of 64.12% achieved by the Hardswish activation function. The test loss for GELU was marginally higher than that of Hardswish, with values of 1.3351 and 1.3122, respectively. However, considering the higher test accuracy, GELU still demonstrates a more consistent performance across both evaluation metrics. Other activation functions, such as ReLU, LeakyReLU, and RReLU, exhibited competitive performance, with test accuracies ranging between 59.81% and 61.84%. Nevertheless, their performance remained inferior to that of GELU, further highlighting its efficacy in the context of the CIFAR-100 dataset.

Similarly, on the STL-10 dataset, the GELU activation function outperformed all other activation functions, achieving the highest test accuracy of 58.48%. LeakyReLU secured the second-best performance with a test accuracy of 56.26%. However, in terms of test loss, GELU was slightly higher at 1.1853, compared to the best value of 1.1650 observed for LeakyReLU. Despite this minor discrepancy, the overall performance of GELU remains superior, as evidenced by its higher test accuracy. Other activation functions, such as ReLU and Hardswish, showcased relatively competitive performance but ultimately fell short of the performance exhibited by GELU.

These additional experiments on the CIFAR-100 and STL-10 datasets reinforce the notion that GELU is a highly effective activation function for deep learning models. Its consistently superior performance across multiple evaluation metrics and datasets attests to its robustness and

adaptability, making it a compelling choice for practitioners seeking optimal activation functions for their deep learning applications. Moreover, the results of this empirical analysis complement our earlier mathematical investigation of GELU's properties, together offering a comprehensive understanding of its performance and suitability in a wide range of deep learning scenarios.

5. Conclusion

In this comprehensive study, we have embarked upon an intricate exploration of the GELU activation function and its mathematical properties, including differentiability, boundness, stationarity, and smoothness. Our analysis elucidates the unique characteristics that contribute to GELU's efficacy in the context of deep learning architectures. GELU's smoothness, differentiability, and well-behaved optimization landscape have cemented its position as an indispensable asset in state-of-the-art models such as BERT and GPT.

Furthermore, we have conducted a rigorous experimental comparison of various activation functions within the context of residual convolutional networks trained on the CIFAR-10, CIFAR-100, and STL-10 datasets. Our findings reinforce the exceptional performance of the GELU activation function, which attains the highest test accuracy and lowest test loss among the activation functions investigated. Other activation functions, such as Hardswish and ReLU6, exhibit commendable performance as well, highlighting their potential applicability in diverse deep learning scenarios.

Looking forward, the mathematical properties of the GELU activation function as derived in this study offer promising directions for the generalization analysis of deep neural networks. Specifically, our results regarding the Lipschitz continuity and smoothness of GELU contribute towards theoretical understanding of the flexibility and adaptivity of deep neural networks, potentially paralleling the recent studies [35]. In addition, our insights about the higher-order derivatives and Holder continuity of GELU can facilitate the analysis of deep learning models in more complex spaces, such as the unit sphere of high-dimensional spaces [36]. It is anticipated that the integration of our mathematical analyses with the theoretical frameworks proposed in these works could pave the way for a more robust theoretical foundation of deep learning. We also recommend future work to investigate the implications of our findings for the approximation ability of new activation functions such as DLU [37].

In conclusion, our in-depth analysis and experimental evaluation substantiate the GELU activation function's prominence in the realm of deep learning. The GELU function's mathematical properties and exemplary performance render it a potent choice for a wide array of applications, providing a foundation for future research and innovation in the field of artificial intelligence. In this paper, we have presented a comprehensive mathematical analysis of the GELU activation function and normalization methods in deep learning, specifically focusing on differentiability,

boundness, feature space continuity, stationarity, and smoothness of the feature space. Our findings provide insights into the reasons behind the success of these methods and their impact on the training dynamics of deep neural networks. We hope that our work contributes to the understanding of GELU activation and normalization techniques and informs the design of future deep learning architectures.

Data Availability

No underlying data were collected or produced in this study.

Disclosure

An earlier version of this paper was published as a preprint on <https://arxiv.org/abs/2305.12073> [38].

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

This work was supported by a research grant funded by Generative Artificial Intelligence System Inc. (GAIS) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00251528).

References

- [1] J. G. Greener, S. M. Kandathil, L. Moffat, and D. T. Jones, "A guide to machine learning for biologists," *Nature Reviews Molecular Cell Biology*, vol. 23, no. 1, pp. 40–55, 2022.
- [2] B. Liu, M. Ding, S. Shaham, W. Rahayu, F. Farokhi, and Z. Lin, "When machine learning meets privacy: a survey and outlook," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–36, 2021.
- [3] M. Lee, "A mathematical investigation of hallucination and creativity in gpt models," *Mathematics*, vol. 11, no. 10, p. 2320, 2023.
- [4] G. Carleo, I. Cirac, K. Cranmer et al., "Machine learning and the physical sciences," *Reviews of Modern Physics*, vol. 91, no. 4, Article ID 045002, 2019.
- [5] M. Lee, "The geometry of feature space in deep learning models: a holistic perspective and comprehensive review," *Mathematics*, vol. 11, no. 10, p. 2375, 2023.
- [6] X. Dong, J. Bao, D. Chen et al., "Cswin transformer: a general vision transformer backbone with cross-shaped windows," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12124–12134, New Orleans, LA, USA, June 2022.
- [7] M. Antonelli, A. Reinke, S. Bakas et al., "The medical segmentation decathlon," *Nature Communications*, vol. 13, no. 1, p. 4128, 2022.
- [8] S. Minaee, Y. Y. Boykov, F. Porikli, A. J. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: a survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3523–3542, 2022.
- [9] I. Kim, M. Lee, and J. Seok, "Icegan: inverse covariance estimating generative adversarial network," *Machine Learning: Science and Technology*, vol. 4, no. 2, Article ID 025008, 2023.

- [10] S. R. Choi and M. Lee, "Estimating the prognosis of low-grade glioma with gene attention using multi-omics and multi-modal schemes," *Biology*, vol. 11, no. 10, p. 1462, 2022.
- [11] M. Lee, "An ensemble deep learning model with a gene attention mechanism for estimating the prognosis of low-grade glioma," *Biology*, vol. 11, no. 4, p. 586, 2022.
- [12] M. Lee and J. Seok, "Estimation with uncertainty via conditional generative adversarial networks," *Sensors*, vol. 21, no. 18, p. 6194, 2021.
- [13] H. Yun, M. Lee, Y. S. Kang, and J. Seok, "Portfolio management via two-stage deep learning with a joint cost," *Expert Systems with Applications*, vol. 143, Article ID 113041, 2020.
- [14] S. R. Dubey, S. K. Singh, and B. B. Chaudhuri, "Activation functions in deep learning: a comprehensive survey and benchmark," *Neurocomputing*, vol. 503, 2022.
- [15] A. Apicella, F. Donnarumma, F. Isgrò, and R. Prevete, "A survey on modern trainable activation functions," *Neural Networks*, vol. 138, pp. 14–32, 2021.
- [16] L. Lu, Y. Shin, Y. Su, and E. K. George, "Dying relu and initialization: theory and numerical examples," 2019, <https://arxiv.org/abs/1903.06733>.
- [17] D. Hendrycks and K. Gimpel, "Gaussian error linear units (gelus)," 2016, <https://arxiv.org/abs/1606.08415>.
- [18] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova, "Bert: pre-training of deep bidirectional transformers for language understanding," 2019, <https://arxiv.org/abs/1810.04805>.
- [19] A. Dosovitskiy, L. Beyer, K. Alexander et al., "An image is worth 16x16 words: transformers for image recognition at scale," *International Conference on Learning Representations*, vol. 15, 2021.
- [20] T. Brown, B. Mann, N. Ryder et al., "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.
- [21] Y. Tian, D. Su, S. Lauria, and X. Liu, "Recent advances on loss functions in deep learning for computer vision," *Neurocomputing*, vol. 497, pp. 129–158, 2022.
- [22] T. Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," *Proceedings of the IEEE International Conference on Computer Vision*, vol. 42, no. 2, pp. 2980–2988, 2017.
- [23] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, "Generative adversarial networks: an overview," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 53–65, 2018.
- [24] A. Shrestha and A. Mahmood, "Review of deep learning algorithms and architectures," *IEEE Access*, vol. 7, pp. 53040–53065, 2019.
- [25] D. P. Kingma and B. Jimmy, "Adam: a method for stochastic optimization," 2014, <https://arxiv.org/abs/1412.6980>.
- [26] O. T. Ilya, N. Houlsby, K. Alexander et al., "Mlp-mixer: an all-mlp architecture for vision," *Advances in Neural Information Processing Systems*, vol. 34, pp. 24261–24272, 2021.
- [27] Z. Dai, H. Liu, V. L. Quoc, and M. Tan, "Coatnet: marrying convolution and attention for all data sizes," *Advances in Neural Information Processing Systems*, vol. 34, pp. 3965–3977, 2021.
- [28] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "Vivit: a video vision transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6836–6846, Montreal, Canada, January 2021.
- [29] A. Radford, J. W. Kim, C. Hallacy et al., "Learning transferable visual models from natural language supervision," in *Proceedings of the International Conference on Machine Learning*, pp. 8748–8763, PMLR, Vienna, Austria, July 2021.
- [30] S. Ioffe and C. Szegedy, "Batch normalization: accelerating deep network training by reducing internal covariate shift," in *Proceedings of the International Conference on Machine Learning*, pp. 448–456, pmlr, Lille, France, July 2015.
- [31] J. Lei, J. Ryan Kiros, and G. E. Hinton, "Layer normalization," 2016, <https://arxiv.org/abs/1607.06450>.
- [32] Y. Wu and K. He, "Group normalization," in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, Tel Aviv, Israel, October 2018.
- [33] K. Asadi, D. Misra, and M. Littman, "Lipschitz continuity in model-based reinforcement learning," in *Proceedings of the International Conference on Machine Learning*, pp. 264–273, PMLR, Vienna, Austria, August 2018.
- [34] G. Cupini, M. Focardi, F. Leonetti, and E. Mascolo, "On the hölder continuity for a class of vectorial problems," *Advances in Nonlinear Analysis*, vol. 9, no. 1, pp. 1008–1025, 2019.
- [35] T. Mao, Z. Shi, and D.-X. Zhou, "Approximating functions with multi-features by deep convolutional neural networks," *Analysis and Applications*, vol. 21, no. 01, pp. 93–125, 2023.
- [36] Z. Fang, H. Feng, S. Huang, and D.-X. Zhou, "Theory of deep convolutional neural networks ii: spherical analysis," *Neural Networks*, vol. 131, pp. 154–162, 2020.
- [37] J. Li, F. Han, and D.-X. Zhou, "A new activation for neural networks and its approximation," 2022, <https://arxiv.org/abs/2210.10264>.
- [38] M. Lee, "Gelu activation function in deep learning: a comprehensive mathematical analysis and performance," 2023, <https://arxiv.org/abs/2305.12073>.