

# SYMMETRIC-APL ACTIVATIONS: TRAINING INSIGHTS AND ROBUSTNESS TO ADVERSARIAL ATTACKS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Deep neural networks with learnable activation functions have shown superior performance over deep neural networks with fixed activation functions for many different problems. The adaptability of learnable activation functions adds expressive power to the model which results in better performance. Here, we propose a new learnable activation function based on Adaptive Piecewise Linear units (APL), which 1) gives equal expressive power to both the positive and negative halves on the input space and 2) is able to approximate any zero-centered continuous non-linearity in a closed interval. We investigate how the shape of the Symmetric-APL function changes during training and perform ablation studies to gain insight into the reason behind these changes. We hypothesize that these activation functions go through two distinct stages: 1) adding gradient information and 2) adding expressive power. Finally, we show that the use of Symmetric-APL activations can significantly increase the robustness of deep neural networks to adversarial attacks. Our experiments on both black-box and open-box adversarial attacks show that commonly-used architectures, namely Lenet, Network-in-Network, and ResNet-18 can be up to 51% more resistant to adversarial fooling by only using the proposed activation functions instead of ReLUs.

## 1 INTRODUCTION

Deep neural networks (DNNs) are a series of linear transformations followed by point-wise non-linear functions. These non-linear functions are also referred to as activation functions. Activation functions directly influence the training dynamics and, therefore, affect the final performance of the DNN. For instance, activations such as Sigmoid and Tanh suffer from what is known as the “vanishing gradient” problem, when the inputs are in the saturation regions of these functions. This can adversely affect network performance and cause slow convergence. Rectified linear units (ReLU) (Nair & Hinton, 2010) have been shown to perform better than Sigmoid and Tanh in many scenarios. Instead of using a fixed activation function, one can learn a parameterized activation function to add expressive power to the model. Towards this goal, significant efforts have been made to introduce different learnable activation functions (Agostinelli et al., 2014; He et al., 2015; Ramachandran et al., 2017; Jin et al., 2016; Li et al., 2016). Although most of these adaptive activation functions outperform fixed activation functions in multiple tasks, there is very little insight into why.

In this paper, we propose the Symmetric Adaptive Piecewise Linear unit (Symmetric-APL; pronounced as Symmetric apple: 'æp.əl) as a new adaptive activation function which builds upon the adaptive piecewise linear unit (APL) (Agostinelli et al., 2014). Similar to APL units, S-APL has multiple hinges which are connected by linear pieces. The S-APL activation is designed to give the same expressive power to both the positive and negative sides of the x-axis. We show that S-APL can approximate any continuous and zero-centered function, which encompasses many types of learned and fixed activations functions. Using this flexible learned activation function, we can then analyze how the shape changes during training and use ablation studies to investigate the role that these changes play in optimization. This helps us better understand learned activation functions and explain why they often provide increases in performance.

Finally, we evaluate the robustness of S-APL activated networks to adversarial attacks. Experiments are done using both black-box and open-box adversarial attacks where we observe S-APL equipped

networks are more robust to attacks than their ReLU counterparts. Experiments show that the robustness improves against differential evolution attacks (Su et al., 2019b) by (up to) 51% on Lenet, 47% on ResNet-18 (He et al., 2016), 43% on pure CNN, and 38% on Net in Net (Lin et al., 2013). For open box adversarial attacks, the robustness also improves against FGSM (Goodfellow et al., 2014) attack by (up to) 36%, 28%, and 25% for Lenet, pure CNN, and ResNet-18 architectures, respectively using CIFAR-10 dataset.

## 2 RELATED WORKS

ReLU and its variants such as Leaky-ReLU (Maas et al., 2013), P-ReLU (He et al., 2015), exponential linear units (ELU) (Clevert et al., 2015), and scaled exponential linear units (SELU) (Klambauer et al., 2017) are believed to solve the vanishing gradient problem and some are shown to better normalize the network. By changing the negative side of x-axis, these variants push the mean activations of each unit closer to zero and accelerate learning.

Some of the pioneering attempts to learn activations in a neural network can be found in Poli (1996), Weingaertner et al. (2002), and Khan et al. (2013) where the authors proposed novel approaches to learning the best activation function per neuron among a pool of candidate activations by employing genetic and evolutionary algorithms. Maxout (Goodfellow et al., 2013) has been introduced as an activation function aimed at enhancing the model averaging properties of dropout (Srivastava et al., 2014). However, not only is it limited to approximating convex functions, but it also requires a significant increase in parameters. Neural architecture search (Ramachandran et al., 2017) uses a combination of exhaustive and reinforcement learning based search. Although the authors discovered multiple novel activation functions, of which  $f(x) = x * sigmoid(\beta x)$  performed slightly better than ReLU, the search space of activation functions is limited. S-ReLU (Jin et al., 2016) was proposed as an adaptive activation function which can mimic both convex and non-convex functions.

It is worth mentioning that, in Lin et al. (2013), the authors proposed the network in network approach where they replace activation functions in convolutional layers with small multi-layer perceptrons. Theoretically, due to Universal Approximation Theorem, this is the most expressive activation; however, it requires many more parameters.

## 3 ADAPTIVE PIECEWISE ACTIVATION FUNCTION

### 3.1 ADAPTIVE PIECEWISE LINEAR UNITS

Agostinelli et al. (2014) proposed an adaptive activation function which consists of multiple linear pieces. The formulation of this adaptive piecewise linear (APL) activation function is shown in Equation 1.

$$h_i(x) = \max(0, x) + \sum_{s=1}^S a_i^s \max(0, -x + b_i^s) \quad (1)$$

Where  $h_i$  is the non-linear transformation for the hidden unit  $i$  and  $S$  is a hyperparameter that determines the number of hinges.  $a_i$  is the slope of the linear piece and  $b_i$  is the location of the hinge.

### 3.2 SYMMETRIC ADAPTIVE PIECEWISE LINEAR UNITS (APL: 'æp.əl)

To overcome the shortcomings of APL activation function, we define a new adaptive piecewise linear activation function that gives the same expressive power to the positive and negative half of the input space and uses fewer parameters by sharing parameters across layers. We formulate the activation  $h_i$  of the hidden unit  $i$  as the summation of symmetric hinge-shaped pieces.

$$h_i(x, S) = \sum_{s=1}^S a^{s+} \max(0, x - b^{s+}) + \sum_{s=1}^S a^{s-} \max(0, -x - b^{s-}) \quad (2)$$

Similar to APL units,  $S$  is a hyperparameter which determines the number of hinges (on each side of the x-axis).  $a^{s^+}$  and  $a^{s^-}$  control the slope of linear pieces in both the positive and negative sides of the x-axis, while  $b^{s^+}, b^{s^-} \in [0, +\infty)$  determine the location of hinge  $s$ . The total number of additional parameters is  $4S$  for a single layer. Because our networks are trained with batch normalization (Ioffe & Szegedy, 2015), the distribution of input to the S-APL function will follow that of a normal distribution. Therefore, we fix the hinges to be a pre-determined number of standard deviations away from the mean, decreasing the number of parameters to  $2S$ .

Based on Equation 2, S-APL activation is zero-centered which leads to more stable statistical behavior during training. Due to increasing the number of parameters, S-APL offers more flexibility than the normal APL, however, in case of using a fixed set of hinges among all hidden units of a layer, this increase is of  $O(\text{constant})$  which is negligible. With the assumption of  $b_i^{s^+} = b_i^{s^-}$ , since  $h_i^s(x) = h_i^s(-x)$ , S-APL allows neurons to activate symmetrically. As it is mentioned in Zhao & Griffin (2016), networks with symmetric activation functions are more robust to adversarial fooling. We will delve into this property in later sections.

S-APL can better handle the limits. It could be inferred from the definition that, in contrast to normal APL, increasing the absolute value of the input  $x$  will also increase the number of terms affecting the  $h_i(x)$ . However, normal APL requires the rightmost linear section in all the component functions to have a unit slope and bias 0, which is not an appropriate constraint and undermines its representation ability.

The following theorem shows that S-APL can approximate any non-linear, zero-centered, and Lipschitz continuous function in a closed interval of real numbers. It is worth mentioning that APL was not flexible enough to do so.

**Theorem 3.1** *For any function  $f : [A, B] \rightarrow \mathbb{R}$  and  $\epsilon \in \mathbb{R}, \exists S \in \mathbb{N}$ , where  $|f(x) - S\text{-APL}(x)| \leq \epsilon$ , assuming:*

- $A$  and  $B$  are finite real numbers.
- $f$  is  $M$ -Lipschitz continuous.

**Proof:** First we define  $g : \mathbb{R} \rightarrow \mathbb{R}$  as follows:

$$g(x, S) = \sum_{i=1}^S \begin{cases} a_i x + b_i & A_i \leq x \leq B_i, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

It's clear that any piecewise linear function, including S-APL can be written in the form of  $g(x)$ . Now, if we divide  $[A, B]$  into  $S$  smaller windows  $[A_i, B_i]$ , due to the Lipschitz continuity for any  $x_1$  and  $x_2 \in [A_i, B_i]$ , we have:

$$|f(x_1) - f(x_2)| \leq \frac{M(B-A)}{S} \quad (4)$$

Due to the continuity of  $f$ ,  $\exists x_i^{\min}$  and  $x_i^{\max}$  such that  $x_i^{\min} = \text{argmin}(f(x))$  and  $x_i^{\max} = \text{argmax}(f(x))$  in  $[A_i, B_i]$ . Using  $x_i^{\min}$  and  $x_i^{\max}$  we can define  $g(x)$  as follows:

$$g(x) = \begin{cases} (x - A_i) \frac{f(x_i^{\max})}{B_i} \frac{f(x_i^{\min})}{A_i} + f(x_i^{\min}) & x_i^{\min} \leq x_i^{\max} \\ (x - A_i) \frac{f(x_i^{\min})}{B_i} \frac{f(x_i^{\max})}{A_i} + f(x_i^{\max}) & x_i^{\max} \leq x_i^{\min} \end{cases} \quad (5)$$

Using the defined  $g$  in equation 4, we have:  
 $\forall x \in [A, B] :$

$$|f(x) - g(x)| \leq \frac{M(B-A)}{S} \quad (6)$$

Thus, in order to satisfy  $|f(x) - g(x)| \leq \epsilon$ :

$$\frac{M(B - A)}{\epsilon} \leq S \tag{7}$$

The inequality 7 implies that if  $S$  is set to be greater than  $\frac{M(B - A)}{\epsilon}$  then the error of approximation would be lower than  $\epsilon$ .  $\square$

#### 4 COMPARISON TO COMMONLY USED ACTIVATION FUNCTIONS

In order to show that S-APL unit is beneficial for deep neural networks, we compare it with well-known activation functions in different architectures. We employed five different neural networks trained on three different datasets, MNIST, CIFAR-10, and CIFAR-100. We used  $S = 4$  and set the hinges at positions  $x = -2.5, -2, -1, 0, +1, +2, +2.5$  (0 considered as two hinges). Slope of  $a_1^{s+}$  is initialized to 1 and the remaining slopes are initialized to 0. With this initialization, the initialized S-APL mimics the shape of ReLU.

The results of the experiments are shown in Table 1. The table shows that S-APL leads to a better validation accuracy in almost all architectures.

Table 1: ReLU, leaky-ReLU, PReLU, tanh, sigmoid, ELU, maxout (nine features), swish:  $x \cdot \text{sigmoid}(\beta x)$  and S-APL activations are compared in three different tasks. The performance of S-APL in comparison to the best of the other activations. For the sake of brevity, D-A refers to Data Augmentation. The values in the tables are error-rates and are reported in percentage.

Activation	MNIST	CIFAR-10		CIFAR-100	
	-	-	D-A	-	D-A
Lenet5 (best: PReLU)	1.11	31.93	27.42	46.31	45.39
Lenet5 (S-APL)	<b>1.03</b>	30.83	27.02	45.74	44.81
MLP (best: swish)	1.55				
MLP (S-APL)	1.71				
pure CNN (best: maxout)		12.01	10.92	35.02	35.33
pure CNN (S-APL)		<b>11.81</b>	<b>10.32</b>	<b>34.42</b>	<b>34.10</b>
Net-in-Net(ReLU)		10.41	8.81	33.72	31.63
Net-in-Net(APL)		9.59	7.51	32.40	30.83
Net-in-Net(S-APL)		<b>9.41</b>	<b>7.20</b>	<b>32.12</b>	<b>30.33</b>

It is necessary to mention that pure CNN architecture consists of nine 2d convolutional layers, followed by one global average pooling prior to the output layer. The first seven convolution layers have kernels of size 3x3 with padded inputs, while the kernel size in the last two layers is reduced to 1x1 and the corresponding input is not padded. After every two convolution layers, a drop out layer with  $p = 0.25$  is applied.

#### 5 HOW S-APL CHANGES DURING TRAINING AND POSSIBLE REASONS WHY

Figures 1 and 2 show how the shape of S-APL changes for each layer during training. From these figures, we can see that during the early stages of training, the shape of S-APL looks similar to that of leaky ReLU: having a negative output for  $x < 0$  and a positive output for  $x \geq 0$ . During the later stages of training, the activation function has a positive output for both  $x < 0$  and  $x \geq 0$ . In addition, the slope of the output decreases as the magnitude of  $x$  increases.

We hypothesize that there are two distinct stages during the optimization of these learned activation functions. In the first stage, the function returns negative values for  $x < 0$  and in the second stage, the function returns positive values for  $x < 0$ . To investigate this hypothesis, we train S-APL under

two different conditions where we: 1) force the output to only be negative for  $x < 0$  (S-APL-negative) and 2) force the output to only be positive for  $x < 0$  (S-APL-positive). We accomplish this by taking the absolute value of all  $a^{s^-}$  and multiplying by  $+1$  or  $-1$ .

Figure 2 shows that, although S-APL-positive has the ability to mimic the final learned shape of the S-APL function, it barely deviates from the ReLU initialization. This shows that the ability to give a negative output for negative inputs is crucial for S-APLs. Figure 2 shows that S-APL-negative does initially go negative in the 4th layer and then returns to the ReLU shape. To gain further insight, we compare the loss of these activation functions in Figure 3. We see that S-APL has a lower loss compared to S-APL-positive and S-APL-negative. Furthermore, S-APL-negative has a lower loss compared to S-APL-positive which has a lower loss than ReLU.

These results lead us to believe that it is possible that the first stage, S-APL-negative, adds gradient information by giving negative output for negative inputs. It is possible that S-APL-positive is not capable of doing so because giving positive outputs for negative inputs creates an identifiability issue. In other words, outputs for  $x < 0$  and  $x \geq 0$  look the same. However, we can see that this property is useful for the S-APL function in the later stages of training. It is possible that this is because the DNN has adapted its parameters such that outputs on both the negative and positive sides capture a similar concept, adding expressive power to the model. Furthermore, one cannot achieve the final S-APL shape without the first stage of having negative outputs for negative inputs. To conclude, it appears that the S-APL units first go through two stages: 1) adding gradient information and 2) adding expressive power and that the second stage depends on the first stage. This information can be used to inform the design of learnable activation functions and even DNN optimization techniques.




Figure 1: Shape of S-APL activation during training a simple network of MLPs on MNIST dataset.

## 6 ROBUSTNESS TO ADVERSARIAL ATTACKS

DNNs have been shown to be vulnerable to many types of adversarial attacks (Szegedy et al., 2013; Goodfellow et al., 2014). Research suggests that activation functions are a major cause of this vulnerability (Zantedeschi et al., 2017). Zhang et al. (2018) bounds a given activation function with a few linear and quadratic functions and allows it to tackle general activation functions. This adds up with applying a different activation for each neuron so the resulting network shows efficiency and robustness to adversarial foolings. Wang et al. (2018) proposed a data-dependent activation function and empirically shows its robustness to both black-box and gradient-based adversarial attacks. Other studies such as Rakin et al. (2018), Dhillon et al. (2018), and Song et al. (2018) focused on other




Figure 2: Shape of S-APL during training a Lenet5 architecture on CIFAR-10 dataset.




Figure 3: Training loss trajectory for different S-APL initializations compared to fixed ReLU and leaky ReLU.

additional properties of activation such as quantization and pruning and showed they can improve the robustness of DNNs against adversarial examples.

Recently, Zhao & Griffin (2016) theoretically showed that DNNs with symmetric activations are less likely to get fooled. The authors proved that *“symmetric units suppress unusual signals of exceptional magnitude which result in robustness to adversarial fooling and higher expressibility.”* Due to the symmetric convergence of S-APL units and also their hinges’ locations ( $b_i^+ = b_i^-$ ) which result in  $h(x) = h(-x)$ , S-APL units are capable of increasing the robustness of DNNs against adversarial attacks.

In this section, relying on the properties of S-APL such as data dependency, piecewise linearity, non-injective behavior, better handling of extreme values, and most importantly, symmetric shape, we show that adding S-APL to a DNNs greatly improves the robustness against adversarial attacks. This claim is verified through a wide range of experiments with the CIFAR-10 dataset under both black-box and open-box methods, including one-pixel-attack and Fast Gradient Sign Method.

Figure 4 provides an intuition for the robustness of S-APL activated network in comparison to ReLU activated ones. For each of the two networks, we take 100 random samples of frog and ship images and visualize the pre-softmax representations using tSNE visualization (Maaten & Hinton, 2008) in Figure 4. As it is depicted, for S-APL activated networks, the two classes are better separated. We

suspect that red samples such as those that are lying on the outer circle in the left plot of Figure 4 might be somewhere on the manifold of the data where the network is more susceptible. As one can see, there is no red point on the outer circle of the right plot which might be showing the robustness of S-APL networks. For other samples on the inner circle and the cloud next to circles, we can see better separation by the S-APL network.




Figure 4: tSNE visualization of the pre-softmax layer’s outputs. Left: a ReLU activated Lenet5 architecture trained on CIFAR-10. Right: same network trained with S-APL activation functions.

### 6.1 BLACK BOX ADVERSARIAL ATTACK

For black-box attacks against a DNN, we assume the adversary has no information about the structure or parameters of the DNN and does not have access to any large training dataset. The adversary’s only capability is to observe labels assigned by the DNN for chosen inputs in a manner analog to a cryptographic oracle. A successful attacking technique of this type is Su et al. (2019a) which is based on differential evolution. We can iteratively generate adversarial images to try to minimize the confidence (probability) of the network’s classification.

The process starts with random modifications of a few pixels to generate adversarial examples. At each step, one feeds several adversarial images to the DNN and only observes the output as the probability of different classes. Examples that lowered the confidence of the true class will be kept to generate the next generation of adversaries. The next generation modifications are then evolved through a certain mutation scheme to generate new adversarial images. By repeating these steps for a few iterations, the adversarial modifications generate more and more misleading images. The last step returns an adversarial modification that reduces the confidence of the true class the most. This means that the confidence of the true class would be reduced so much that a new incorrect category now has the highest classification confidence. This technique is called untargeted attack since we don’t specify the labels of the modified images.

In the following experiment, we modify one, three, and five pixels of images to generate adversarial examples. The mutation scheme we used for this experiment is as follows:

$$x_i^{l+1} = x_{r_1}^l + 0.5(x_{r_2}^l + x_{r_3}^l) \quad (8)$$

Where  $r_1$ ,  $r_2$ , and  $r_3$  are three non-equal random indices of the modifications at level  $l$ .  $x_i^{l+1}$  will be an element of a new candidate modification.

To evaluate the effect of S-APL activation on the robustness of DNNs, we employ commonly-used architectures, namely, Lenet5, Network-in-Network (Lin et al., 2013), pure CNN (ours), and ResNet-18 (He et al., 2016). Each architecture is trained once with ReLU activation functions and once with S-APL activation functions. The results in Table 2 show that Lenet5 and ResNet-18 architectures are 51% and 47.8% more resistant to back-box adversarial attacks. As authors in Zhao & Griffin (2016) theoretically proved, that one can be robust to adversarial examples with symmetric behavior.

After observing adversarial samples which are deceiving to both networks, we found that S-APL activated network still assigns high confidence to the true labels of the perturbed images. In other words, the difference between the confidence of the true label and the confidence of adversarially wrong label is much smaller in S-APL network than ReLU network. More precisely, we measure

Table 2: 100 images are randomly chosen from CIFAR-10 test set to generate adversarial examples. We attacked each architecture five times and report the average number of successful attacks. The maximum number of iterations for all attacks is set to 40.

Model	Lenet5		Net in Net		pure CNN		ResNet18	
	ReLU	S-APL	ReLU	S-APL	ReLU	S-APL	ReLU	S-APL
one-pixel	61	<b>29</b>	29	18	16	9	23	<b>12</b>
three-pixels	88	58	69	39	53	34	63	37
five-pixels	92	67	80	58	64	41	71	49
$avg( Z_{\text{true}} - Z_{\text{adv}} )$	0.721	<b>0.401</b>	0.411	<b>0.281</b>	0.483	<b>0.243</b>	0.501	<b>0.334</b>

the average of  $|Z(x^\theta)_{\text{true label}} - Z(x^\theta)_{\text{adversarial label}}|$  over all adversarial samples where each network is fooled.  $Z(\cdot)$  is the post-softmax output and  $x^\theta$  is the adversarial sample. The  $avg(|Z(x^\theta)_{\text{true label}} - Z(x^\theta)_{\text{adversarial label}}|)$  for each model is also included in Table 2. Due to the reported averages, we can conclude that S-APL is much more robust against being adversarially deceived.




Figure 5: First row: samples of successful attacks (one pixel) to ReLU activated network. Second row: samples of successful attacks to S-APL activated network. The differences in confidence of the true label and adversarially classified label is shown for each sample.

## 6.2 OPEN BOX ADVERSARIAL ATTACK

To further explore the robustness of S-APL activated networks, in this section, we consider one of the popular benchmarks of open box adversarial attacks: Fast Gradient Sign Method (FGSM) (Goodfellow et al., 2014). For this attacking strategy, we consider four different architectures and compare the rate of successful attacks for each of the networks with both ReLU and S-APL activation functions. The dataset and architectures are same as those were used for black-box adversarial attacks.

### 6.2.1 FGSM

Fast Gradient Sign Method generates an adversarial image  $x^\theta$  from the original image  $x$  by maximizing the loss  $L(x^\theta, y)$ , where  $y$  is the true label of the image  $x$ . This maximization problem is subjected to  $\|x - x^\theta\|_1 \leq \epsilon$  where  $\epsilon$  is considered as the *attack strength*. Approximating with the first term of Taylor expansion, we have:



$$L(x^0, y) = L(x, y) + \nabla_x L(x, y)^T \cdot (x - x^0) \quad (9)$$

So the adversarial image  $x^0$  would be:

$$x^0 = x + \epsilon \cdot \text{sign}(\nabla_x L(x, \theta)) \quad (10)$$

Where the  $\epsilon$  is considered as the power of the attack. This form of FGSM is considered to be an untargeted attack where there is no pre-specified label  $t$  that  $x^0$  should be classified as. This method only yields an adversary  $x^0$  which will not be classified as  $y$ .

To see how S-APL increases the robustness against FGSM attack, we employ three architectures all trained to classify CIFAR-10 dataset. For different range of  $\epsilon$  the results are summarized in Table 3. In the best case, for Lenet architecture and  $\epsilon = 0.02$ , S-APL activated network is 36% more robust

Table 3: Similar to the black-box attack experiment, 100 images are randomly chosen from CIFAR-10 test set to generate adversarial examples. For each architecture, four different  $\epsilon$  is used to FGSM attack. All S-APL activations are set as **symmetric-shared S-APL** with four hinges on each side of the x-axis.

Model	Lenet5		pure CNN		ResNet18	
	ReLU	S-APL	ReLU	S-APL	ReLU	S-APL
$\epsilon = 0.02$	57	<b>36</b>	35	25	36	27
$\epsilon = 0.04$	64	44	36	28	38	29
$\epsilon = 0.06$	69	62	38	33	39	35
$\epsilon = 0.08$	73	68	42	38	43	40

than the ReLU activated one. The robustness also improves by 28% and 25% for pure CNN and ResNet-18 respectively. We can see the overall increase in the robustness over all networks and  $\epsilon$  values. However, one can conclude that for smaller values of  $\epsilon$  the robustness of S-APL activated network is higher in comparison to ReLU activated networks. As we increase the attack strength ( $\epsilon$ ), the robustness of S-APL networks gets closer to the robustness of ReLU activated networks.

## 7 CONCLUSION AND FUTURE WORKS

We extend the idea of piecewise linear activation functions by introducing S-APL as a learnable and potentially symmetric activation function. We investigate the properties of S-APL and their evolution during training to understand how these functions accelerate learning and increase expressive power in DNNs. Finally, we show that S-APL networks show more robustness to adversarial attacks than ReLU networks. One can potentially design a hybrid defense method that combines other defense techniques and S-APL units.

## REFERENCES

- Forest Agostinelli, Matthew Hoffman, Peter Sadowski, and Pierre Baldi. Learning activation functions to improve deep neural networks, 2014.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus), 2015.
- Guneet S. Dhillon, Kamyar Azizzadenesheli, Zachary C. Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense, 2018.
- Ian J Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Maxout networks. *arXiv preprint arXiv:1302.4389*, 2013.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec 2015. doi: 10.1109/iccv.2015.123. URL <http://dx.doi.org/10.1109/ICCV.2015.123>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Xiaojie Jin, Chunyan Xu, Jiashi Feng, Yunchao Wei, Junjun Xiong, and Shuicheng Yan. Deep learning with s-shaped rectified linear activation units. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Maryam Mahsal Khan, Arbab Masood Ahmad, Gul Muhammad Khan, and Julian F Miller. Fast learning neural networks using cartesian genetic programming. *Neurocomputing*, 121:274–289, 2013.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Advances in neural information processing systems*, pp. 971–980, 2017.
- Hongyang Li, Wanli Ouyang, and Xiaogang Wang. Multi-bias non-linear activation in deep neural networks. In *International conference on machine learning*, pp. 221–229, 2016.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, pp. 3, 2013.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807–814, 2010.
- Riccardo Poli. *Parallel distributed genetic programming*. University of Birmingham, Cognitive Science Research Centre, 1996.
- Adnan Siraj Rakin, Jinfeng Yi, Boqing Gong, and Deliang Fan. Defend deep neural networks against adversarial examples via fixed and dynamic quantized activation functions, 2018.
- Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- Sibo Song, Yueru Chen, Ngai-Man Cheung, and C. C. Jay Kuo. Defense against adversarial attacks with saak transform, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, pp. 1–1, 2019a. ISSN 1941-0026. doi: 10.1109/tevc.2019.2890858. URL <http://dx.doi.org/10.1109/TEVC.2019.2890858>.
- Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 2019b.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

- Bao Wang, Alex T. Lin, Zuoqiang Shi, Wei Zhu, Penghang Yin, Andrea L. Bertozzi, and Stanley J. Osher. Adversarial defense via data dependent activation function and total variation minimization, 2018.
- Daniel Weingaertner, Victor K Tatai, Ricardo R Gudwin, and Fernando J Von Zuben. Hierarchical evolution of heterogeneous neural networks. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No. 02TH8600)*, volume 2, pp. 1775–1780. IEEE, 2002.
- Valentina Zantedeschi, Maria-Irina Nicolae, and Ambrish Rawat. Efficient defenses against adversarial attacks. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 39–49. ACM, 2017.
- Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions, 2018.
- Qiyang Zhao and Lewis D Griffin. Suppressing the unusual: towards robust cnns using symmetric activation functions, 2016.