Stochastic Adaptive Activation Function

Anonymous Author(s)
Affiliation
Address
email

Abstract

Human neurons and neurotransmission mechanisms have been realized in a deep neural network by the theoretical implementations of activation functions. However, recent studies have reported that the threshold potential of neurons exhibited different values by locations and types of individual neurons, and the activation functions have limitations in representing it. Therefore, this study proposes a simple yet effective activation function that exhibits different thresholds and adaptive activations according to the positions of units and the contexts of inputs. Furthermore, the devised activation function mathematically exhibits a more generalized form of Swish activation function, and thus we denoted it as Adaptive SwisH (ASH). ASH highlights informative features, which exhibit large values in the top percentiles in an input, whereas it rectifies low values. Most importantly, ASH exhibits trainable, adaptive, and context-aware properties compared to other activation functions. Furthermore, ASH represents general formula of the previously studied activation function and provides a reasonable mathematical background for the superior performances. To validate the effectiveness and robustness of ASH, we implemented ASH into many deep learning models for various tasks, including classification, detection, segmentation, and image generation. The experimental analysis demonstrates that our activation function could provide benefits of accurate prediction as well as earlier convergence in many deep learning applications. [1]

1 Introduction

Searching for the optimal activation function has been challenging in artificial intelligence (Maas et al., 2013; Ramachandran et al., 2017; Clevert et al., 2015). Early activation functions have been studied to compensate for the non-linearity of the artificial neural networks or ameliorate the gradient vanishing problem (Hertz et al., 1997; Hochreiter, 1998). Recently, novel activation functions have been suggested for the zero-centered or parametric properties that improve the training efficiency of deep neural networks (Maas et al., 2013; Clevert et al., 2015). Currently, the activation functions regarding the stability of deep neural networks or probabilistic distribution of inputs have been proposed (Hendrycks, Gimpel, 2016; Misra, 2019). The advances in the activation functions have opened the era of deep neural networks in the creative tasks of detecting/segmenting target objects in sophisticated images or generating images beyond the simple classifiers (Simonyan, Zisserman, 2014; Zhao et al., 2019; Badrinarayanan et al., 2017; Goodfellow et al., 2014).

The development of activation functions has evolved to behave more like a human neuron (Sharma et al., 2017; Lee et al., 2017). However, Izhikevich (2003); Evans et al. (2018) reported that the neurotransmission mechanism, including a membrane, action, and threshold potential of human neurons, is subject to the location or the connection type of the neurons. In addition, humans perceive objects with surrounding contexts using $N : N$ mapping of visions to human neurons rather than 1:1

[1] Our code is available at https://github.com/Anonymous/Repo

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

mapping of a pixel to an input node in neural networks (Liu et al., 2018). The connections between neurons can be realized by the linear combinations of layers in a deep neural network, but deep neural networks have limitations in realizing contextual perception. It implies that there is room for improvement in deep neural networks and convolutional neural networks despite the superior performance in analyzing images (Jinsakul et al., 2019; Misra, 2019). Therefore, the development of a deep neural network is leaned to mimic human perception by realizing the mechanism of human neurons (Aggarwal, others, 2018; Lindsay, 2021).

Currently, the primary issue is that activation functions exhibit passivity, in this paper, indicating that the activation function determines the output only concerning the value of one element rather than the entire contexts. For instance, the Rectified Linear Unit (ReLU), defined as $f(x) = \max(x, 0)$, determines the output values related to x (Fukushima, Miyake, 1982), whereas the *softmax* function generates output values as the ratio of the input value to the totals (Goodfellow et al., 2017); ReLU exhibits passivity, whereas *softmax* does not. Suppose an image can be classified by considering 80% of the total portion. The current activation functions exhibit the limitation in classifying this image since only elements (pixels) of the image are used to rectify the image rather than the ratio. Additionally, another issue is that the activation function is invariant. Even though the parametric activation function trains its parameters during training, the values of the parametric activation functions are invariant in the inference phase (Xu et al., 2015; Bingham, Miikkulainen, 2022). Therefore, the limited rectification could be realized by the invariant parameters or thresholds regardless of new inputs from different domains (e.g., test set).

Contribution To realize the mechanism of human neurons that rectify inputs considering their contexts, we propose ASH activation function. The main contributions of this study are to suggest a simple yet effective activation function, ASH, and to implement ASH in a mathematically effective form. Going beyond the passive activation functions, ASH activation function is designed as (1) an active activation function that provides outputs regarding the context of inputs and (2) a conditional activation function that employs an adaptive threshold. Unlike ReLU or Leaky ReLU, the threshold value of ASH is adaptively changing considering the contextual information.

$$f(x) = \begin{cases} x & \text{if } x \geq \theta, \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

In particular, the threshold value (θ) is adaptively changing according to the input distribution without heavy calculation, and thus ASH provides outputs considering the contexts of inputs adaptively. By applying ASH, we obtained the following theoretical and experimental results:

- We conducted mathematical modeling for ASH in an effective form to be trainable and parametric, and thus ASH exhibits parametric and adaptive properties. The baseline threshold of ASH is trainable during the training phase, and the threshold value is adaptively changing according to the contexts of inputs without heavy calculation.
- We theoretically verified ASH adaptively changes its threshold alongside the stochastic distribution of inputs. This implies ASH provides outputs regarding the entire contexts of inputs, as well as leads to better feature extraction.
- We theoretically verified that ASH exhibited general formula for Swish activation function and provided the mathematical explanations for the superior performance of Swish, which was empirically searched in the previous works.
- We experimentally showed that ASH improves the performance of deep learning models for various tasks as well as shortens training time, in terms of early convergence

Related works Activation functions affect the performance of the training process to determine functional subspace in a deep neural network (DNN) (Hayou et al., 2019). In particular, the non-linearity using activation functions has been introduced to avoid the linear transform causing simple feature extractions in the DNN (Misra, 2019; Jarrett et al., 2009). Therefore, the DNN with non-linearity has realized any complex tasks in the current (Leshno et al., 1993). In the early era, Rectified Linear Unit (ReLU) replaced the classical activation functions such as sigmoid and tanh (Nair, Hinton, 2010), due to its simple and computational efficiency compared to other activation functions. During decades, many activation functions have been proposed such as Leaky ReLU (Maas et al., 2013), Exponential Linear Unit (ELU) (Clevert et al., 2015), Gaussian Error Linear Unit (GELU)

(Hendrycks, Gimpel, 2016), Scaled Exponential Linear Unit (SELU) (Klambauer et al., 2017), and Swish (Ramachandran et al., 2017) for better performance and stability of learning parameters in DNNS. Those activation functions have solved the dying ReLU problem, which exhibits zero value in the negative region and improved DNNs more smoothly for stable optimization. In particular, ELU and SELU have realized internal normalization in the layer using the zero-mean property. Furthermore, GELU exploited a Gaussian error that could implement adaptive dropout to apply a higher probabilistic intuition. In recent, more advanced activation functions have exhibited trainable properties to achieve flexibility of optimization in many deep learning networks (Arjovsky et al., 2016; Molina et al., 2019).

Problem statement For adaptive thresholding, ASH exploits stochastic selection methodology like quick selection (Hoare, 1961). In particular, for better feature extraction, the informative elements, which exhibit large values, should be identified like an attention mechanism. In contrast, some elements, which exhibit low relevance, are required to be rectified. To this end, ASH is designed to identify informative elements but rectify others as 0.

Let $X \in R^{H \times W \times C}$ be a tensor (here, feature-map) disregarding batch, but with height (H), width (W), and channel (C), and let $\bar{X} \ni X$ be a set of feature-maps. Let $\mathcal{A}$ be an activation function such that $\mathcal{A} : \bar{X} \rightarrow \bar{X}$. Then, we can extract i^{th} element from X, and denote it as $x^{(i)}$. Here, the goal of this study is to design an activation function represented as follows:

$$\mathcal{A}(x^{(i)}) = \begin{cases} x^{(i)} & \text{if } x^{(i)} \text{ is ranked in top } k\% \text{ percentile of } X, \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Note that, a novel activation function $\mathcal{A}$ is represented similar to ReLU, whereas the threshold is not invariant compared to ReLU, but is subjected to the distribution of input X. To simplify, let $C(x^{(i)}, k; X)$ be a condition whether $x^{(i)}$ is ranked in top $k\%$ percentile of X, and the negation of C is denoted as $\neg C$. In particular, the elements that satisfy $C(x^{(i)}, k; X)$ are from the first largest element to the $(0.01kN)$-th largest element in X, where N indicates the number of elements in X. Here, the elements can be extracted using a simple algorithm like a quick selection. Suppose a set $\hat{X}$ that includes all elements in X such that $\hat{X} = \{x^{(1)}, x^{(2)}, ..., x^{(N)}\}$. Then, we can construct subsets of $\hat{X}$ as $\hat{X}_C = \{x^{(i)} \in X | C(x^{(i)}, k; X)\}$ and $(\hat{X}_C)^c = \{x^{(i)} \in X | \neg C(x^{(i)}, k; X)\}$. Then, Equation (2) can be simplified as follows:

$$\mathcal{A}(x^{(i)}) = \begin{cases} x^{(i)} & \text{if } x^{(i)} \in \hat{X}_C, \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

To summarize, the activation function is designed to sample the top-k percentile from input, where the criteria is the value (less or greater) of the element. The sampling examples are illustrated in Fig. 1 compared to ReLU.

(a) Feature-map (b) ReLU (c) ASH (80%) (d) ASH (50%) (e) ASH (30%)

Figure 1: Input feature-map and outputs by activation functions. ASH ($k\%$) indicates that ASH activation function samples top-$k\%$ elements from input feature-map. The sampled elements by the activation functions are colored as blue.

2 Method

This study aims at designing an activation function for a stochastic sampling of the top-k percentile elements from inputs. However, sorting or sampling methods like quick selection requires high

computational cost. In contrast, sampling the top-k percentile could be realized by a z-score-based method in a simple calculation despite the prerequisites of normal distribution (also known as Gaussian distribution). Since we reviewed that the outputs of neurons are normally distributed, we employ the stochastic sampling to design ASH activation function.

In this section, we (1) reviewed that the outputs of neurons are distributed in a normal distribution, (2) conducted modeling for stochastic sampling using a z-score, (3) formulated ASH activation function, (4) verified ASH is parametric and trainable, and (5) searched general application of ASH activation function.

2.1 Background

Gaussian distribution Currently, many deep learning models employ a normalization method due to its stability in training deep learning models (Ioffe, Szegedy, 2015; Ulyanov et al., 2016; Wu, He, 2018). In the previous study, Ioffe, Szegedy (2015) reported that the output of the convolutional layer, $x = Wu + b$, is more likely to have a symmetric, non-sparse distribution, that is "more Gaussian". Since most deep learning models are based on convolutional operations, the outputs of neurons are supposed to be normally distributed (Gaussian distribution). Therefore, it is concluded that the inputs of the activation functions are normally distributed when activation functions follow convolutional layers, such that $x \sim N(\mu_x, \sigma_x^2)$, where x is an input feature-map of an activation function, μ_x and σ_x are mean and standard deviation of x, respectively. Therefore, we obtained the following.

Proposition 1. The outputs of neurons in convolutional neural networks are normally distributed.

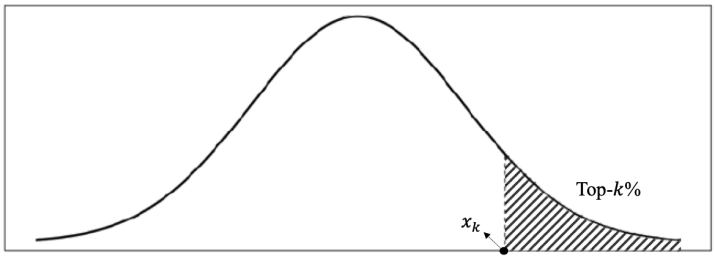

Figure 2: Top-k% sampling from normal distribution

Sampling from normal distribution This study aims at sampling the top-k% elements, which exhibit greater values, from an input feature-map. Since the area under normal distribution indicates the percentile, sampling the top-k% percentile from a normal distribution is theoretically identical to the statement calculating the area under curve shown in Fig. 2. Let F be a tensor, a multi-dimensional array or a matrix. Then, F is supposed to be normally distributed, and thus we could sample the elements ($f^{(i)}$) in the top-k% percentile from F using the following equation:

$$P(f^{(i)} \geq z'_k) = k, \quad k \in [0,1] \quad z'_k \in [-\infty, \infty] \tag{4}$$

However, heavy computational costs are required to find the trivial solution from the probability density function of normal distribution, defined as $\frac{1}{\sigma_F\sqrt{2\pi}} e^{-(x-\mu_F)^2/2\sigma_F^2}$. Therefore, we leaned to *probability theory* to simplify the computation rather than *calculus*. To this end, we employed the standard normal distribution (Z-score normalization) for Equation (4), and we obtained the following:

$$P(Z^{(i)} \geq z_k) = k, \quad Z^{(i)} = \frac{f^{(i)} - \mu_F}{\sigma_F} \quad s.t. \quad Z^{(i)} \sim N(0,1) \tag{5}$$

where $z_k = (z'_k - \mu_F)/\sigma_F$, which is the Z-normalized value from z'_k, and thus z_k is subjected to k, indicating percentile to sample, in terms of Z-Table (Larsen, Marx, 2005). Then, we go Z-table and easily find the proper value for z_k, intuitively. Therefore, the condition, $Z^{(i)} = (f^{(i)} - \mu_F)/\sigma_F \geq z_k \Leftrightarrow f(i) \geq \mu_F + z_k\sigma_F$, is mathematically identical to sample the top-k% elements from F. To summarize, we obtained the following.

Proposition 2. Element $x^{(i)} \in X$, which is normally distributed, is in the top-$k\%$ percentile of X if $x^{(i)} \geq \mu_X + z_k\sigma_X$, where z_k is a z-value subjected to k in Z-Table (Larsen, Marx, 2005).

Differentiation The mechanism of convolutional neural network (CNN) has been studied in many previous works (Rumelhart et al., 1986; Bottou, 2010; Zhang, 2016; Hu et al., 2018). Training CNN models is subjected to the backpropagation derived from the partial derivatives of loss function by the individual convolutional parameter. Let L be a loss function for deep learning model M and let W be one of the variables in M. Then, the derivative of L in terms of W is represented as $\frac{\partial L}{\partial W}$, and W is updated as $W \leftarrow W - \eta\frac{\partial L}{\partial W}$ with learning rate η. On the other hand, suppose variable θ be the threshold and the function $f(x)$ is αx if $x \geq \theta$, otherwise 0, such that the partial derivative of f is represented as following:

$$\frac{\partial f}{\partial x} = \begin{cases} \alpha, \text{ if } x \geq \theta \\ 0, \text{ otherwise} \end{cases}, \quad \frac{\partial f}{\partial \alpha} = \begin{cases} x, \text{ if } x \geq \theta \\ 0, \text{ otherwise} \end{cases}, \quad \frac{\partial f}{\partial \theta} = 0 \tag{6}$$

Here, α is arithmetically combined with $f(x)$, whereas θ does not. Therefore, α is trainable, and θ is not trainable in this context. In basic *calculus*, it is trivial that if the loss function is not dependent on the variable, the partial derivative is zero, and thus the variable cannot be trained or optimized; invariant. Therefore, we obtain the following:

Lemma 1. The derivative of a variable in the conditional statement is zero, and thus that variable cannot be optimized.

2.2 ASH Activation Function

Let X be an input of ASH activation function (*Æ*) and be a tensor of which elements are normally distributed. Besides, let $x^{(i)} \in X$ be the i-th element of X. Then, by **Proposition 2**, ASH activation function, which samples the top-$k\%$ elements from input, is represented as following:

$$Æ(x^{(i)}) = \begin{cases} x^{(i)} & \text{if } x^{(i)} \geq \mu_X + z_k\sigma_X, \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where μ_X and σ_X are the mean and the standard deviation of all elements in X, and z_k is the z-score concerning percentile (k) to sample (i.e. $z = 1.96$ if $k = 2.5\%$, See Z-Table). Equation (7) exhibits that ASH activation function is represented in a simple yet effective form with low computational costs.

Intuitively, we assumed that the activation level (percentile) is supposed to be different by each neuron and the tasks of the deep learning model like human neurons. However, in Equation (7), the condition ($x^{(i)} \geq \mu_X + z_k\sigma_X$) is invariant by **Lemma 1**. Note that, $\mu_X + z_k\sigma_X$ is variable and changeable concerning X, but the sampled portion ($k\%$ percentile related to z_k) from X is invariant, and thus ASH in Equation (7) is not parametric and trainable. To make ASH be trainable and parametric, let Equation (7) be substituted using a proxy function as $Æ(x^{(i)}) = x^{(i)}f(x^{(i)})$, such that the proxy function $f(x^{(i)})$ is represented as the following:

$$f(x^{(i)}) = \begin{cases} 1 & \text{if } x^{(i)} - \mu_X - z_k\sigma_X \geq 0, \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

To simplify, suppose Heaviside step function (Weisstein, 2002), defined as $H(x) = \frac{d}{dx}\max(0, x)$. Then, $f(x^{(i)}) = H(x^{(i)} - \mu_X - z_k\sigma_X)$. Therefore, we obtain the arithmetical form to formulate ASH activation function as the following:

$$Æ(x^{(i)}) = x^{(i)}H(x^{(i)} - \mu_X - z_k\sigma_X) \tag{9}$$

Despite the arithmetic formula, ASH activation function in Equation (9) is still independent of z_k, and thus the z_k is still not trainable. However, it is well known that the Heaviside step function is analytically approximated as $2H(x) = 1 + 1\tanh(\alpha x)$ with a large value of α (Iliev et al., 2017),

and thus ASH activation function is approximated using the smooth Heaviside step function as the following:

$$\begin{aligned} Æ(x^{(i)}) &= x^{(i)} H(x^{(i)} - \mu_X - z_k \sigma_X) \\ &= \frac{1}{2} x^{(i)} + \frac{1}{2} x^{(i)} \tanh(\alpha(x^{(i)} - \mu_X - z_k \sigma_X)) \\ &= \frac{x^{(i)}}{1 + e^{-2\alpha(x^{(i)} - \mu_X - z_k \sigma_X)}} \end{aligned} \tag{10}$$

Since z_k is arithmetically placed, z_k indicating sampling percentiles is trainable, and thus ASH activation function is also trainable and parametric. By optimizing z_k, ASH activation functions exhibit different thresholds for sampling percentiles. Therefore, it is concluded that ASH exhibits different activation levels by a stochastic sampling of inputs and different thresholds like human neurons, synapses, and their potential. Like human neurons, the mechanism of ASH is the following:

(1) In the training phase, each ASH activation function optimizes its z_k and fine-tunes the threshold for the percentile sampling of inputs. Thus, it implies that ASH activation function realizes the arbitrary threshold potentials in human neurons (Clevert et al., 2015; Evans et al., 2018). Some examples of z_k related to Equation (7) are:

Example 1. Small value of z_k, even negative, implies the dense activation, and the dying ReLU problem could be improved.

Example 2. Large value of z_k implies the sparse activation, and the sparsity could be employed.

Furthermore, (2) in the training or inference phase, ASH activation function rectifies the inputs using the learnt threshold value and contexts of inputs. In particular, to sample the top-$k\%$ percentile, ASH employs the mean and the standard deviation of input, representing the contexts of the entire input. Therefore, it implies that ASH realizes the adaptive activation considering the contexts of inputs. Some examples of the adaptive activation related to Equation (7) are:

Example 3. Small threshold value (θ_s) is employed to calculate input X that exhibits large mean and standard deviation ($X > \theta_s$).

Example 4. Large threshold value (θ_l) is employed to calculate input X' that exhibits large mean and standard deviation ($X' > \theta_l \gg \theta_s$).

2.3 Generalized Activation Function

We found the innovation while representing Equation (10) using the sigmoid function $S(x) = \frac{1}{1+e^{-x}}$ as the following:

$$\begin{aligned} Æ(x^{(i)}) &= x^{(i)} S\big(-2\alpha(x^{(i)} - \mu_X - z_k \sigma_X)\big) \\ &= x^{(i)} S(a x^{(i)} + b)) \end{aligned} \tag{11}$$

In the previous work, Ramachandran et al. (2017) introduced the leverage of automatic search techniques to discover the best performance activation function. The experiments empirically discovered that the Swish activation function is the best performance activation function, defined as $xS(x)$ (Ramachandran et al., 2017). Intuitively, the definition of the Swish activation function is the same with Equation (11), and Equation (11) represents more generalized formula. Therefore, ASH (Adaptive SwisH) activation function provides the theoretical explanations for why Swish was the best performance activation function in the empirical evaluations. Therefore, we obtain the following.

Lemma 3. ASH activation function exhibits general formula for Swish activation function.

Interestingly, the activation function designed for stochastic adaptive sampling is converged to the generalized Swish activation function. The extreme impression is that the stochastic percentile sampling by the activation function that mimics real neurons expresses the general formula of the swish activation function formerly known as state-of-the-art. Therefore, the stochastic percentile sampling could partially be applied to the Swish activation function. In addition, it could be supposed that the Swish activation function achieved superior performance in the previous study due to the utilization of stochastic percentile sampling.

This paper initially addressed the activation function that enables stochastic percentile sampling in a mathematically effective manner, but we found that the mathematical expression of ASH supports the theoretical background of the Swish activation function. Therefore, this paper provides the theoretical backgrounds and rationales for the Swish activation function, which was empirically searched. It is a significant innovation to provide the mathematical theorem that the Swish activation function is derived from a stochastically designed activation function.

3 Main Result

Like the previous study (Ramachandran et al., 2017), we compare ASH against several additional baseline activation functions on various models for a variety of tasks using various datasets. Since many activation functions have been introduced, we employ commonly used activation functions: ReLU, Leaky ReLU (LReLU) (Maas et al., 2013), Parametric ReLU (PReLU) (He et al., 2015), Softplus (Nair, Hinton, 2010), Exponential Linear Unit (ELU) (Clevert et al., 2015), Scaled Exponential Linear Unit (SELU) (Klambauer et al., 2017), Gaussian Error Linear Unit (GELU) (Hendrycks, Gimpel, 2016). In the experiments, every hyper-parameter between ASH and other activation functions is restricted to be the same to exhibit the contribution of ASH compared to other activation functions. In tables, the highest accuracies are highlighted as **bold**.

3.1 Classification Task

We first compare ASH to all baseline activation functions on the datasets of CIFAR-10 (Krizhevsky et al., 2009), CIFAR-100 (Krizhevsky et al., 2009), and ImageNet (Russakovsky et al., 2015). We employ the environments from the previous study (Ramachandran et al., 2017), and re-implement the baseline models of the ResNet-164 (He et al., 2016), Wide ResNet28-10 (Zagoruyko, Komodakis, 2016), and DenseNet-100-12 (Huang et al., 2017). Due to different environments, small differences exist from the previous paper, but we believe that accuracy trends are the same. We first evaluate ASH activation function against other activation functions on the ImageNet 2012 classification dataset. Since ImageNet is the widely utilized dataset in the classification tasks, we employ the ImageNet2012 dataset to measure the performances of ASH and other activation functions. We then evaluate all activation functions using CIFAR-10 and CIFAR-100 datasets widely utilized in the benchmark.

Model	Top-1 Acc. (%)			Top-5 Acc. (%)		
ReLU	76.4 ± 0.09	75.6 ± 0.10	77.1 ± 0.11	91.2 ± 0.09	90.7 ± 0.06	90.7 ± 0.06
LReLU	77.6 ± 0.07	78.0 ± 0.03	76.6 ± 0.07	91.6 ± 0.10	91.2 ± 0.07	92.3 ± 0.07
PLeLU	77.0 ± 0.13	78.7 ± 0.03	78.0 ± 0.09	92.9 ± 0.03	92.3 ± 0.14	92.2 ± 0.12
Softplus	76.8 ± 0.11	77.3 ± 0.03	76.0 ± 0.05	91.5 ± 0.12	93.7 ± 0.05	93.8 ± 0.11
ELU	71.6 ± 0.09	73.7 ± 0.09	74.9 ± 0.10	85.6 ± 0.06	89.8 ± 0.13	90.2 ± 0.08
SELU	75.4 ± 0.13	78.1 ± 0.14	76.8 ± 0.06	91.6 ± 0.09	93.5 ± 0.10	90.9 ± 0.04
GELU	76.8 ± 0.13	77.9 ± 0.05	78.0 ± 0.12	90.2 ± 0.11	92.6 ± 0.04	91.9 ± 0.09
Swish	77.5 ± 0.07	76.6 ± 0.06	76.5 ± 0.05	92.2 ± 0.12	90.9 ± 0.07	92.2 ± 0.07
ASH	**78.5 ± 0.06**	**78.6± 0.07**	**78.7± 0.10**	**94.0± 0.08**	**94.7± 0.07**	**94.1± 0.08**

Table 1. ImageNet dataset. Three models are averaged. The values are mean and 95% Confidence Interval (C.I.)

Model	ResNet	WRN	DenseNet
ReLU	94.4 ± 0.04	95.6 ± 0.03	95.7 ± 0.02
LReLU	94.5 ± 0.05	95.6 ± 0.04	94.7 ± 0.09
PLeLU	94.7 ± 0.08	95.4 ± 0.03	95.1 ± 0.08
Softplus	94.3 ± 0.10	94.2 ± 0.08	95.2 ± 0.07
ELU	93.5 ± 0.10	93.8 ± 0.09	94.5 ± 0.11
SELU	94.5 ± 0.05	95.8 ± 0.07	94.9 ± 0.10
GELU	95.2 ± 0.04	95.7 ± 0.06	94.8 ± 0.10
Swish	95.5 ± 0.09	95.6 ± 0.08	95.2 ± 0.03
ASH	**95.7 ± 0.08**	**96.7 ± 0.04**	**96.0 ± 0.11**

Table 2. CIFAR-10 with mean values and 95% C.I.

Model	ResNet	WRN	DenseNet
ReLU	74.5 ± 0.10	78.4 ± 0.04	84.0 ± 0.09
LReLU	75.3 ± 0.06	77.9 ± 0.07	82.2 ± 0.07
PLeLU	74.7 ± 0.06	77.6 ± 0.06	82.1 ± 0.07
Softplus	76.1 ± 0.05	78.6 ± 0.06	84.1 ± 0.02
ELU	75.0 ± 0.08	76.4 ± 0.09	80.8 ± 0.05
SELU	73.4 ± 0.05	74.4 ± 0.08	81.4 ± 0.06
GELU	75.0 ± 0.05	78.2 ± 0.10	84.0 ± 0.02
Swish	75.7 ± 0.10	78.9 ± 0.05	84.0 ± 0.03
ASH	**76.5 ± 0.08**	**79.2 ± 0.06**	**84.6 ± 0.06**

Table 3. CIFAR-100 with mean values and 95% C.I.

Here, the ImageNet dataset evaluations are averaged from the resultant accuracy of three deep learning models. The results in Tables 1-3 exhibit the outstanding performance of ASH activation function in improving predictive accuracy. Since the deep learning models for the classification tasks demands sparsity, it is expected that ASH activation function could improve the accuracy compared to other activation functions.

3.2 Detection Task

We then compare ASH to all baseline activation functions on the datasets of COCO (Lin et al., 2014) and PASCAL VOC (Everingham et al., 2010) for the detection task. Both datasets are widely utilized in the benchmark of the detection task. We employ the same environments from the classification task and implement the baseline models of the Mask-R-CNN (He et al., 2017), SSD (Liu et al., 2016), and YOLOv4 (Bochkovskiy et al., 2020). In the detection task, the deep learning models output the bounding boxes for the location of the target objects. We exploit mAP@50 for the evaluation metric due to its popular utilization in the detection tasks.

Model	MR-CNN	SSD	YOLOv5
ReLU	68.5 ± 0.02	70.1 ± 0.07	72.1 ± 0.04
LReLU	68.9 ± 0.06	70.6 ± 0.10	72.6 ± 0.03
PLeLU	69.4 ± 0.10	71.0 ± 0.11	73.1 ± 0.11
Softplus	69.4 ± 0.03	71.0 ± 0.08	73.0 ± 0.09
ELU	69.4 ± 0.10	71.1 ± 0.10	73.0 ± 0.06
SELU	69.7 ± 0.06	71.2 ± 0.04	73.3 ± 0.10
GELU	70.0 ± 0.08	71.6 ± 0.05	73.7 ± 0.04
Swish	70.4 ± 0.08	72.0 ± 0.03	74.0 ± 0.08
ASH	**71.1 ± 0.06**	**72.7 ± 0.02**	**74.8 ± 0.09**

Table 4. COCO with mean values and 95% C.I.

Model	MR-CNN	SSD	YOLOv5
ReLU	65.8 ± 0.03	67.3 ± 0.07	69.3 ± 0.09
LReLU	66.9 ± 0.07	68.4 ± 0.07	70.5 ± 0.07
PLeLU	67.8 ± 0.03	69.2 ± 0.07	71.2 ± 0.09
Softplus	67.9 ± 0.05	69.4 ± 0.03	71.4 ± 0.03
ELU	67.8 ± 0.05	69.3 ± 0.06	71.4 ± 0.06
SELU	68.4 ± 0.07	69.9 ± 0.03	72.1 ± 0.11
GELU	68.9 ± 0.07	70.5 ± 0.03	72.4 ± 0.11
Swish	69.0 ± 0.04	70.6 ± 0.02	72.6 ± 0.08
ASH	**70.5 ± 0.06**	**72.1 ± 0.04**	**74.1 ± 0.11**

Table 5. PASCAL VOC with mean values and 95% C.I.

The quantitative results in Tables 4 and 5 exhibit the outstanding performance of ASH activation function compared to other activation functions. Higher mAP indicates the predictive bounding boxes are significantly correct with the annotations, ASH activation function could provide superior performance in detecting target objects in various datasets while combining many deep learning models. Since the deep learning models for the detection task demands the locality to generate the bounding boxes, it is expected that the z_k could be small, and thus relatively much activation could be realized in ASH activation function compared to the models for the classification tasks.

3.3 Segmentation Task

Finally, we compare ASH to all baseline activation functions on the datasets of ADE20K (Zhou et al., 2017) and PASCAL VOC (Everingham et al., 2010) for the segmentation task. Both datasets include many target objects in one scene, and thus they are widely utilized in the benchmark of the segmentation task. We employ the same environments from the classification and detection tasks, and implement the baseline models of the U-Net (Ronneberger et al., 2015), DeepLabV3+(DLV3+) (Chen et al., 2018), and EfficientNet (Tan, Le, 2019). Like other general benchmarks, we exploit Intersection over Union (IoU) and mean IoU (mIoU) values as evaluation metrics due to their popular utilization in segmentation tasks.

Like previous tasks, the quantitative results in Tables 6 and 7 highlight the outstanding performance of ASH activation function compared to other activation functions. Since the locality is significantly important to segment target objects from the background in the segmentation task, we could intuitively understand that ASH activation function improves the locality during the feature extraction. The experimental results show that superior segmentation performance improvement could be realized using ASH activation function, which significantly helps to localize the target objects.

Model	U-Net	DLV3+	EfficientNet
ReLU	49.4 ± 0.04	50.7 ± 0.06	52.2 ± 0.03
LReLU	49.8 ± 0.07	51.0 ± 0.09	52.3 ± 0.05
PLeLU	49.9 ± 0.02	51.0 ± 0.03	52.5 ± 0.06
Softplus	50.1 ± 0.11	51.2 ± 0.10	52.8 ± 0.03
ELU	50.3 ± 0.04	51.4 ± 0.09	52.8 ± 0.08
SELU	50.9 ± 0.04	52.1 ± 0.03	53.5 ± 0.06
GELU	50.9 ± 0.07	52.2 ± 0.08	53.5 ± 0.05
Swish	51.3 ± 0.05	52.4 ± 0.02	53.9 ± 0.07
ASH	**53.4 ± 0.05**	**54.7 ± 0.08**	**56.3 ± 0.09**

Table 6. ADE20K with mean values and 95% C.I.

Model	U-Net	DLV3+	EfficientNet
ReLU	74.3 ± 0.05	76.0 ± 0.08	78.1 ± 0.07
LReLU	76.2 ± 0.06	78.0 ± 0.05	80.3 ± 0.10
PLeLU	77.0 ± 0.07	78.9 ± 0.06	81.0 ± 0.08
Softplus	77.1 ± 0.09	78.8 ± 0.03	81.2 ± 0.10
ELU	77.2 ± 0.10	79.0 ± 0.05	81.3 ± 0.02
SELU	78.2 ± 0.06	80.1 ± 0.03	82.4 ± 0.03
GELU	78.9 ± 0.06	80.7 ± 0.05	83.2 ± 0.08
Swish	78.8 ± 0.11	80.7 ± 0.04	82.9 ± 0.04
ASH	**81.2 ± 0.04**	**83.2 ± 0.04**	**85.5 ± 0.05**

Table 7. PASCAL VOC with mean values and 95% C.I.

3.4 Early Convergence and Optimization

We empirically explore the effectiveness of ASH activation function on the training time by monitoring the validation loss values for all of the activation functions. The result shows that using ASH activation function could improve the training time, in terms of early convergence from the early epoch with low loss values. Fig. 3 demonstrates that the loss values of ASH activation function exhibit a more steep slope compared to those of other activation functions. Therefore, since the ASH activation function reaches a convergence value significantly faster than other activation functions, we could empirically achieve the results that ASH has a superior effect in reducing the training time.

Through the experiments, we explore the outstanding performance of ASH activation function compared to other activation functions, including the improvement in accuracy, sparsity, training time and localization property. Despite the lack of proof of those properties mathematically, the experimental results demonstrate the outstanding performance of ASH activation function. The supporting mathematical analysis and proofs could be more discussed as potential future work. Additionally, the supporting experiments and the results of other tasks of image generation or language translation are exhibited in the *Supplementary Material*.

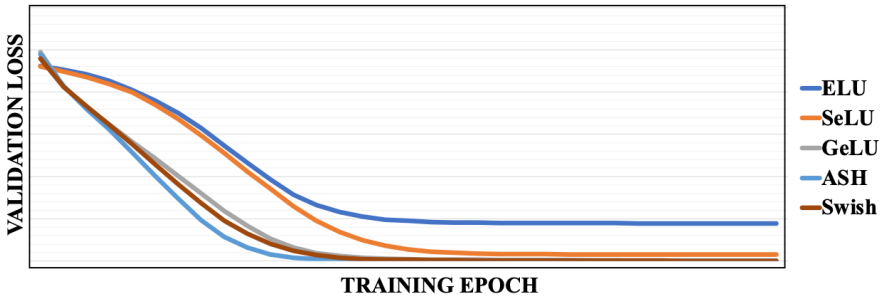

Figure 3: Validation loss value alongside the training epoch for the activation functions. The validation losses are averaged from the results of all experiments. Here, x-axis indicates the percentage of training epoch ($0\% \sim 100\%$), and they were averaged. In addition, the y-axis indicates the range of (0, 0.8).

4 Conclusion

In this paper, we devised a novel activation function to rectify inputs using an adaptive threshold considering entire contexts of inputs more like human neurons. To this end, we designed the activation function to extract elements in top-$k\%$ percentile from the input feature-map. Since sorting algorithm-based selections or quick selection algorithm demands a heavy computational cost, we employed the stochastic technique utilizing normal distribution to realize stochastic percentile sampling. From the mathematical derivation, we implemented ASH activation function in simple yet effective formula ($f(x) = x \cdot \text{sigmoid}(ax + b)$) with low computational cost in sampling top-$k\%$ percentile from the input. In addition, we implemented ASH activation function, realizing (1) the adaptive threshold by employing the z-score-based trainable variables and (2) the perception of entire contexts in rectifying an input by utilizing the mean and standard deviation of the input. Meanwhile, ASH activation function represented the generalized form of the Swish activation function that was empirically searched in the previous study. Therefore, this paper exhibited the additional novel contribution of the mathematical proofs for the state-of-the-art performance of the Swish activation function. The experiments using many deep learning models for various tasks (classification, detection, and segmentation) exhibited superior improvements in the accuracy, the localization, and training time to be converged while employing ASH activation function.

References

Aggarwal Charu C, others . Neural networks and deep learning // Springer. 2018. 10. 978–3.

Arjovsky Martin, Shah Amar, Bengio Yoshua. Unitary evolution recurrent neural networks // International conference on machine learning. 2016. 1120–1128.

Badrinarayanan Vijay, Kendall Alex, Cipolla Roberto. Segnet: A deep convolutional encoder-decoder architecture for image segmentation // IEEE transactions on pattern analysis and machine intelligence. 2017. 39, 12. 2481–2495.

Bingham Garrett, Miikkulainen Risto. Discovering parametric activation functions // Neural Networks. 2022.

Bochkovskiy Alexey, Wang Chien-Yao, Liao Hong-Yuan Mark. Yolov4: Optimal speed and accuracy of object detection // arXiv preprint arXiv:2004.10934. 2020.

Bottou Léon. Large-scale machine learning with stochastic gradient descent // Proceedings of COMPSTAT'2010. 2010. 177–186.

Chen Liang-Chieh, Zhu Yukun, Papandreou George, Schroff Florian, Adam Hartwig. Encoder-decoder with atrous separable convolution for semantic image segmentation // Proceedings of the European conference on computer vision (ECCV). 2018. 801–818.

Clevert Djork-Arné, Unterthiner Thomas, Hochreiter Sepp. Fast and accurate deep network learning by exponential linear units (elus) // arXiv preprint arXiv:1511.07289. 2015.

Evans Dominic A, Stempel A Vanessa, Vale Ruben, Ruehle Sabine, Lefler Yaara, Branco Tiago. A synaptic threshold mechanism for computing escape decisions // Nature. 2018. 558, 7711. 590–594.

Everingham Mark, Van Gool Luc, Williams Christopher KI, Winn John, Zisserman Andrew. The pascal visual object classes (voc) challenge // International journal of computer vision. 2010. 88, 2. 303–338.

Fukushima Kunihiko, Miyake Sei. Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition // Competition and cooperation in neural nets. 1982. 267–285.

Goodfellow Ian, Bengio Yoshua, Courville Aaron. Deep learning (adaptive computation and machine learning series) // Cambridge Massachusetts. 2017. 321–359.

Goodfellow Ian, Pouget-Abadie Jean, Mirza Mehdi, Xu Bing, Warde-Farley David, Ozair Sherjil, Courville Aaron, Bengio Yoshua. Generative adversarial nets // Advances in neural information processing systems. 2014. 27.

Hayou Soufiane, Doucet Arnaud, Rousseau Judith. On the impact of the activation function on deep neural networks training // International conference on machine learning. 2019. 2672–2680.

He Kaiming, Gkioxari Georgia, Dollár Piotr, Girshick Ross. Mask r-cnn // Proceedings of the IEEE international conference on computer vision. 2017. 2961–2969.

He Kaiming, Zhang Xiangyu, Ren Shaoqing, Sun Jian. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification // Proceedings of the IEEE international conference on computer vision. 2015. 1026–1034.

He Kaiming, Zhang Xiangyu, Ren Shaoqing, Sun Jian. Identity mappings in deep residual networks // European conference on computer vision. 2016. 630–645.

Hendrycks Dan, Gimpel Kevin. Gaussian error linear units (gelus) // arXiv preprint arXiv:1606.08415. 2016.

Hertz John, Krogh Anders, Lautrup Benny, Lehmann Torsten. Nonlinear backpropagation: doing backpropagation without derivatives of the activation function // IEEE Transactions on neural networks. 1997. 8, 6. 1321–1327.

Hoare Charles AR. Algorithm 65: find // Communications of the ACM. 1961. 4, 7. 321–322.

Hochreiter Sepp. Recurrent neural net learning and vanishing gradient // International Journal Of Uncertainity, Fuzziness and Knowledge-Based Systems. 1998. 6, 2. 107–116.

Hu Zheng, Li Yongping, Yang Zhiyong. Improving convolutional neural network using pseudo derivative ReLU // 2018 5th International Conference on Systems and Informatics (ICSAI). 2018. 283–287.

Huang Gao, Liu Zhuang, Van Der Maaten Laurens, Weinberger Kilian Q. Densely connected convolutional networks // Proceedings of the IEEE conference on computer vision and pattern recognition. 2017. 4700–4708.

Iliev A., Kyurkchiev N., Markov S. On the approximation of the step function by some sigmoid functions // Mathematics and Computers in Simulation. 2017. 133. 223–234. Biomath 2014 and Biomath 2015.

Ioffe Sergey, Szegedy Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift // International conference on machine learning. 2015. 448–456.

Izhikevich Eugene M. Simple model of spiking neurons // IEEE Transactions on neural networks. 2003. 14, 6. 1569–1572.

Jarrett Kevin, Kavukcuoglu Koray, Ranzato Marc'Aurelio, LeCun Yann. What is the best multi-stage architecture for object recognition? // 2009 IEEE 12th international conference on computer vision. 2009. 2146–2153.

Jinsakul Natinai, Tsai Cheng-Fa, Tsai Chia-En, Wu Pensee. Enhancement of deep learning in image classification performance using xception with the swish activation function for colorectal polyp preliminary screening // Mathematics. 2019. 7, 12. 1170.

Klambauer Günter, Unterthiner Thomas, Mayr Andreas, Hochreiter Sepp. Self-normalizing neural networks // Advances in neural information processing systems. 2017. 30.

Learning multiple layers of features from tiny images. // . 2009.

Larsen Richard J, Marx Morris L. An introduction to mathematical statistics. 2005.

Lee June-Goo, Jun Sanghoon, Cho Young-Won, Lee Hyunna, Kim Guk Bae, Seo Joon Beom, Kim Namkug. Deep learning in medical imaging: general overview // Korean journal of radiology. 2017. 18, 4. 570–584.

Leshno Moshe, Lin Vladimir Ya, Pinkus Allan, Schocken Shimon. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function // Neural networks. 1993. 6, 6. 861–867.

Lin Tsung-Yi, Maire Michael, Belongie Serge, Hays James, Perona Pietro, Ramanan Deva, Dollár Piotr, Zitnick C Lawrence. Microsoft coco: Common objects in context // European conference on computer vision. 2014. 740–755.

Lindsay Grace W. Convolutional neural networks as a model of the visual system: Past, present, and future // Journal of cognitive neuroscience. 2021. 33, 10. 2017–2031.

Liu Wei, Anguelov Dragomir, Erhan Dumitru, Szegedy Christian, Reed Scott, Fu Cheng-Yang, Berg Alexander C. Ssd: Single shot multibox detector // European conference on computer vision. 2016. 21–37.

Liu Yu, Chen Xun, Wang Zengfu, Wang Z Jane, Ward Rabab K, Wang Xuesong. Deep learning for pixel-level image fusion: Recent advances and future prospects // Information Fusion. 2018. 42. 158–173.

Maas Andrew L, Hannun Awni Y, Ng Andrew Y, others . Rectifier nonlinearities improve neural network acoustic models // Proc. icml. 30, 1. 2013. 3.

Misra Diganta. Mish: A self regularized non-monotonic activation function // arXiv preprint arXiv:1908.08681. 2019.

Molina Alejandro, Schramowski Patrick, Kersting Kristian. Pad\'e Activation Units: End-to-end Learning of Flexible Activation Functions in Deep Networks // arXiv preprint arXiv:1907.06732. 2019.

Nair Vinod, Hinton Geoffrey E. Rectified linear units improve restricted boltzmann machines // Icml. 2010.

Ramachandran Prajit, Zoph Barret, Le Quoc V. Searching for activation functions // arXiv preprint arXiv:1710.05941. 2017.

Ronneberger Olaf, Fischer Philipp, Brox Thomas. U-net: Convolutional networks for biomedical image segmentation // International Conference on Medical image computing and computer-assisted intervention. 2015. 234–241.

Rumelhart David E, Hinton Geoffrey E, Williams Ronald J. Learning representations by back-propagating errors // nature. 1986. 323, 6088. 533–536.

Russakovsky Olga, Deng Jia, Su Hao, Krause Jonathan, Satheesh Sanjeev, Ma Sean, Huang Zhiheng, Karpathy Andrej, Khosla Aditya, Bernstein Michael, others . Imagenet large scale visual recognition challenge // International journal of computer vision. 2015. 115, 3. 211–252.

Sharma Sagar, Sharma Simone, Athaiya Anidhya. Activation functions in neural networks // towards data science. 2017. 6, 12. 310–316.

Simonyan Karen, Zisserman Andrew. Very deep convolutional networks for large-scale image recognition // arXiv preprint arXiv:1409.1556. 2014.

Tan Mingxing, Le Quoc. Efficientnet: Rethinking model scaling for convolutional neural networks // International conference on machine learning. 2019. 6105–6114.

Ulyanov Dmitry, Vedaldi Andrea, Lempitsky Victor. Instance normalization: The missing ingredient for fast stylization // arXiv preprint arXiv:1607.08022. 2016.

Weisstein Eric W. Heaviside step function // https://mathworld. wolfram. com/. 2002.

Wu Yuxin, He Kaiming. Group normalization // Proceedings of the European conference on computer vision (ECCV). 2018. 3–19.

Xu Bing, Wang Naiyan, Chen Tianqi, Li Mu. Empirical evaluation of rectified activations in convolutional network // arXiv preprint arXiv:1505.00853. 2015.

Zagoruyko Sergey, Komodakis Nikos. Wide residual networks // arXiv preprint arXiv:1605.07146. 2016.

Zhang Zhifei. Derivation of backpropagation in convolutional neural network (cnn) // University of Tennessee, Knoxville, TN. 2016.

Zhao Zhong-Qiu, Zheng Peng, Xu Shou-tao, Wu Xindong. Object detection with deep learning: A review // IEEE transactions on neural networks and learning systems. 2019. 30, 11. 3212–3232.

Zhou Bolei, Zhao Hang, Puig Xavier, Fidler Sanja, Barriuso Adela, Torralba Antonio. Scene parsing through ade20k dataset // Proceedings of the IEEE conference on computer vision and pattern recognition. 2017. 633–641.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [N/A]
- Did you include the license to the code and datasets? [N/A]
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Introduction
   (b) Did you describe the limitations of your work? [Yes] See Section 2.1, Appendix
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2.1
   (b) Did you include complete proofs of all theoretical results? [Yes] See Section 2
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Abstract
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] statistical analysis were provided using confidence interval
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]