NEURAL NETWORKS WITH TRAINABLE MATRIX ACTIVATION FUNCTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

The training process of neural networks usually optimize weights and bias parameters of linear transformations, while nonlinear activation functions are pre-specified and fixed. This work develops a systematic approach to constructing matrix activation functions whose entries are generalized from ReLU. The activation is based on matrix-vector multiplications using only scalar multiplications and comparisons. The proposed activation functions depends on parameters that are trained along with the weights and bias vectors. Neural networks based on this approach are simple and efficient and are shown to be robust in numerical experiments.

1 INTRODUCTION

In recent decades, deep neural networks (DNNs) have achieved significant successes in many fields such as computer vision and natural language processing (Voulodimos et al., 2018), (Otter et al., 2018). The DNN surrogate model is constructed using recursive composition of linear transformations and nonlinear activation functions. To achieve good performance, it is essential to choose activation functions suitable for specific applications. In practice, Rectified Linear Unit (ReLU) is one of the most popular activation functions for its simplicity and efficiency. A drawback of ReLU is the presence of vanishing gradient in the training process, known as dying ReLU problem (Lu et al., 2019). Several relatively new activation approaches are proposed to overcome this problem, e.g., the simple Leaky ReLU, and Piecewise Linear Unit (PLU) (Nicolae, 2018), Softplus (Glorot et al., 2011), Exponential Linear Unit (ELU) (Clevert et al., 2016), Scaled Exponential Linear Unit (SELU) (Klambauer et al., 2017), Gaussian Error Linear Unit (GELU) (Hendrycks & Gimpel, 2016).

Although the aforementioned activation functions are shown to be competitive in benchmark tests, they are still fixed nonlinear functions. In a DNN structure, it is often hard to determine a priori the optimal activation function for a specific application. In this paper, we shall generalize ReLU and introduce arbitrary trainable matrix activation functions. The effectiveness of the proposed approach are validated using function approximation examples and well-known benchmark datasets such as `MNIST` and `CIFAR-10`. There are a few classical works on adaptively tuning of parameters in the training process, e.g., the parametric ReLU (He et al., 2015b). However, our adaptive matrix activation functions are shown to be competitive and more robust in those experiments.

1.1 PRELIMINARIES

We consider the general learning process based on a given training set $\{(x_n, f_n)\}_{n=1}^N$, where the inputs $\{x_n\}_{n=1}^N \subset \mathbb{R}^d$ and outputs $\{f_n\}_{n=1}^N \subset \mathbb{R}^J$ are implicitly related via an unknown target function $f : \mathbb{R}^d \to \mathbb{R}^J$ with $f_n = f(x_n)$. The ReLU activation function is a piecewise linear function given by

$$\sigma(t) = \max\{t, 0\}, \quad \text{for} \quad t \in \mathbb{R}. \tag{1}$$

In the literature σ is acting component-wise on an input vector. In a DNN, let L be the number of layers and n_ℓ denote the number of neurons at the ℓ-th layer for $0 \le \ell \le L$ with $n_0 = d$ and $n_L = J$. Let $\mathcal{W} = (W_1, W_2, \ldots, W_L) \in \prod_{\ell=1}^L \mathbb{R}^{n_\ell \times n_{\ell-1}}$ denote the tuple of admissible weight matrices and $\mathcal{B} = (b_1, b_2, \ldots, b_L) \in \prod_{\ell=1}^L \mathbb{R}^{n_\ell}$ the tuple of admissible bias vectors. The ReLU

DNN approximation to f at the ℓ-th layer is recursively defined as

$$\eta_\ell(x) := \sigma(W_\ell \eta_{\ell-1}(x) + b_\ell) \in \mathbb{R}^{n_\ell}, \quad \eta_0(x) = x \in \mathbb{R}^d. \tag{2}$$

The traditional training process for such a DNN is to find optimal $\mathcal{W}_* \in \prod_{\ell=1}^L \mathbb{R}^{n_\ell \times n_{\ell-1}}$, $\mathcal{B}_* \in \prod_{\ell=1}^L \mathbb{R}^{n_\ell}$, (and thus optimal $\eta_L = \eta_{L,\mathcal{W}_*,\mathcal{B}_*}$) such that

$$(\mathcal{W}_*, \mathcal{B}_*) = \arg\min_{\mathcal{W},\mathcal{B}} E(\mathcal{W}, \mathcal{B}), \quad \text{where} \quad E(\mathcal{W}, \mathcal{B}) = \frac{1}{N} \sum_{n=1}^N \left| f_n - \eta_{L,\mathcal{W},\mathcal{B}}(x_n) \right|^2. \tag{3}$$

In other words, $\eta_{L,\mathcal{W}_*,\mathcal{B}_*}$ best fits the data with respect to the discrete ℓ^2 norm within the function class $\{\eta_{L,\mathcal{W},\mathcal{B}}\}$. In practice, the sum of squares norm in E could be replaced with more convenient norms.

2 TRAINABLE MATRIX ACTIVATION FUNCTION

Having a closer look at ReLU σ, we observe that the activation $\sigma(\xi_\ell(x))$ with $\xi_\ell := W_\ell \eta_{\ell-1} + b_\ell$ could be given as a matrix-vector multiplication $\sigma(\xi_\ell(x)) = D_\ell(x)\xi_\ell(x)$, where D_ℓ is a *diagonal* matrix-valued function mapping from $\mathbb{R}^{n_\ell}$ to $\mathbb{R}^{n_\ell \times n_\ell}$ with entries from the discrete set $\{0, 1\}$. This is a simple but quite useful observation. There is no reason to restrict on $\{0, 1\}$ and we thus look for a larger set of values over which the diagonal entries of D_ℓ are running or sampled. With slight abuse of notation, our new DNN approximation to f is calculated using the following recurrence relation

$$\eta_0(x) = x \in \mathbb{R}^d, \quad \xi_\ell(x) = W_\ell \eta_{\ell-1}(x) + b_\ell, \quad \eta_\ell = (D_\ell \circ \xi_\ell)\xi_\ell, \quad \ell = 1, \ldots, L. \tag{4}$$

Here each D_ℓ is diagonal and is of the form

$$D_\ell(y) = \mathrm{diag}(\alpha_{\ell,1}(y_1), \alpha_{\ell,2}(y_2), \ldots, \alpha_{\ell,n_\ell}(y_{n_\ell})), \quad y \in \mathbb{R}^{n_\ell}, \tag{5}$$

where $\alpha_{\ell,i}(y_i)$ is a nonlinear function to be determined. Since piecewise constant functions can approximate a continuous function within arbitrarily high accuracy, we specify $\alpha_{\ell,i}$ with $1 \le i \le n_\ell$ as

$$\alpha_{\ell,i}(s) = \begin{cases} t_{\ell,i,0}, & s \in (-\infty, s_{\ell,i,1}], \\ t_{\ell,i,1}, & s \in (s_{\ell,i,1}, s_{\ell,i,2}], \\ \vdots & \\ t_{\ell,i,m_{\ell,i}-1}, & s \in (s_{\ell,i,m_{\ell,i}-1}, s_{\ell,i,m_{\ell,i}}], \\ t_{\ell,i,m_{\ell,i}}, & s \in (s_{\ell,i,m_{\ell,i}}, \infty), \end{cases} \tag{6}$$

where $m_{\ell,i}$ is a positive integer and $\{t_{\ell,i,j}\}_{j=0}^{m_{\ell,i}}$ and $\{s_{\ell,i,j}\}_{j=1}^{m_{\ell,i}}$ are constants. We may suppress the indices ℓ, i in $\alpha_{\ell,i}$, $m_{\ell,i}$, $t_{\ell,i,j}$, $s_{\ell,i,j}$ and write them as α, m, t_j, s_j when those quantities are uniform across layers and neurons. If $m = 1$, $s_1 = 0$, $t_0 = 0$, $t_1 = 1$, then the DNN equation 4 is exactly the ReLU DNN. If $m = 1$, $s_1 = 0$, $t_1 = 1$ and t_0 is a fixed small negative number, equation 4 reduces to the DNN based on Leaky ReLU. If $m = 2$, $s_1 = 0$, $s_2 = 1$, $t_0 = t_2 = 0$, $t_1 = 1$, then $\alpha = \alpha_{\ell,i}$ actually represents a discontinuous activation function.

In our case, we shall fixed some parameters from $\cup_{\ell=1}^L \cup_{i=1}^{n_\ell} \{t_{\ell,i,j}\}_{j=0}^{m_{\ell,i}}$ and $\cup_{\ell=1}^L \cup_{i=1}^{n_\ell} \{s_{\ell,i,j}\}_{j=1}^{m_{\ell,i}}$ and let the rest of them vary in the training process. The resulting DNN might use different activation functions for different neurons and layers, adapted to the target function f. Since ReLU and Leaky ReLU are included by our DNN as special cases, the proposed DNN is supposed to behave no worse than the traditional ones in practice. In the following, we call equation 4 with the activation approach in equation 6 a DNN based on the "trainable matrix activation function (TMAF)".

Starting from the diagonal activation D_ℓ, we can go one step further to construct more general activation matrices. First we note that D_ℓ could be viewed as a nonlinear operator $T_\ell : [C(\mathbb{R}^d)]^{n_\ell} \to [C(\mathbb{R}^d)]^{n_\ell}$, where

$$[T_\ell(g)](x) = D_\ell(g(x))g(x), \quad g \in [C(\mathbb{R}^d)]^{n_\ell}, \quad x \in \mathbb{R}^d.$$

There seems to be no convincing reason to consider only diagonal operators. A more ambitious possibility is to consider a trainable nonlinear activation *operator* determined by more general matrices, for example, by following tri-diagonal operator

$$[T_\ell(g)](x) = \begin{pmatrix} \alpha_{\ell,1} & \beta_{\ell,2} & 0 & \cdots & 0 \\ \gamma_{\ell,1} & \alpha_{\ell,2} & \beta_{\ell,3} & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_{\ell,n_\ell-1} & \beta_{\ell,n_\ell} \\ 0 & 0 & \cdots & \gamma_{\ell,n_\ell-1} & \alpha_{\ell,n_\ell} \end{pmatrix} g(x), \quad x \in \mathbb{R}^d. \tag{7}$$

The diagonal $\{\alpha_{\ell,i}\}$ is given in equation 6 while the off-diagonals $\beta_{\ell,i}$, $\gamma_{\ell,i}$ are piecewise constant functions in the i-th coordinate y_i of $y \in \mathbb{R}^{n_\ell}$ defined in a fashion similar to $\alpha_{\ell,i}$. Theoretically speaking, even trainable full matrix activation is possible despite of potentially huge training cost. In summary, the corresponding DNN based on trainable nonlinear activation operators $\{T_\ell\}_{\ell=1}^L$ reads

$$\eta_0(x) = x \in \mathbb{R}^d, \quad \xi_\ell(x) = W_\ell \eta_{\ell-1}(x) + b_\ell, \quad \eta_\ell := T_\ell(\xi_\ell), \quad \ell = 1, \ldots, L. \tag{8}$$

The evaluation of D_ℓ and T_ℓ are cheap because they require only scalar multiplications and comparisons. When calling a general-purpose packages such as PyTorch in the training process, it is observed that the computational time of D_ℓ and T_ℓ is comparable to the classical ReLU.

REMARK

Our observation also applies to an activation function σ other than ReLU. For example, we may rescale $\sigma(x)$ to obtain $\sigma(\omega_{i,\ell} x)$ for constants $\{\omega_{i,\ell}\}$ varying layer by layer and neuron by neuron. Then $\sigma(\omega_{i,\ell} x)$ are used to form a matrix activation function and a TMAF DNN, where $\{\omega_{i,\ell}\}$ are trained according to given data and are adapted to the target function. This observation may be useful for in some applications.

3 NUMERICAL RESULTS

In this section, we demonstrate the feasibility and efficiency of TMAF by comparing it with the traditional ReLU-type activation function. In principle, all parameters in equation 6 are allowed to be trained while we shall only let function values $\{t_j\}$ vary for simplicity in the following. In each experiment, we use the same neural network structure, as well as learning rates, stochastic gradient descent (SGD) optimization and number NE of epochs (SGD iterations). In particular, the learning rate 1e-4 is used for epochs 1 to $\frac{\mathrm{NE}}{2}$ and 1e-5 is used for epochs $\frac{\mathrm{NE}}{2} + 1$ to NE. The optimization method we used is ADAM (Kingma & Ba (2015)). We provide two sets of numerical examples:

- Function approximations by TMAF network and ReLU-type networks;
- Classification problems for `MNIST` and `CIFAR-10` sets solved by TMAF and ReLU networks.

For the first class of examples we use the ℓ^2-loss function as defined in equation 3. For the classification problems we consider the *cross-entropy* that is widely used as a loss function in classification models. The cross entropy is defined using a training set which consists of p images, each with N pixels. Thus, we have a matrix $Z \in \mathbb{R}^{N \times p}$ and each column corresponds to an image with N pixels. Each image belongs to a fixed class c_j from the set of image classes $\{c_k\}_{k=1}^p$, where $c_j \in \{1, \ldots, M\}$. The network structure maps $Z \in \mathbb{R}^{N \times p}$ to $X \in \mathbb{R}^{M \times p}$, and each column x_j of X is an output of the network evaluation at the corresponding column z_j of Z. More precisely,

$$\begin{aligned} &Z = (z_1, \ldots, z_p), \quad X = (x_1, \ldots, x_p), \quad c_j = \mathrm{class}(z_j), \\ &x_j := \eta_{L,\mathcal{W},\mathcal{B}}(z_j), \quad x_j \in \mathbb{R}^M, \quad z_j \in \mathbb{R}^N, \quad j = 1, \ldots, p. \end{aligned}$$

The cross entropy loss function then is defined by

$$\begin{aligned} &\mathcal{C}(\mathcal{W}, \mathcal{B}) = \sum_{k=1}^p -\log\left(\frac{\exp(x_{c_k,k})}{\sum_{j=1}^M \exp(x_{j,k})}\right), \\ &(\mathcal{W}_*, \mathcal{B}_*) = \arg\max_{\mathcal{W},\mathcal{B}} \mathcal{C}(\mathcal{W}, \mathcal{B}). \end{aligned}$$

	Approximation error							
	Single hiden layer				Two hiden layers			
n	1	2	3	4	5	6	7	8
ReLU	0.089	0.34	0.39	0.41	0.14	0.41	0.45	0.47
TMAF	0.015	0.016	0.13	0.18	0.07	0.155	0.173	0.25

Table 1: Approximation errors for $\sin(\pi x_1 + \cdots + \pi x_n)$ by neural networks

To evaluate the loss function at a given image $z \in \mathbb{R}^N$, we first evaluate the network at z with the given $(\mathcal{W}, \mathcal{B}) = (\mathcal{W}_*, \mathcal{B}_*)$. We then define the class $c(z)$ of z and the loss $\text{loss}(z)$ at z as follows:

$$\text{c}(z) = \arg\max_{1 \le j \le M} \{\exp(a_j)\}, \quad \text{where} \quad a = \eta_{L,\mathcal{W}_*,\mathcal{B}_*}(z),$$

$$\text{loss}(z) = -\log\left(\frac{\exp(a_{c(z)})}{\sum_{j=1}^{M}\exp(a_j)}\right).$$

3.1 Approximation of a smooth function

As our first example, we use neural networks to approximate

$$f(x_1, \cdots, x_n) = \sin(\pi x_1 + \cdots + \pi x_n), \quad x_k \in [-2, 2], \quad k = 1, \ldots, n.$$

The training datasets consist of 20000 input-output data pairs where the input data are randomly sampled from the hypercube $[-2, 2]^n$. The networks given by equation 2 and equation 4 have single or double hidden layers with 20 neurons per layer. For TMAF D_ℓ in equation 5, the function $\alpha = \alpha_{\ell,i}$ in equation 6 uses intervals $(-\infty, -1.4)$, $(-1.4, -0.92]$,$(-0.92, -0.56]$,$(-0.56, -0.26]$,$(-0.26, 0]$,$(0, 0.26]$,$(0.26, 0.56]$,$(0.56, 0.92]$,$(0.92, 1.4]$, $(1.4, \infty)$. The approximation results are shown in Table 1 and Figure 1–3. It is observed that TMAF is the most accurate activation approach in these examples.

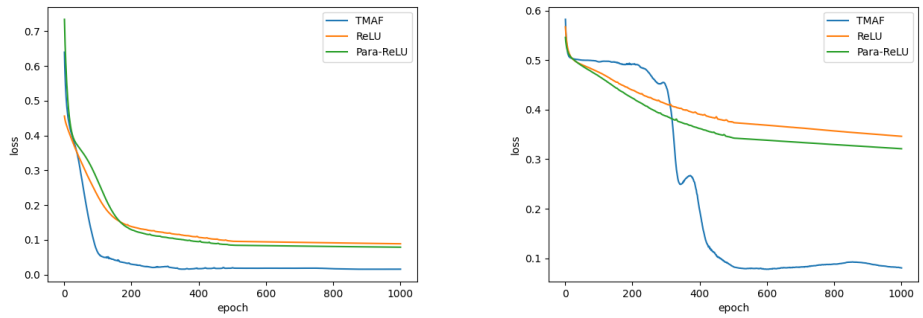

(a) $n = 1$ (b) $n = 2$

Figure 1: Training errors for $\sin(\pi x_1 + \cdots + \pi x_n)$, single hidden layer

3.2 Approximation of an oscillatory function

The next example is on approximating the following function having high and low frequency components

$$f(x) = \sin(100\pi x) + \cos(50\pi x) + \sin(\pi x), \tag{9}$$

see Figure 4 for an illustration. In fact, the function in equation 9 is rather difficult to capture by traditional approximation methods as it is highly oscillatory. The training datasets are 20000 input-output data pairs where the input data are randomly sampled from the interval $[-1, 1]$. We test the diagonal TMAF given in equation 5 and the function $\alpha = \alpha_{\ell,i}$. The definition of TMAF in equation 6 uses intervals $(-\infty, -5)$, $(-5 + kh, -5 + (k+1)h]$, $(5, \infty)$ with $h = 0.1$, $0 \le$

	Error
RELU	0.97
Diag-TMAF	0.033
Tri-diag TMAF	0.029

Table 2: Error comparison for $f(x) = \sin(100\pi x) + \cos(50\pi x) + \sin(\pi x)$

$k \leq 99$. We also consider the tri-diagonal TMAF equation 7, where $\{\alpha_{\ell,i}\}$ is the same as the diagonal TMAF, $\{\beta_{\ell,i}\}$ and $\{\gamma_{\ell,i}\}$ are all piecewise constants based on intervals $(-\infty, -5 + \underline{h})$, $(-5+kh+\underline{h}, -5+(k+1)h+\underline{h}]$, $(5+\underline{h}, \infty)$ and $(-\infty, -5+2\underline{h})$, $(-5+kh+2\underline{h}, -5+(k+1)h+2\underline{h}]$, $(5 + 2\underline{h}, \infty)$ with $\underline{h} = 0.1/3$, $0 \leq k \leq 99$, respectively. Numerical results could be found in Figures 4, 5 and Table 2.

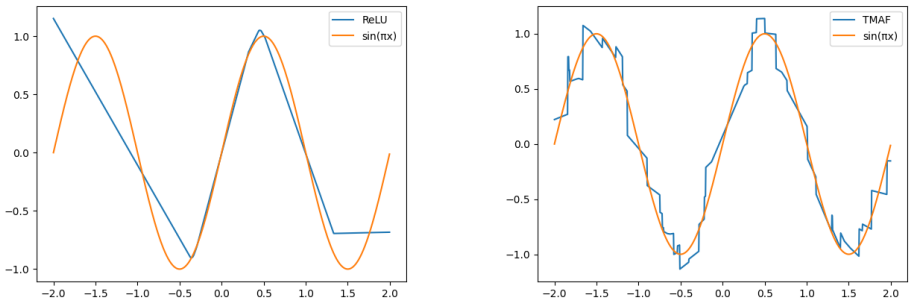

(a) ReLU (b) TMAF

Figure 2: Neural network approximations to $\sin(\pi x)$, single hidden layer

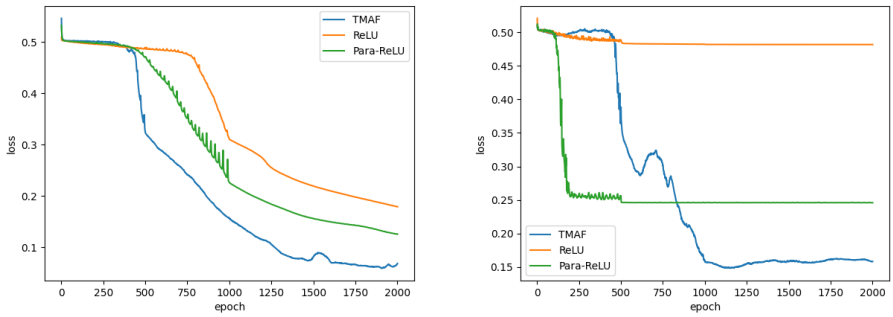

(a) $n = 5$ (b) $n = 6$

Figure 3: Training errors for $\sin(\pi x_1 + \cdots + \pi x_n)$, two hidden layers.

For this challenging problem, we note that the diagonal TMAF and tri-diagonal TMAF produce high-quality approximations while ReLU and parametric ReLU are not able to approximate the highly oscillating function within reasonable accuracy. It is observed from Figure 5 that ReLU actually captures the low frequency part of equation 9. To capture the high frequency, ReLU clearly has to use more neurons and thus much more weight and bias parameters. On the other hand, increasing the number of intervals in TMAF only lead to a few more training parameters.

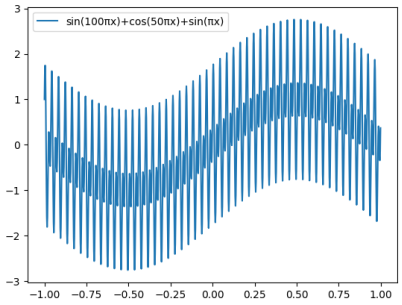

(a) Exact oscillating function

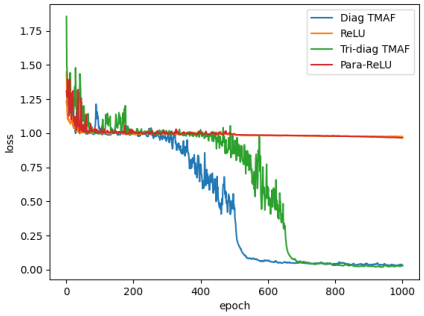

(b) Training loss comparison

Figure 4: Plot of $f(x) = \sin(100\pi x) + \cos(50\pi x) + \sin(\pi x)$ and training loss comparison

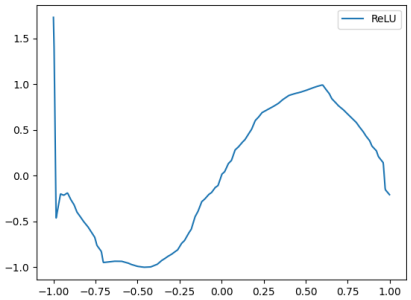

(a) ReLU approximation

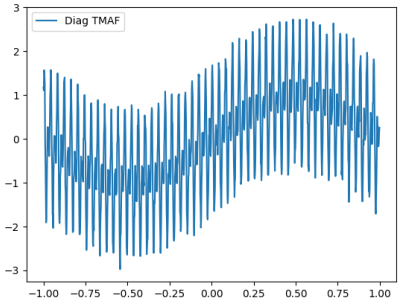

(b) TMAF approximation

Figure 5: Approximations to $f(x) = \sin(100\pi x) + \cos(50\pi x) + \sin(\pi x)$ by neural networks

3.3 Classification of MNIST and CIFAR-10 data sets

We now test TMAF by classifying images in the MNIST and CIFAR-10 data sets. For TMAF D_ℓ in equation 5, the function $\alpha = \alpha_{\ell,i}$ equation 6 uses intervals $(-\infty, -5)$, $(-5+k, -4+k]$, $(5, \infty)$ with $0 \leq k \leq 9$.

For the MNIST set, we implement single and double layer fully connected networks equation 2 and equation 4 with 10 neurons per layer (except at the first layer $n_0 = 764$), and ReLU or diagonal TMAF equation 5 activation. Numerical results are shown in Figures 6a, 6b, 7a, 7b and Table 3. We note the the TMAF with single hidden layer ensures higher evaluation accuracy than ReLU, see Table 3.

For the CIFAR-10 dataset, we use the ResNet18 network structure with 18 layers and number of neurons provided by (He et al., 2015a). The activation functions are still ReLU and the diagonal TMAF equation 5. Numerical results are presented in Figures 8a, 8b and Table 3. Those parameters given in (Paszke et al., 2019) are already tuned well with respect to ReLU. Nevertheless, TMAF still produces smaller errors in the training process and returns better classification results in the evaluation stage.

It is possible to improve the performance of TMAF applied to those benchmark datasets. The key point is to select suitable intervals in $\alpha_{\ell,i}$ to optimize the performance. A simple strategy is to let those intervals in equation 6 be varying and adjusted in the training process, which will be investigated in our future research.

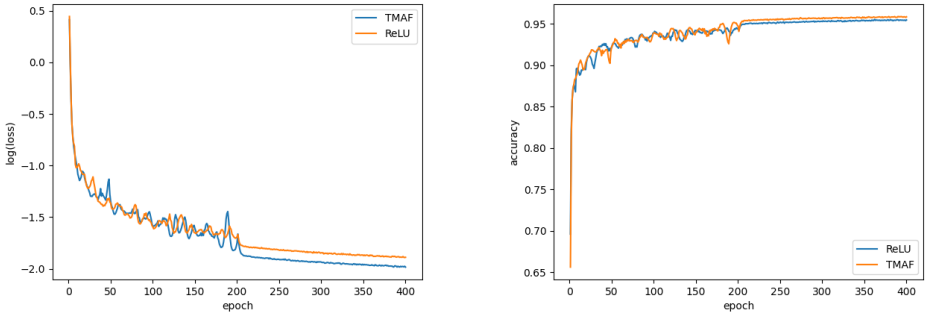

(a) Training loss (b) Classification accuracy

Figure 6: MNIST: Single hidden layer

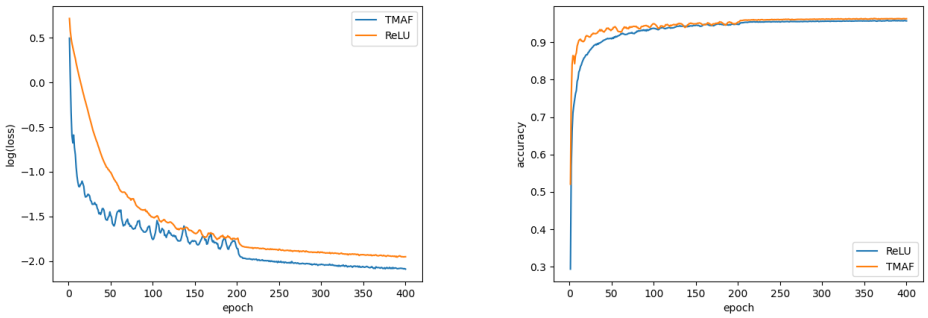

(a) Training loss (b) Classification accuracy

Figure 7: MNIST: Two hidden layers

REFERENCES

Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL http://arxiv.org/abs/1511.07289.

Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In Geoffrey J. Gordon, David B. Dunson, and Miroslav Dudík (eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011*, volume 15 of *JMLR Proceedings*, pp. 315–323. JMLR.org, 2011. URL http://proceedings.mlr.press/v15/glorot11a/glorot11a.pdf.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint*, arXiv: 1512.03385, 2015a. URL http://arxiv.org/arXiv:

Dataset	Evaluation Accuracy	
	ReLU	TMAF
MNIST (1 hidden layer)	86.1%	92.1%
MNIST (2 hidden layers)	91.8%	92.2%
CIFAR-10 (Resnet18)	92.8%	93.2%

Table 3: Evaluation accuracy for the MNIST and CIFAR-10

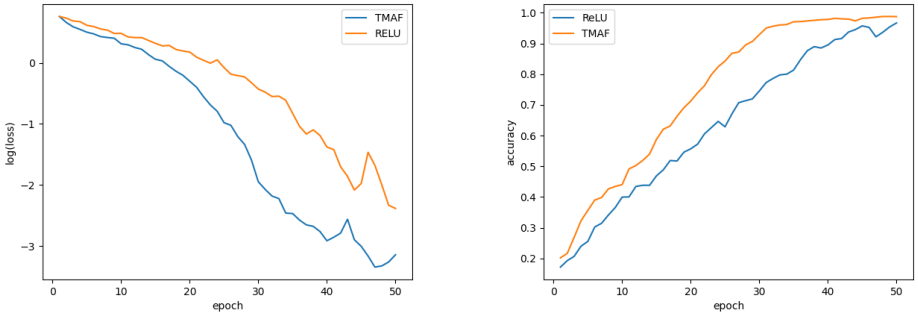

(a) Training loss

(b) Classification accuracy

Figure 8: Comparison between ReLU and TMAF for `CIFAR-10`

1512.03385.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint*, arXiv: 1502.01852, 2015b. URL http://arxiv.org/abs/1502.01852.

Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *arXiv preprint*, arXiv: 1606.08415, 2016. URL http://arxiv.org/abs/1606.08415.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 971–980, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/5d44ee6f2c3f71b73125876103c8f6c4-Abstract.html.

Lu Lu, Yeonjong Shin, Yanhui Su, and George E. Karniadakis. Dying relu and initialization: Theory and numerical examples. *arXiv preprint*, arXiv: 1903.06733, 2019. URL http://arxiv.org/abs/1903.06733.

Andrei Nicolae. PLU: the piecewise linear unit activation function. *arXiv preprint*, arXiv: 1809.09534, 2018. URL http://arxiv.org/abs/1809.09534.

Daniel W. Otter, Julian R. Medina, and Jugal K. Kalita. A survey of the usages of deep learning in natural language processing. *arXiv preprint*, arXiv: 1807.10854, 2018. URL http://arxiv.org/abs/1807.10854.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.

Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. Deep learning for computer vision: A brief review. *Computational Intelligence and Neuroscience*, 2018:7068349, February 2018. ISSN 1687-5265. URL https://doi.org/10.1155/2018/7068349.