

Deep Learning with S-Shaped Rectified Linear Activation Units

Xiaojie Jin¹ Chunyan Xu⁴ Jiashi Feng² Yunchao Wei² Junjun Xiong³ Shuicheng Yan²

¹NUS Graduate School for Integrative Science and Engineering, NUS

²Department of ECE, NUS ³Beijing Samsung Telecom R&D Center

⁴School of CSE, Nanjing University of Science and Technology

xiaojie.jin@nus.edu.sg, xuchunyan01@gmail.com, elefjia@nus.edu.sg,
wychao1987@gmail.com, jjun.xiong@samsung.com, eleyans@nus.edu.sg

Abstract

Rectified linear activation units are important components for state-of-the-art deep convolutional networks. In this paper, we propose a novel S-shaped rectified linear activation unit (SReLU) to learn both convex and non-convex functions, imitating the multiple function forms given by the two fundamental laws, namely the Webner-Fechner law and the Stevens law, in psychophysics and neural sciences. Specifically, SReLU consists of three piecewise linear functions, which are formulated by four learnable parameters. The SReLU is learned jointly with the training of the whole deep network through back propagation. During the training phase, to initialize SReLU in different layers, we propose a “freezing” method to degenerate SReLU into a predefined leaky rectified linear unit in the initial several training epochs and then adaptively learn the good initial values. SReLU can be universally used in the existing deep networks with negligible additional parameters and computation cost. Experiments with two popular CNN architectures, Network in Network and GoogLeNet on scale-various benchmarks including CIFAR10, CIFAR100, MNIST and ImageNet demonstrate that SReLU achieves remarkable improvement compared to other activation functions.

Introduction

Convolutional neural networks (CNNs) have made great progress in various fields, such as object classification (Krizhevsky, Sutskever, and Hinton 2012), detection (Girshick et al. 2013) and character recognition (Cireşan et al. 2011). One of the key factors contributing to the success of the modern deep learning models is using the non-saturated activation function (*e.g.* ReLU) to replace its saturated counterpart (*e.g.* sigmoid and tanh), which not only solves the problem of “exploding/vanishing gradient” but also makes the deep networks converge fast. Among all the proposed non-saturated activation functions, the Rectified Linear Unit (ReLU) (Nair and Hinton 2010) is widely viewed as one of the several reasons for the remarkable performance of deep networks (Krizhevsky, Sutskever, and Hinton 2012).

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

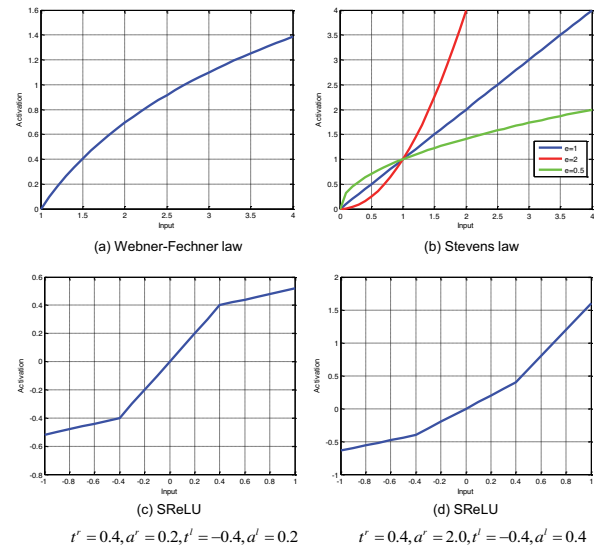


Figure 1: The function forms of Webner-Fechner law and Stevens law along with the proposed SReLU. (a) shows the logarithm function. (b) shows the power function with different exponents. (c) and (d) are different forms of SReLU by changing the parameters. The positive part of (c) and (d) are derived by imitating the logarithm function (a) and power function (b), respectively.

Recently, there are some other activation functions proposed to boost the performance of CNNs. Leaky ReLU (LReLU) (Maas, Hannun, and Ng 2013) assigns the negative part with a non-zero slope. (He et al. 2015) proposed the parametric rectified linear unit (PReLU), which requires learning the negative part instead of using predefined values. Adaptive piecewise linear activation (APL) proposed in (Agostinelli et al. 2014) sums up several hinge-shared linear functions. (Goodfellow et al. 2013) proposed the “max-out” activation function, which approximates arbitrary convex functions by computing the maximum of k linear functions for each neuron as the output.

Although the activation functions mentioned above have been reported to achieve good performance in CNNs, they all suffer from a weaknesses, *i.e.*, their limited ability

to learn non-linear transformation. For example, none of ReLU, LReLU, PReLU and maxout can learn the non-convex functions since they are essentially all convex functions. Although APL can approximate non-convex function, it requires the rightmost linear function in all the component functions to have a unit slope and bias 0, which is an inappropriate constraint and undermines its representation ability.

Inspired by the fundamental Webner-Fechner law (Weber 1851) and Stevens law (Stevens 1957) in psychophysics and neural sciences, we propose a novel kind of activation unit, namely the S-shaped rectified linear unit (SReLU). Two examples of SReLU’s function forms are shown in Figure 1(c)(d). Briefly speaking, both the Webner-Fechner law and the Stevens law describe the relationship between the magnitude of a physical stimulus and its perceived intensity or strength (Johnson, Hsiao, and Yoshioka 2002). The Webner-Fechner law holds that the perceived magnitude s is a logarithmic function of the stimulus intensity p multiplied by a modality and a dimension specific constant k . That is,

$$s = k \log p. \quad (1)$$

And the Stevens law explains the relationship through a power function, *i.e.*,

$$s = kp^e, \quad (2)$$

where all the parameters have the same definitions as in the Webner-Fechner law, except for an additional parameter e which is an exponent depending on the type of the stimulus. The function forms proposed by the two laws are shown in Figure 1(a)(b). These laws are usually valid for general sensory phenomena and can account for many properties of sensory neurons (Randall et al. 2002). More detailed discussions will be presented in Related Work.

Roughly, SReLU consists of three piecewise linear functions constrained by four learnable parameters as shown in Eqn. (7). The usage of SReLU brings two advantages to the deep network. Firstly, SReLU can learn both convex and non-convex functions, without imposing any constraints on its learnable parameters, thus the deep network with SReLU has a stronger feature learning capability. Secondly, since SReLU utilizes piecewise linear functions rather than saturated functions, thus it shares the same advantages of the non-saturated activation functions: it does not suffer from the “exploding/vanishing gradient” problem and has a high computational speed during the forward and back-propagation of deep networks. To verify the effectiveness of SReLU, we test it with two popular deep architectures, Network in Network and GoogLeNet, on four datasets with different scales, including CIFAR10, CIFAR100 and MNIST and ImageNet. The experimental results have shown remarkable improvement over other activation functions.

Related Work

In this section, we first review some activation units including ReLU, LReLU, PReLU, APL and maxout. Then we introduce two basic laws in psychophysics and neural sciences: Webner-Fechner law (Fechner 1965) and Stevens law (Stevens 1957), as well as our motivation.




Figure 2: Piecewise linear activation functions: ReLU, LReLU, PReLU, APL and maxout.

Rectified Units

- Rectified Linear Unit (ReLU) and Its Generalizations
ReLU (Nair and Hinton 2010) is defined as

$$h(x_i) = \max(0, x_i) \quad (3)$$

where x_i is the input and $h(x_i)$ is the output. LReLU (Maas, Hannun, and Ng 2013) assigns a slope to its negative input. It is defined as

$$h(x_i) = \min(0, a_i x_i) + \max(0, x_i) \quad (4)$$

where $a_i \in (0, 1)$ is a predefined slope. PReLU is only different from LReLU in that the former needs to learn the slope parameter a_i via back propagation during the training phase.

- Adaptive Piecewise Linear Units (APL)
APL is defined as a sum of hinge-shared functions:

$$h(x_i) = \max(0, x_i) + \sum_s^S a_i^s \max(0, -x_i + b_i^s), \quad (5)$$

where S is the number of hinges, and the variables $a_i^s, b_i^s, i \in 1, \dots, S$ are parameters of linear functions.

One disadvantage of APL is it explicitly forces the rightmost line to have unit slope 1 and bias 0. Although it is stated that if the output of APL serves as the input to a linear function $wh(x_i) + z$, the linear function will restore the freedom of the rightmost line which is lost due to the constraint, we argue that this does not always hold because in many cases for deep networks, the function taking the output of APL as the input is non-linear or unrestorable, such as local response normalization (Krizhevsky, Sutskever, and Hinton 2012) and dropout (Krizhevsky, Sutskever, and Hinton 2012).

- Maxout Unit

Maxout unit takes as the input the output of multiple linear functions and returns the largest:

$$h(x_i) = \max_{k \in \{1, \dots, K\}} w^k \cdot x_i + b^k. \quad (6)$$

In theory, maxout can approximate any convex function (Goodfellow et al. 2013), but unfortunately, it lacks the ability to learn non-convex functions. Moreover, a large number of extra parameters introduced by the K linear functions of each hidden maxout unit result in large storage memory cost and considerable training time, which affect the training efficiency of very deep CNNs, e.g. GoogLeNet (Szegedy et al. 2014).

Basic Laws in Psychophysics and Neural Sciences

The Webner-Fechner law (Fechner 1965) and the Stevens law (Stevens 1957) are two basic laws in psychophysics (Randall et al. 2002) (Johnson, Hsiao, and Yoshioka 2002) and neural sciences (Dayan and Abbott 2001). Webner first observed through experiments that the amount of change needed for sensory detection to occur increases with the initial intensity of a stimulus, and is proportional to it (Weber 1851). Based on Webner’s work, Fechner proposed the Webner-Fechner law which developed the theory by stating that the subjective sense of intensity is related to the physical intensity of a stimulus by a logarithmic function, which is formulated as Eqn. (1) and shown in Figure 1(a). Stevens refuted the Webner-Fechner law by arguing that the subjective intensity is related to the physical intensity of a stimulus by a power function (Johnson, Hsiao, and Yoshioka 2002), which is formulated as Eqn. (2) and shown in Figure 1(b). The two laws have been verified through lots of experiments (Nieder 2005). For example, in vision, the amount of change in brightness with respect to the present brightness accords with the Webner-Fechner law, *i.e.*, Eqn. (2). (Stevens 1957) shows various examples, one of which is that the perception of pain and taste follows the Stevens law but with different exponent values. In neural sciences, these two laws also explain many properties of sensory neurons and the response characteristics of receptor cells (Nieder 2005) (Dayan and Abbott 2001). The more detailed discussion is beyond the range of this paper.

Motivated by the previous research on these two laws, we propose SReLU which imitates the logarithm function and the power function given by the Webner-Fechner law and the Stevens law, respectively, and uses piecewise linear functions to approximate non-linear convex and non-convex functions. Through experiments, we find that our method can be universally used for current deep networks, and significantly boosts the performance.

S-shaped Rectified Linear Units (SReLU)

In this section, we introduce in detail our proposed SReLU. Firstly, we present the definition and the training process of SReLU. Secondly, we propose a method to initialize the parameters of SReLU as a good starting point for training.

Finally, we discuss the relationship of SReLU with other activation functions.

Definition of SReLU

SReLU is essentially defined as a combination of three linear functions, which perform mapping $\mathbb{R} \rightarrow \mathbb{R}$ with the following formulation:

$$h(x_i) = \begin{cases} t_i^r + a_i^r(x_i - t_i^r), & x_i \geq t_i^r \\ x_i, & t_i^r > x_i > t_i^l \\ t_i^l + a_i^l(x_i - t_i^l), & x_i \leq t_i^l \end{cases} \quad (7)$$

where $\{t_i^r, a_i^r, t_i^l, a_i^l\}$ are four learnable parameters used to model an individual SReLU activation unit. The subscript i indicates that we allow SReLU to vary in different channels. As shown in Figure 1(c)(d), in the positive direction, a_i^r is the slope of the right line when the inputs exceed the threshold t_i^r . Symmetrically, t_i^l is used to represent another threshold in the negative direction. When the inputs are smaller than t_i^l , the outputs are calculated by the left line. When the inputs of SReLU fall into the range of (t_i^l, t_i^r) , the outputs are linear functions with unit slope 1 and bias 0.

By designing SReLU in this way, we hope that it can imitate the formulations of multiple non-linear functions, including the logarithm function (Eqn. (1)) and the power function (Eqn. (2)) given by the Webner-Fechner law and the Stevens law, respectively. As shown in Figure 1(c), when $t_i^r > 1, a_i^r > 0$, the positive part of SReLU imitates the power function with the exponent e larger than 1; when $1 > t_i^r > 0, a_i^r > 0$, the positive part of SReLU imitates the logarithm function; when $t_i^r = 1, a_i^r > 0$, SReLU follows the power function with the exponent 1. For the negative part of SReLU, we have a similar observation except for the inverse representation of the logarithm function and the power function as analyzed for its positive counterpart. The reason for setting the middle line to be a linear function with slope 1 and bias 0 when the input is within the range (t_i^l, t_i^r) is that it can better approximate both Eqn. (1) and Eqn. (2) using such a function, because the change of the outputs with respect to the inputs is slow when the inputs are in small magnitudes.

Unlike APL which restricts the form of rightmost line, we do not apply any constraints or regularization to the parameters, thus both the thresh parameters and slope parameters can be learned freely as the training goes on. It is noteworthy that no divergence of deep networks occurs although SReLU is allowed to be trained without any constraints in all of our experiments. As shown in Table 3, the learned parameters are all in reasonable condition.

In our method, we learn an independent SReLU following each channel of kernels. Thus the number of the parameters for SReLU in the deep networks is only $4N$, where N is the overall number of kernel channels in the whole network. Compared with the large number of parameters in CNNs, e.g. 5 million parameters in GoogLeNet (Szegedy et al. 2014), such an increase in the number of parameters (21.7K in GoogLeNet with SReLU, as shown in Table 5) is negligible. This is a good property of SReLU, because on one hand, we

Table 1: Comparison of error rates between the channel-shared variant and the channel-wise variant of SReLU on CIFAR-10 without data augmentation.

Model	Error Rates
NIN + ReLU [Lin et. al.]	10.43%
NIN + SReLU (channel-shared)	9.01%
NIN + SReLU (channel-wise)	8.41%

avoid the overfitting effectively by increasing only a negligible number of parameters, and on the other hand we keep the memory size and the computing time almost unchanged. Similar to PReLU (He et al. 2015), we also try the channel-shared variant of SReLU. In this case the number of SReLU is equal to the overall number of layers in the deep network. In Tabel 1, we compare the performance of these two variants of SReLU on CIFAR-10 without data augmentation and find that the channel-wise version performs slightly better than the channel-shared version.

With respect to the training of SReLU, we use the gradient descent algorithm and jointly train the parameters of SReLU with the deep networks. The update rule of $\{t_r, a_r, t_l, a_l\}$ is derived by the chain rule:

$$\frac{\partial L}{\partial t_i^r} = \sum_{x_i} \frac{\partial L}{\partial h(x_i)} \frac{\partial h(x_i)}{\partial o_i} \quad (8)$$

where $o_i \in \{t_i^r, a_i^r, t_i^l, a_i^l\}$ and L represents the objective function of the deep network. The term $\frac{\partial L}{\partial h(x_i)}$ is the gradient back-propagated from the higher layer of SReLU. The summation \sum_{x_i} is applied in all positions of the feature map. For the channel-shared variant, the gradient of o_i is $\frac{\partial L}{\partial o_i} = \sum_i \sum_{x_i} \frac{\partial L}{\partial h(x_i)} \frac{\partial h(x_i)}{\partial o_i}$, where \sum_{x_i} is the sum over all channels in each layer. Specifically, the gradient for each parameter of SReLU is given by

$$\begin{aligned} \frac{\partial h(x_i)}{\partial t_i^r} &= I\{x_i \geq t_i^r\}(1 - a_i^r) \\ \frac{\partial h(x_i)}{\partial a_i^r} &= I\{x_i \geq t_i^r\}(x_i - t_i^r) \\ \frac{\partial h(x_i)}{\partial t_i^l} &= I\{x_i \leq t_i^l\}(1 - a_i^l) \\ \frac{\partial h(x_i)}{\partial a_i^l} &= I\{x_i \leq t_i^l\}(x_i - t_i^l) \end{aligned} \quad (9)$$

where $I\{\cdot\}$ is an indicator function and $I\{\cdot\} = 1$ when the expression inside holds true, otherwise $I\{\cdot\} = 0$. By this way, the gradient of the input is

$$h(x_i) = \begin{cases} a_i^r, & x_i \geq t_i^r \\ 1, & t_i^r > x_i > t_i^l \\ a_i^l, & x_i \leq t_i^l. \end{cases} \quad (10)$$

The rule for updating o_i by momentum method is:

$$\Delta o_i := \mu \Delta o_i + \varepsilon \frac{\partial L}{\partial o_i}. \quad (11)$$

Here μ is the momentum and ε is the learning rate. Because the weight decay term tends to pull the parameters to zero, we do not use weight decay (l_2 regularization) for o_i .




Figure 3: The distribution of the magnitude of the input to SReLU following convolution layers in GoogLeNet. The indexes of convolution layers follow a low-level to high-level order. The magnitudes shown here are calculated by averaging the activations of all SReLU in each layer.

Adaptive Initialization of SReLU

One problem we are faced in training SReLU is how to initialize the parameters in SReLU. An intuitive way is to set the parameters manually. However, such an initialization method is cumbersome. Furthermore, if the manually set initialization values are not appropriate, *e.g.* too large or too small compared with the real value of its input, SReLU may not work well. For example, if t_i^r is set to be very large, based on Eqn. (9), nearly all the inputs for the SReLU will lie in the left part of t_i^r , which will cause t_i^r and a_i^r to be insufficiently learned. In current deep networks, the magnitude of the inputs in each layer varies a lot (see Figure 3), making it more difficult to manually set parameters. To deal with this problem, we propose to firstly initialize each o_i to be $\{\tilde{t}_i, 1, 0, \tilde{a}_i\}$ in all layers, where \tilde{t}_i is any positive real number and $\tilde{a}_i \in (0, 1)$, and we “freeze” the update of the parameters of SReLU during the initial several training epochs. By this method, SReLU is degenerated into a conventional LReLU at the beginning of the training. Then upon the end of the “freezing” phase, we set t_i^r to be the largest k^{th} value of each SReLU’s input from all training data, *i.e.*,

$$t_i^r = \text{supp}(\mathbf{X}_i, k) \quad (12)$$

where $\text{supp}(\mathbf{X}, k)$ calculates the k^{th} largest value from the set \mathbf{X} , and \mathbf{X}_i represents all the input values of an individual SReLU. Our initialization method offers following two advantages. Firstly, it learns adaptively the initial values of t_i to fit better to the real distributions of the training data, thus providing a good starting point for the training of SReLU. Secondly, it enables SReLU to re-use the per-trained model with LReLU, thus it can reduce the training time compared with training the whole network from the scratch.

Comparison with Other Activation Functions

In this part, we compare our method with five published non-linear activation functions: ReLU, LReLU, PReLU, APL and maxout.

By checking Eqn. (3), Eqn. (4) and Eqn. (7), it can be easily concluded that ReLU, LReLU and PReLU can be

seen as special cases of SReLU. Specifically, when $t_i^r \geq 0$, $a_i^r = 1$, $t_i^l = 0$, $a_i^l = 0$, SReLU is degenerated into ReLU; when $t_i^r \geq 0$, $a_i^r = 1$, $t_i^l = 0$, $a_i^l > 0$, SReLU is transformed to LReLU and PReLU. However, ReLU, LReLU and PReLU can only approximate convex functions, while SReLU is able to approximate both convex and non-convex functions. Compared with APL, when the inputs have large magnitudes and lie in the rightmost region of the activate function, SReLU allows its parameters to take more flexible values and gives output features with adaptive scaling over the inputs. This is similar to the Webner-Fechner law that has logarithm function form to suppress the outputs for the input with too large magnitude. SReLU models such suppression effect by learning the slope of its rightmost line adaptively. In contrast, APL constrains the output to be same as input even when the inputs have very large magnitudes. This is the key difference between SReLU and APL and also the main reason why SReLU consistently outperforms APL. The experimental results shown in Table 2 clearly demonstrate this point. Without data augmentation and the proposed initialization strategy, NIN + SReLU outperforms NIN + APL by 0.98% and 3.04% on CIFAR-10 and CIFAR-100, respectively. Compared to maxout, which can only approximate convex functions and introduces a large number of extra parameters, SReLU needs much less parameters, therefore SReLU is more suitable for training very deep networks, *e.g.* GoogLeNet.

Experiments and Analysis

Overall Settings

To evaluate our method thoroughly, we conduct experiments on four datasets with different scales, including CIFAR-10, CIFAR-100 (Krizhevsky and Hinton 2009), MNIST (LeCun et al. 1998) and a much larger dataset, ImageNet (Deng et al. 2009) with two popular deep networks, *i.e.*, NIN (Lin, Chen, and Yan 2013) and GoogLeNet (Szegedy et al. 2014). NIN is used on CIFAR-10, CIFAR-100 and MNIST and GoogLeNet is used on ImageNet. NIN replaces the single linear convolution layers in the conventional CNNs by multilayer perceptrons, and uses the global average pooling layer to generate feature maps for each category. Compared to NIN, GoogLeNet is much larger with 22 layers built on *Inception* model, which can be seen as a deeper and wider extension of NIN. Both these two networks have achieved state-of-the-art performance on the datasets we use.

Since we mainly focus on testing the effects of SReLU on the performance of deep networks, in all our experiments, we only replace the ReLU in the original networks with SReLU and keep the other parts of networks unchanged. For the setting of hyperparameters (such as learning rate, weight decay and dropout ratio, etc.), we follow the published configurations of original networks. To compare LReLU with our method, we try different slope values in Eqn. (4) and picks the one that gets the best performance on validation set. For PReLU in our experiments, we follow the initialization methods presented in (He et al. 2015). For every dataset, we randomly sample 20% of the total training data as the validation set to configure the needed hyperparam-

Table 2: Error rates on CIFAR-10 and CIFAR-100. In the column for comparing the no. of parameters, the number after “+” is the extra number of parameters (in KB) introduced by corresponding methods. For the row of NIN + APL, 5.68K and 2.84K correspond to the extra parameters for CIFAR-10 and CIFAR-100, respectively.

Model	No. of Param.(B)	CIFAR-10	CIFAR-100
Without Data Augmentation			
Maxout	>5M	11.68%	38.57%
Prob maxout	>5M	11.35%	38.14%
APL	>5M	11.38%	34.54%
DSN	0.97M	9.78%	34.57%
Tree based priors	-	-	36.85%
NIN	0.97M	10.41%	35.68%
NIN + ReLU	0.97M	9.67%	35.96%
NIN + LReLU	0.97M	9.75%	36.00%
NIN + PReLU(ours)	0.97M + 1.42K	9.74%	35.95%
NIN + APL	0.97M + 5.68K/2.84K	9.59%	34.40%
NIN + SReLU ¹ (ours)	0.97M + 5.68K	8.61%	31.36%
NIN + SReLU (ours)	0.97M + 5.68K	8.41%	31.10%
With Data Augmentation			
Maxout	>5M	9.38%	-
Prob maxout	>5M	9.39%	-
APL	>5M	9.89%	33.88%
DSN	0.97M	8.22%	-
NIN	0.97M	8.81%	-
NIN + ReLU	0.97M	7.73%	32.75%
NIN + LReLU	0.97M	7.69%	32.70%
NIN + PReLU (ours)	0.97M + 1.42K	7.68%	32.67%
NIN + APL	0.97M + 5.68K/2.84K	7.51%	30.83%
NIN + SReLU (ours)	0.97M + 5.68K	6.98%	29.91%

ters in different methods. After fixing hyperparameters, we train the model from the scratch with the whole training data. For SReLU, we use $a^l = 0.2$ and $k = \lceil 0.9 \times |\mathbf{X}_i| \rceil$ for all datasets. In all experiments, we *ONLY* use single model and single view test.

We choose Caffe (Jia et al. 2014) as the platform to conduct our experiments. To reduce the training time, four NVIDIA TITAN GPUs are employed in parallel for training. Other hardware information of the PCs we use includes Intel Core i7 3.3GHz CPU, 64G RAM and 2T hard disk. The codes of SReLU are available at <https://github.com/AIROBOTAI/caffe/tree/SReLU>.

CIFAR

The CIFAR-10 and CIFAR-100 datasets contain color images with size of 32x32 from 10 and 100 classes, respectively. Both of them have 50,000 training images and 10,000 testing images. The preprocessing methods follow the way used in (Goodfellow et al. 2013). The comparison results of SReLU with other methods (including maxout (Goodfellow et al. 2013), prob maxout (Springenberg and Riedmiller 2013), APL (Agostinelli et al. 2014), DSN (Lee et al. 2014), tree based priors (Srivastava and Salakhutdinov 2013), NIN (Lin, Chen, and Yan 2013), etc.) on these two datasets either when the data augmentation is applied or not

¹Manually set initialization parameters in SReLU

are shown in Tabel 2, from which we can see that our proposed SReLU achieves the best performance against all the compared methods.

When no data augmentation is used, compared with ReLU, LReLU and PReLU, our method reduces the error significantly by 1.26%, 1.34%, 1.33% on CIFAR-10, respectively. On CIFAR-100, the error reduction is 4.86%, 4.90%, 4.85%, respectively. SReLU also demonstrates superiority by surpassing other activation functions including APL and maxout. When compared with other deep network methods, such as tree based priors and DSN, our method also beats them by a remarkable gap, demonstrating a promising ability to help boost the performance of deep models. We also compare the number of parameters used in each method, from which we notice that SReLU only incurs a very slight increase (5.68K) to the total number of parameters (0.97M in original NIN). APL uses the same number of additional parameters as SReLU on CIFAR-10, but its performance in either case of applying data augmentation or not is inferior to our method. The convergence curve of SReLU with other methods on CIFAR-10 and CIFAR-100 are shown in Figure 4(a) and Figure 4(b), respectively.

To observe the learned parameters of SReLU, we list in Table 3 the parameters' values after the training phase. Since the SReLU we use are channel-wise, we simply calculate the average of the input for all SReLU in the same layer. It is interesting to observe that SReLU in different layers learn meaningful parameters in coincide with our motivations. For example, the SReLU following conv1 and ccp1 learns a^r less than 1 (0.81 and 0.77, respectively) on CIFAR-10, while SReLU following conv3 and ccp5 on CIFAR-100 learns a^r larger than 1 (1.42 and 1.36, respectively). SReLU following conv2 on CIFAR-10 learns a^r nearly equal to 1 (1.01). These experimental results verify that SReLU has a strong ability to learn various forms of nonlinear functions, which can either be convex or non-convex. Moreover, in Table 3, we can see that t^r is of very large value in higher layers. It's because that the inputs of SReLU have higher average values than the ones in lower layers. Therefore, SReLU in higher layers learns larger t^r for adapting to inputs. This demonstrates the strong adaptive ability of SReLU to distribution of its inputs.

In the experiments on the augmented version of CIFAR-10 and CIFAR-100, we simply use random horizontal reflection during training for both datasets. In this case, SReLU still consistently outperforms other methods.

MNIST

MNIST (LeCun et al. 1998) contains 70,000 28x28 gray scale images of numerical digits from 0 to 9, divided as 60,000 images for training and 10,000 images for testing. In this dataset, we do not apply any preprocessing to the data and only compare models without data augmentation. The experiment results on this dataset are shown in Tabel 4, from which we see SReLU performs better than other methods.

Table 3: The parameters' values of SReLU after training with NIN on CIFAR-10 and CIFAR-100, respectively. The layers listed in the tabel are all convolution layers. "conv" layers are with kernel sizes larger than 1 and "cccp" layers are with kernel sizes equal to 1. Each layer in the tabel is followed by channel-wise SReLU. For more details of the layers in NIN, please refer to (Lin, Chen, and Yan 2013)

layers	CIFAR-10 / CIFAR-100			
	t^r	t^l	a^r	a^l
conv1	0.91 / 0.73	-0.48 / -0.68	0.81 / 0.62	-0.25 / -0.22
cccp1	1.06 / 0.52	-0.36 / -0.34	0.77 / 0.38	-0.04 / 0.04
cccp2	1.27 / 0.37	-0.20 / -0.26	0.47 / 0.51	0.39 / 0.44
conv2	5.32 / 4.02	-0.31 / -0.51	1.01 / 0.88	0.07 / 0.06
cccp3	6.95 / 4.73	-0.21 / -0.79	0.92 / 0.64	-0.01 / 0.05
cccp4	8.18 / 5.79	-0.08 / -0.13	0.77 / 0.56	0.61 / 0.45
conv3	25.17 / 23.72	-0.15 / -0.61	1.21 / 1.42	0.05 / 0.07
cccp5	31.09 / 36.44	-0.47 / -0.46	0.97 / 1.36	-0.16 / -0.02
cccp6	72.03 / 66.13	-0.13 / -0.21	1.53 / 1.23	-0.44 / -0.35

Table 4: Error rates on MNIST without data augmentation.

Model	No. of Param.(B)	Error Rates
Stochastic Pooling	-	0.47%
Maxout	0.42M	0.47%
DSN	0.35M	0.35%
NIN + ReLU	0.35M	0.47%
NIN + LReLU (ours)	0.35M	0.42%
NIN + PReLU (ours)	0.35M + 1.42K	0.41%
NIN + SReLU (ours)	0.35M + 5.68K	0.35%

Table 5: Error rates on ImageNet. Tests are by single model single view.

Model	No. of Param.(B)	Error Rates
GoogLeNet ²	5M	11.1%
GoogLeNet + SReLU (ours)	5M + 21.6K	9.86%

ImageNet

To further evaluate our method on large-scale datasets, we perform a much more challenging image classification task on 1000-class ImageNet dataset, which contains about 1.2 million training images, 50,000 validation images and 100,000 test images. Our baseline model is GoogLeNet model, which achieved the best performance on image classification in ILSVRC 2014 (Russakovsky et al. 2015). We run experiments using the publicly available configurations in Caffe (Jia et al. 2014). For this dataset, no additional preprocessing method is used except subtracting the image mean from each input raw image.

Table 5 compares the performance of GoogLeNet using SReLU and the original GoogLeNet released by Caffe. The GoogLeNet with SReLU achieves significant improvement (1.24%) on this challenging dataset compared with the original GoogLeNet using ReLU, at the cost of only 21.6K additional parameters (versus the total number of 5M parameters in the original GoogLeNet).

²https://github.com/BVLC/caffe/tree/master/models/bvlc_googlenet




Figure 4: (a) The convergence curves of SReLU and other methods on CIFAR-10. (b) the convergence curves of SReLU and other methods on CIFAR-100.

Conclusion

In this paper, inspired by the fundamental laws in psychophysics and neural sciences, we proposed a novel S-shaped rectified linear unit (SReLU) to be used in deep networks. Compared to other activation functions, SReLU is able to learn both convex and non-convex functions, and can be universally used in existing deep networks. Experiments on four datasets including CIFAR-10, CIFAR-100, MNIST and ImageNet with NIN and GoogLeNet demonstrate that SReLU effectively boosts the performance of deep networks. In our future work, we will exploit the applications of SReLU in other domains beyond vision, such as NLP.

References

Agostinelli, F.; Hoffman, M.; Sadowski, P. J.; and Baldi, P. 2014. Learning activation functions to improve deep neural networks. *CoRR* abs/1412.6830.

Cireřan, D. C.; Meier, U.; Gambardella, L. M.; and Schmidhuber, J. 2011. Convolutional neural network committees for handwritten character classification. In *Document Analysis and Recognition (ICDAR), 2011 International Conference on*, 1135–1139. IEEE.

Dayan, P., and Abbott, L. F. 2001. *Theoretical neuroscience*, volume 806. Cambridge, MA: MIT Press.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 248–255. IEEE.

Fechner, G. 1965. Elements of psychophysics.

Girshick, R. B.; Donahue, J.; Darrell, T.; and Malik, J. 2013. Rich feature hierarchies for accurate object detection and semantic segmentation. *CoRR* abs/1311.2524.

Goodfellow, I. J.; Warde-Farley, D.; Mirza, M.; Courville, A.; and Bengio, Y. 2013. Maxout networks. *arXiv preprint arXiv:1302.4389*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *CoRR* abs/1502.01852.

Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R. B.; Guadarrama, S.; and Darrell, T. 2014.

Caffe: Convolutional architecture for fast feature embedding. *CoRR* abs/1408.5093.

Johnson, K. O.; Hsiao, S. S.; and Yoshioka, T. 2002. Book review: neural coding and the basic law of psychophysics. *The Neuroscientist* 8(2):111–121.

Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.

Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2014. Deeply-supervised nets. *arXiv preprint arXiv:1409.5185*.

Lin, M.; Chen, Q.; and Yan, S. 2013. Network in network. *CoRR* abs/1312.4400.

Maas, A. L.; Hannun, A. Y.; and Ng, A. Y. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30.

Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 807–814.

Nieder, A. 2005. Counting on neurons: the neurobiology of numerical competence. *Nature Reviews Neuroscience* 6(3):177–190.

Randall, D.; Burggren, W. W.; French, K.; and Eckert, R. 2002. *Eckert animal physiology*. Macmillan.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 1–42.

Springenberg, J. T., and Riedmiller, M. 2013. Improving deep neural networks with probabilistic maxout units. *arXiv preprint arXiv:1312.6116*.

Srivastava, N., and Salakhutdinov, R. R. 2013. Discriminative transfer learning with tree-based priors. In Burges, C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems 26*. Curran Associates, Inc. 2094–2102.

Stevens, S. S. 1957. On the psychophysical law. *Psychological review* 64(3):153.

Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2014. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*.

Weber, E. 1851. Annotationes anatomicae et physiologicae [anatomical and physiological annotations]. *Leipzig: CF Koehler*.