

Differential Equation Units: Learning Functional Forms of Activation Functions from Data

MohamadAli Torkamani,¹ Shiv Shankar,²
Amirmohammad Rooshenas,² Phillip Wallis³

¹Amazon.com*

²University of Massachusetts Amherst

³Microsoft Dynamics 365 AI

¹alitor@amazon.com

²{sshankar, pedram}@cs.umass.edu

³phwallis@microsoft.com

Abstract

Most deep neural networks use simple, fixed activation functions, such as sigmoids or rectified linear units, regardless of domain or network structure. We introduce differential equation units (DEUs), an improvement to modern neural networks, which enables each neuron to learn a particular nonlinear activation function from a family of solutions to an ordinary differential equation. Specifically, each neuron may change its functional form during training based on the behavior of the other parts of the network. We show that using neurons with DEU activation functions results in a more compact network capable of achieving comparable, if not superior, performance when compared to much larger networks.

Introduction

Driven in large part by advancements in storage, processing, and parallel computing, deep neural networks (DNNs) have become capable of outperforming other methods across a wide range of highly complex tasks.

The advent of new activation functions such as rectified linear units (ReLU) (Nair and Hinton 2010), exponential linear units (ELU) (Clevert, Unterthiner, and Hochreiter 2015), and scaled exponential linear units (SELU) (Klambauer et al. 2017) address a network’s ability to effectively learn complicated functions, thereby allowing them to perform better on complicated tasks. The choice of an activation function is typically determined empirically by tuning, or due to necessity. For example, in modern deep networks, ReLU activation functions are often favored over sigmoid functions, which used to be a popular choice in the earlier days of neural networks. A reason for this preference is that the ReLU function is non-saturating and does not have the vanishing gradient problem when used in deep structures (Hochreiter 1998).

In all of the aforementioned activation functions, the functional form of the activation function is fixed. However, depending on data, different forms of activation functions may

be more suitable to describe data. In this paper, we introduce differential equation units (DEUs), where the activation function of each neuron is the nonlinear, possibly periodic solution of a parameterized second order, linear, ordinary differential equation. By learning the parameters of the differential equations using gradient descent on the loss function, the solutions to the differential equations change their functional form, thus adapting their forms to extract more complex features from data.

From an applicability perspective, our approach is similar to Maxout networks (Goodfellow et al. 2013) and adaptive piece-wise linear units (PLUs) (Agostinelli et al. 2014; Ramachandran, Zoph, and Le 2017).

Figure 1 shows nested noisy circles for which we trained neural networks with Maxout, ReLU, and our proposed DEU activation functions. For using ReLU, we needed at least two layers of four hidden units to separate the circles, while one layer of either two Maxout or two DEUs is enough to learn an appropriate decision boundary. While, the number of parameters learned by Maxout and PLU is proportional to the number of input weights to a neuron, and the number of linear units in that neuron, for each DEU, we learn *only five additional* parameters that give us highly flexible activation functions in return; altogether, resulting in more compact representations. Moreover, DEUs can represent a broad range of functions including harmonic functions, while Maxout is limited to piecewise approximation of convex functions (see Figure 2).

DEUs can also transform their forms during training in response to the behavior of other neurons in the network in order to describe data with interaction to the other part of the network, thus the neurons in a neural network may adopt different forms as their activation functions. The variety of activation function forms throughout a network enables it to encode more information, thus requiring fewer neurons to achieve the same performance compared to networks with fixed activation functions.

The main contribution of this paper is to explore the benefits of learning the form of activation functions through the introduction of differential equation units (DEUs). We also propose a learning algorithm to train the parameters of DEUs,

*The work was done when the author was affiliated with the University of Oregon
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Learned decision boundaries by neural networks with ReLU, Maxout, and DEU activation functions for classifying noisy circles.

and empirically show that neural networks with DEUs can achieve high performance with compact representations and are effective for solving real-world problems.

Differential Equation Units

Inspired by functional analysis and calculus of variations, instead of using a fixed activation function for each layer in a network, we propose a novel solution for learning an activation function for each neuron. Experimentally, we show that by allowing each neuron to learn its own activation function, the network as a whole can perform on par with (or even outperform) much larger baseline networks.

The main idea behind DEUs is to find the parameters of an ordinary differential equation (ODE) for each neuron in the network, whose solution would be used as the activation function of the corresponding neuron. As a result, each neuron learns a personalized activation function. We select (learn) the parameters of the differential equation from a low dimensional space (viz. five dimensions). By minimizing the networks loss function, our learning algorithm smoothly updates the parameters of the ODE, which results in an uncountably¹ extensive range of possible activation functions.

We parameterize the activation function of each neuron using a linear, second order ordinary differential equation $ay''(t) + by'(t) + cy(t) = g(t)$, parameterized by five coefficients (a, b, c, c_1, c_2) , where a, b , and c are scalars used to parameterize the ODE, c_1 and c_2 represent the initial conditions of the ODE's solution, and $g(t)$ is a regulatory function that we call the *core activation function*. The coefficients are the only additional parameters that we learn for each neuron and are trained through the backpropagation algorithm. To simplify the math and because it is a standard practice in control theory, we have set $g(t) = u(t)$, where $u(t)$ is the Heaviside (unit) step function. Therefore, the ODE that we have chosen has the following form:

$$ay''(t) + by'(t) + cy(t) = u(t),$$

$$\text{where } u(t) = \begin{cases} 0 & x \leq 0 \\ 1 & x > 0 \end{cases} \quad (1)$$

This model is often used to denote the exchange of energy between mass and stiffness elements in a mechanical system, or between capacitors and inductors in an electrical system (Ogata and Yang 2002). Interestingly, by using the solutions of this formulation as activation functions we can gain a few key properties: approximation or reduction to some of

¹Up to computational precision limitations.

the standard activation functions such as sigmoid or ReLU; the ability to capture oscillatory forms; and, exponential decay or growth.

The learning algorithm

For fixed a, b and c , the solution of the differential equation is:

$$y(t) = f(t; a, b, c) + c_1 f_1(t; a, b, c) + c_2 f_2(t; a, b, c), \quad (2)$$

for some functions f, f_1, f_2 . Here, $y(t)$ lies on an affine space parameterized by scalars c_1 and c_2 that represent the initial conditions of the solution (for desired values of y_0 and y'_0 at some $t = t_0$ such that $y(t_0) = y_0$ and $\frac{\partial y(t)}{\partial t}|_{t=t_0} = y'_0$).

Closed-form solutions First, we solve the differential equations parametrically and take derivatives of the closed-form solutions: $\frac{\partial y}{\partial t}$ with respect to its input t , and $\frac{\partial y}{\partial a}, \frac{\partial y}{\partial b}$, and $\frac{\partial y}{\partial c}$ with respect to parameter a, b , and c . Moreover, the derivatives with respect to c_1 and c_2 are f_1 and f_2 , respectively. This is done once. We have solved the equations and taken their derivatives using the software package Maple (2018). Maple also generates optimized code for the solutions, by breaking down equations in order to reuse computations.² Although we used Maple here, this task could have been done by pen and paper.

Training the parameters The parameters of the activation function (the ODE parameters and the appropriate initial conditions of its solution) are jointly trained with the neural networks' parameters using back-propagation.

We adopt regular backpropagation to update the values of parameters a, b, c, c_1 and c_2 for each neuron, along with using $\frac{\partial y}{\partial t}$ for updating network parameters w (input weights to the neuron), and to backpropagate the error to the lower layers.

We initialize network parameters using current best practices with respect to the layer type (e.g. linear layer, convolutional layer, etc.). We initialize parameters a, b, c for all neurons with a random positive number less than one, and strictly greater than zero. We initialize $c_1 = c_2 = 0.0$. To learn the parameters $\theta = [a, b, c, c_1, c_2]^T$ along with the weights w on input values to each neuron, we deploy a gradient descent algorithm. Both the weights w , as well as θ are learned using the conventional backpropagation algorithm with Adam updates (Kingma and Ba 2014). During training, we treat a, b, c, c_1 and c_2 like biases to the neuron (i.e., with input weight of 1.0) and update their values based on the direction of the corresponding gradients in each mini-batch.

Singularity of solutions If one or two of the coefficients a, b or c become zero, then the solution of the differential equation falls into a singularity subspace that is different than the affine function space of neighboring positive or negative values for those coefficients. For example, for $b = 0$ and $a * c > 0$, the solution will be $y(t) = \sin(\sqrt{\frac{c}{a}}t)c_2 + \cos(\sqrt{\frac{c}{a}}t)c_1 - \frac{u(t)}{c}(\cos(\sqrt{\frac{c}{a}}t) - 1)$, but for $b = c = 0$

²The code is available at <https://github.com/rooshenas/deu>

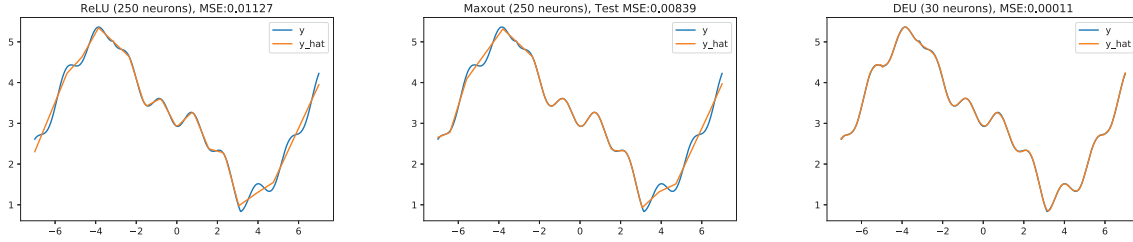


Figure 2: The problem of fitting a challenging function ($y = (\sin(t) - \cos(2t)^2)/2 + 4 * (1 + \arccos(\sin(t/2)))/3$). A neural network with 30 DEUs in one hidden layer fits the function much more accurately than neural networks of 250 neurons with ReLU and Maxout activation function.

Table 1: Subspaces and Singularities

Subspace	Solution
$a = 0, b = 0, c \neq 0$	$\sigma(t)/c$ (i.e., sigmoid when $c = 1$)
$a = 0, b \neq 0, c = 0$	$xu(x)/b + c_1$ (i.e., ReLU when $b = 1$ and $c_1 = 0$)
$a = 0, b \neq 0, c \neq 0$	$c_1 e^{-(cx)/b} - u(x)e^{-(cx)/b}/c + u(x)/c$
$a \neq 0, b = 0, c = 0$	$\frac{x^2 u(x)}{2a} + c_2 x + c_1$
$a < 0, b = 0, c < 0$ or $a > 0, b = 0, c > 0$	$c_2 \sin(\sqrt{\frac{c}{a}}x) + c_1 \cos(\sqrt{\frac{c}{a}}x) - \frac{u(x)}{c} (\cos(\sqrt{\frac{c}{a}}x) - 1)$
$a < 0, b = 0, c > 0$ or $a > 0, b = 0, c < 0$	$c_1 e^{\sqrt{-\frac{c}{a}}x} + c_2 e^{-\sqrt{-\frac{c}{a}}x} + \frac{u(x)e^{-\sqrt{-\frac{c}{a}}x}}{2c} + \frac{u(x)e^{\sqrt{-\frac{c}{a}}x}}{2c} - \frac{u(x)}{c}$
$a \neq 0, b \neq 0, c = 0$	$\frac{u(x)a}{b^2} e^{-\frac{bx}{a}} - \frac{u(x)a}{b^2} - c_1 \frac{a}{b} e^{-\frac{bx}{a}} + \frac{u(x)x}{b} + c_2$
$a \neq 0, b \neq 0, c \neq 0$	Four forms based on the sign of $\Delta = b^2 - 4ac$, a and c

and $a \neq 0$, the solution has the form of $y(t) = \frac{1}{2} \frac{u(t)t^2}{a} + c_1 t + c_2$. In this example, we observe that moving from $c > 0$ to $c = 0$ changes the resulting activation function from a pure oscillatory form to a parametric (leaky) rectified quadratic activation function. More formally, for a parametric variable quantity $p \in \{a, b, c, b^2 - 4ac\}$, if $p = 0$, then the solution of the differential equation may be different than the solution of the neighboring differential equations with $p \neq 0$ (see Table 1 for the complete set of singular subspaces). Moreover, with $p \rightarrow 0$, then we may have $y(t; p) \rightarrow \infty$ for certain values of t . In particular, when exponential forms are involved in the solution, this phenomenon can cause an extreme increase in the magnitude of the output value of the DEU. Therefore, we introduce a hyperparameter ϵ which is used as a threshold to clamp small values to zero. This simple trick avoids numerical computation errors as well as exponentiating large positive numbers when the denominator of the exponentiated fraction is very small.

Our learning algorithm allows the activation functions to jump over singular subspaces. However, if the absolute value of a, b , or c falls below ϵ , then the derivative with respect to that parameter becomes zero. The value of the parameter remains zero during the training if we use the regular derivative. In order to allow an activation function to “escape” a singular subspace, we introduce the concept of “outward gravitation” in the next subsection.

We do not allow $a = b = c = 0$, and in this rare case, we set b to ϵ . During the learning process at most two of a, b , and c can be zero. The sign of a, b, c , and $b^2 - 4ac$ might also change the solution of the ODE, which create “subspaces” that are individually solved in closed-form.

When $b^2 - 4ac$ is close to zero and $ac > 0$, the generic solution may become exponentially large. Therefore if $-\epsilon < b^2 - 4ac < \epsilon$ and $\text{sign}(a) == \text{sign}(c)$, we explicitly assume $a = \frac{b}{2} \text{sign}(a)$ $c = \frac{b}{2} \text{sign}(c)$ in our implementation to stabilize the solution and to avoid large function values (i.e., we force $b^2 - 4ac = 0$ in the solution of the ODE.).

Approximation of Dirac’s delta function The derivatives of the activation function with respect to t include Dirac’s delta function $\delta(t)$, which is the derivative of the Heaviside function. In the parametric derivatives, we substitute the delta function with its approximation $s \frac{e^{-s*t}}{(1+e^{-s*t})^2}$, which is the derivative of $\sigma(s*t) = \frac{1}{(1+e^{-s*t})}$. This approximation is a commonly used in practice for the delta function (Zahedi and Tornberg 2010). The larger values of s result in more accurate approximation of the delta function. In all of our experiments, we set $\epsilon = .01$, and $s = 100$ although further tuning might improve the results.

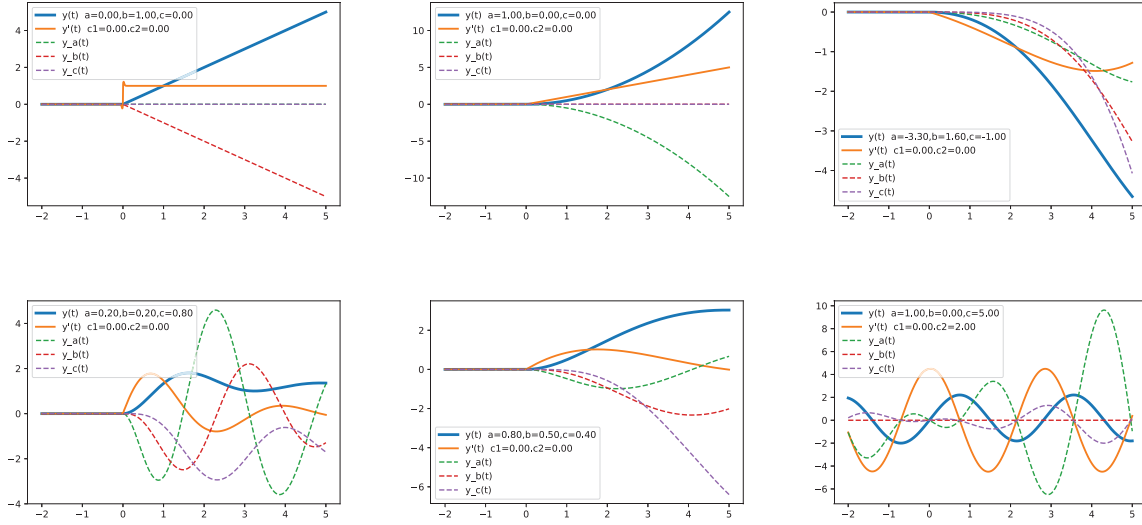


Figure 3: A sample set of DEU activation functions and their derivatives. The bold blue line is the activation function, and the orange solid line is its derivative with respect to t . The dashed lines are its derivatives with respect to a , b and c . First and second on the top row from left are ReLU and ReQU. The bump in the derivative of ReLU is an artifact of approximating Dirac’s delta.

Outward gravitation We need to allow a DEU to jump out of a singular subspace when it benefits the network. This way the activation function can recover from accidentally falling into singularity points, regardless of the initialization or how the order of training samples has changed the geodesic path that an activation function follows to achieve the appropriate functional form. To allow this, for a singular solution, we introduce a hypothetical non-singular differential equation that has the same initial conditions as the corresponding solution function at a reference point t^* .

If $f(t)$ and $f'(t)$ are the current activation function and its derivative w.r.t to the input, then the idea of outward gravitation is to find the nearest activation function $h(t)$ outside the singular subspace, such that $h(t^*) = f(t^*)$ and $h'(t^*) = f'(t^*)$ at the reference point t^* .

To solve this problem, for example, if $|b| < \epsilon$, to allow the activation function to jump to a subspace with $|b| \geq \epsilon$, we want the derivative of the current activation function with respect to b , as if b was present in the equation. To decide on this derivative direction, we use the solution function of the closest non-singular subspace.

For example, assuming $b = 0.005$ and $\epsilon = 0.01$, we consider $\tilde{b} = 0.01$, and then use this \tilde{b} along with the existing a and c to compute the derivative of the activation function with respect to b . We choose initial conditions \tilde{c}_1, \tilde{c}_2 so that the value and derivative of the activation function w.r.t. input remains the same. So the whole network won’t feel the change, but the value of b can change if it helps the network.

Implementation Due to the individual activation functions associated parameters per neuron, each neuron can have a different activation. For example, one neuron could have a sinusoidal activation function while another has a quadratic

form of ReLU. The direct way to implement our method would involve iterating over all neurons and computing the corresponding activation values individually, and would result in significantly higher latency than common activation functions such as ReLU. However, considering the ODE’s closed form solutions, such computations are parallelizable on modern computing units like GPUs. We first compute a mask to assign each neuron to the subspace in which its current activation function resides. Next, we iterate over all possible functional subspaces, and compute in parallel the output of each neuron assuming that its activation function lies in the current subspace while ignoring the validity of the solutions or parameters. Finally, we use the subspace masks computed earlier to select the correct output of the neuron. Similar technique can be applied during the backward pass to compute the gradients in parallel using the masks computed during the forward pass. The pseudo-code is detailed in Algorithm 1.

Algorithm 1 Parallelized DEU

```

1: procedure DEU(input)
2:   output  $\leftarrow$  0
3:   for each singularity space  $S$  do
4:     mask =  $\mathbf{1}_{[n \in S]}$   $\forall n \in neurons$ 
5:     if mask > 0 then
6:       output  $\leftarrow$  output + mask *  $f_S(input)$ 
7:     end if
8:   end for
9:   Return output
10: end procedure

```

Neural networks with DEUs are universal approximators

Feedforward neural networks with monotonically-increasing activation functions are universal approximators (Hornik, Stinchcombe, and White 1989; Barron 1993). Similarly networks with radial basis activation functions are also shown to be universal approximators (Park and Sandberg 1991). A Fourier approximation based argument shows that a neural network with DEU is also a universal approximator.

A geometric interpretation The solution set of a differential equation forms a functional manifold that is affine with respect to c_1 and c_2 , but is nonlinear in a , b , and c . Clearly, this manifold has a trivially low dimensional representation in \mathbb{R}^5 (i.e., $\{a, b, c, c_1, c_2\}$). Gradient descent modifies the functionals in this low dimensional space, and the corresponding functional on the solution manifold is used as the learned activation function. Figure 4 attempts to visually explain how, for example, a ReLU activation function transforms to a cosine activation function.




Figure 4: Left: the solutions of $ay'' + by' + cy = u(t)$ lie on a manifold of functions. Right: every point on this manifold can be equivalently represented as a point on a 5-dimensional space of $\{a, b, c, c_1, c_2\}$ (three shown here). The red arrow shows a path that an initialized function takes to gradually transform itself to a different one.

For synthetic datasets, we empirically show that only one DEU neuron that is initialized with ReLU ($a = c = c_1 = c_2 = 0$, $b = 1$) can transform itself to a sine function. In this scenario, the learned differential equation coefficients $\{a, b, c\}$ and initial condition coefficients $\{c_1, c_2\}$ perfectly represented a sine function after training. In contrast, the learned model by ordinary fixed activation FFNNs were much less accurate with significantly larger networks.

Reduction to common activation functions Two common activation functions are the sigmoid $\sigma(t) = \frac{1}{1+e^{-t}}$, and the rectified linear unit $\text{ReLU}(t) = \max(0, t)$. The sigmoid function is a smooth approximation of the Heavyside step function, and ReLU can be approximated by integrating sigmoid of $s * t$ for a large enough s : $\max(0, t) \approx \int_{-\infty}^t \frac{1}{1+e^{-sz}} dz = \log(1 + e^{st})/s$. Equivalently, $y(t) = \log(1 + e^t) + c_1 \approx \text{ReLU}(t) + c_1$ will be a solution of the following first order linear differential equation: $y'(t) = \frac{1}{1+e^{-st}} \approx u(t)$

We can set the core activation function $g(t)$ to $\sigma(t) = \frac{1}{1+e^{-st}}$, or to the step function $u(t)$. For $g(t) = \sigma(t)$ when

$a \neq 0$, the solutions of the differential equation will involve the Gauss hypergeometric and Li_2 functions, which are expensive to evaluate. Fortunately, if we set the right hand side of the ODE to $u(t)$, then the particular solutions will only involve common functions such as linear, exponential, quadratic, trigonometric, and hyperbolic functions.

In practice, if the learning algorithm decides that a and b should be zero, we use $g(t) = \frac{1}{1+e^{-st}}$ (i.e. $y(t) = \frac{1}{c*(1+e^{-st})}$). Otherwise, we use the step function to avoid complex-valued solutions that involve special mathematical functions. With these conditions in place, if $a = 0$, $b = 0$, and $c = 1$, we recover the sigmoid function; if $a = 0$, $b = 1$, and $c = 0$, we recover the ReLU function; if $a = 1$, $b = 0$, and $c = 0$, we obtain a parametric rectified quadratic form $y = \text{ReLU}(t)^2 + c_1t + c_2$ (similar to parametric ReLU (He et al. 2015; Xu et al. 2015)), which is the solution of $y''(t) = u(t)$. When $b^2 - 4ac < 0$, we observe an oscillatory behaviour. Depending on the sign of b , this can be decaying or exploding, but when $b = 0$, we observe a purely oscillating activation function.

The above-mentioned cases are only a few examples of solutions that could be chosen. The point to emphasize here is that an extensive range of functions can be generated by simply varying these few parameters (Figure 3 illustrates several examples).

Related Work

Recently, Chen et al. (2018) propose using ODE solvers in neural networks. However, the nature of our work is very different from Neural ODE. While that work maps iterative parts of the network to an ODE and uses external solver to find the solution and estimating the gradient, here we use the flexible functions that are the solutions of ODE as activation function. At the training time there is no ODE solver involved, but only functional forms of the solutions.

Oscillatory neurons are generally believed to be important for information processing in animal brains. Efforts have been made to explore the usefulness of periodic oscillations in neural networks since the 1990s, especially for medical applications (Minami, Nakajima, and Toyoshima 1999). However overall their applicability has not yet been appreciated (Sopena, Romero, and Alquezar 1999). In recent times, however, researchers have begun re-exploring the potential of periodic functions as activations (Parascandolo, Huttunen, and Virtanen 2016), and demonstrated their potentials for probabilistic time-series forecasting (Hatalis and Kishore 2017) and lane departure prediction (Tan, Chen, and Wang 2017). Furthermore recent work has demonstrated how pattern recognition can be achieved on physical oscillatory circuits (Velichko, Belyaev, and Boriskov 2019).

For vision, speech and other applications applications on mobile, or other resource-constrained devices, research has been ongoing to make compact networks. Projection-Net (Ravi 2017) and MobileNet (Howard et al. 2017) are both examples of methods that use compact DNN representations with the goal of on-device applications. In Projection-Net, a compact *projection* network is trained in parallel to the primary network, and is used for the on-device network

Table 2: Test accuracy of different models on the MNIST and Fashion-MNIST image classification task.

Model	Size	MNIST	Fashion-MNIST
MLP-ReLU	1411k	98.1	89.0
CNN-ReLU	30k	99.2	90.2
MLP-SELU	1293k	95.5	87.5
CNN-SELU	21k	98.8	89.6
MLP-PReLU	1293k	97.4	88.7
CNN-PReLU	21k	98.9	89.6
MLP-DEU	1292k	98.3	89.8
CNN-DEU	21k	99.2	91.5
MLP-Maxout	2043k	98.5	89.4
CNN-Maxout	26k	99.4	91.3

tasks. MobileNet, on the other hand, proposes a streamlined architecture in order to achieve network compactness. In these approaches, the network compactness is achieved at the expense of performance. We propose a different method for learning compact, powerful, stand-alone networks: we allow each neuron to learn its individual activation function enabling a compact neural network to achieve higher performance.

Results and Discussion

We have conducted several experiments to evaluate the performance and compactness of DEU networks.

Classification

We evaluate DEU on different models considering the classification performance and model size. We first use MNIST and Fashion-MNIST as our datasets to assess the behavior of DEUs with respect to the commonly used ReLU activation function, as well as Maxout and SELU. The neural network is a 2-layer MLP with 1024 and 512 dimensional hidden layers. While the CNN used is a 2-layer model made by stacking 32 and 16 dimensional convolutional filters atop one another followed by average pooling. DEUs are competitive or better than normal networks for these tasks while having substantially smaller number of parameters (see Table 2).

Next we perform a more direct comparison of the effect of DEU on classification performance against ReLU, PReLU (He et al. 2015), Maxout (Goodfellow et al. 2013), and Swish (Ramachandran, Zoph, and Le 2017) activation functions on the CIFAR-10 dataset. PReLU is similar to ReLU with a parametric leakage and Swish has the form of $f(x) = x * sigmoid(\beta x)$ with a learnable parameter β .

For these experiments, we keep the network architecture fixed to ResNet-18 (He et al. 2016a) and use the hyperparameter settings as in He et al. (2016a). We further show that this improvement persists across other model designs. First we use a preactivation ResNet (He et al. 2016b), which is a ResNet-like architecture with a slightly smaller size. Second, to assess suitability for reducing the model size, we experiment with a stunted ResNet-18. The stunted model is

created by removing all 256 and 512 dimensional filters from the standard ResNet-18 model. The result of this comparison (presented in Table 3) indicates that DEUs constantly work better than ReLU, PReLU, SELU, and Swish. Although Maxout is slightly better than DEU, it is using much more parameters (13272k vs 11174k for Resnet, 13267k vs 11174k for Preact-ResNet, and 801k vs 678k for ResNet-Stunted), which makes the comparison biased toward Maxout.

Convergence comparison on MNIST Figure 5 shows the classification error on MNIST across different activations as the training progresses. It is clear that DEUs are better at almost all steps. Also as one might expect, they are better in the initial epochs due to the greater expressiveness of the activation function.




Figure 5: Convergence comparison of different activation function on MNIST.

Computational cost Prima-facie our method seems extremely compute-intensive. However as described earlier, with judicious design one can parallelize the overall model to be quite fast. In the worst case, our method will be $|S|$ times slower than the same architecture with standard activation like ReLU, where $|S|$ is the number of singular subspace solution of the ODE. Nevertheless, in practice all subspaces are unlikely to occur together. In most of our experiments on real data, we observed that three to five different solutions appear. Furthermore, we evaluate computation time of Resnet-18 models with DEUs and ReLUs on CIFAR-10 using a batch size of 128 on a Tesla K40c GPU. The model with DEUs




Figure 6: Convergence comparison of Diabetes.

Table 3: Test accuracy using different ResNet architectures and activation functions on the CIFAR-10 image classification task.

Architecture	Size	ReLU	PReLU	SELU	Swish	Maxout	DEU
ResNet-18	11174k	91.25	92.1	92.2	91.9	92.5	92.5
Preact-ResNet	11170k	92.1	92.2	92.3	92.0	92.4	92.3
ResNet-Stunted	678k	89.3	89.4	90.5	90.1	91.1	90.7

takes 640ms total for combined forward and backward pass per batch, while the one with ReLUs requires 180ms per step. During prediction on test data, the time taken per batch is 111ms for the DEU-based model and 65ms for ReLU-based model.

To further explore the effect of additional parameters for training activation functions, we train a 2-layer fully connected network for MNIST. We kept the size of the first layer fixed and changed the size of the second layer from 16 to 1024. We train each configuration with different activation functions of DEU, Maxout with three affine maps, PReLU, and ReLU. Table 4 reports the training time for one batch in millisecond averaged over the entire epoch. The batch size is 256, and the training is done using Geforce GTX TITAN. As we can see training Maxout gets more expensive than DEU for larger fan-in.

Table 4: Average single-batch training time (in milliseconds) of a 2-layer fully connected network for MNIST varying the size of the second layers as well as its the activation functions.

Size	DEU	Maxout	PReLU	ReLU
16	47	40	37	37
32	47	40	38	37
64	49	43	38	37
128	53	47	40	39
256	55	54	43	44
512	68	89	50	47
1024	186	209	50	47

Regression

We compare the performance of neural networks with one hidden-layer of one, two, four, eight, and 16 neurons with DEU, ReLU, LeakyReLU (LReLU), SELU, Swish and Maxout activation functions applied to a standard diabetes regression dataset.³ We use 3-fold cross validation and report the average performance.

Figure 6 shows that the neural networks with DEUs achieve specifically better performance with more compact networks. We see that other activation functions do not surpass the single DEU neuron performance until they are eight or more neurons.

We further test our model on multiple regression tasks involving timeseries forecasting. For these tasks, we use the LSTNet model (Lai et al. 2018) and evaluate the functions based on the root relative squared error (RSE, lower is better)

³(<https://www4.stat.ncsu.edu/boos/var.select/diabetes.html>)

Table 5: Evaluation summary (in RSE and CORR) of different activations on different datasets: **Traffic**. A collection of 48 months (2015-2016) hourly data from the California Department of Transportation (6 and 12 hours look ahead prediction); **Solar-Energy**. Solar power production records in the year of 2006 sampled every 10 minutes from 137 PV plants in Alabama State (6 and 12 hours look ahead prediction); **Electricity**. The electricity consumption in kWh recorded every 15 minutes from 2012 to 2014 (12 hours look ahead prediction)

Func.	Traffic		Solar output		Elect.	
	6	12	6	12	12	
DEU	RSE	0.487	0.500	0.269	0.327	0.100
	CORR	0.870	0.863	0.965	0.947	0.912
ReLU	RSE	0.499	0.520	0.270	0.433	0.104
	CORR	0.869	0.851	0.965	0.906	0.900
Swish	RSE	0.483	0.505	0.270	0.329	0.104
	CORR	0.872	0.862	0.965	0.944	0.908
Maxout	RSE	0.493	0.501	0.265	0.328	0.107
	CORR	0.863	0.868	0.967	0.945	0.911

and correlation (CORR, higher is better) metrics. The results are presented in Table 5. DEU gives improvement in all cases and substantial improvements in some datasets.

Conclusion

In this paper we introduce differential equation units (DEUs), as novel activation functions based on the solutions of second-order ordinary differential equations (ODEs). DEUs can adapt their function form based on the features of data during training by learning the parameters of ODEs using gradient descent. We have showcased the ability of neural networks with DEUs to learn complicated concepts with a compact network representation. We have demonstrated DEUs' potential to outperform conventional activation function across a number of tasks, and with a reduced network size. Modern DNNs achieve performance gains in large by increasing the size of the network, which is not a sustainable trend. In response, we believe that this line of research can open future directions to explore more complex activation functions such as using the solutions of partial differential equations in order to compactly represent complex functions.

References

Agostinelli, F.; Hoffman, M.; Sadowski, P.; and Baldi, P. 2014. Learning activation functions to improve deep neural networks. *arXiv preprint arXiv:1412.6830*.

- Barron, A. R. 1993. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory* 39(3):930–945.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. In *Advances in neural information processing systems*, 6571–6583.
- Clevert, D.-A.; Unterthiner, T.; and Hochreiter, S. 2015. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*.
- Goodfellow, I. J.; Warde-Farley, D.; Mirza, M.; Courville, A.; and Bengio, Y. 2013. Maxout networks. In *Proceedings of the 30th International Conference on Machine Learning* 1319–1327.
- Hatalis, K., and Kishore, S. 2017. A composite quantile fourier neural network for multi-horizon probabilistic forecasting. *CoRR* 1–13.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016a. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016b. Identity mappings in deep residual networks. *CoRR* abs/1603.05027.
- Hochreiter, S. 1998. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6(02):107–116.
- Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2(5):359–366.
- Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klambauer, G.; Unterthiner, T.; Mayr, A.; and Hochreiter, S. 2017. Self-normalizing neural networks. *arXiv preprint arXiv:1706.02515*.
- Lai, G.; Chang, W.-C.; Yang, Y.; and Liu, H. 2018. Modeling long- and short-term temporal patterns with deep neural networks. *SIGIR 2018*.
- Maple. 2018. Maplesoft, a division of waterloo maple inc.
- Minami, K.; Nakajima, H.; and Toyoshima, T. 1999. Real-time discrimination of ventricular tachyarrhythmia with fourier-transform neural network. In *IEEE transactions on bio-medical engineering*, volume 46, 179–85.
- 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.
- Ogata, K., and Yang, Y. 2002. *Modern control engineering*, volume 4. Prentice hall India.
- Parascandolo, G.; Huttunen, H.; and Virtanen, T. 2016. Taming the waves: sine as activation function in deep neural networks.
- Park, J., and Sandberg, I. W. 1991. Universal approximation using radial-basis-function networks. *Neural computation* 3(2):246–257.
- Ramachandran, P.; Zoph, B.; and Le, Q. 2017. Searching for activation functions. *arXiv preprint arXiv:1710.05941*.
- Ravi, S. 2017. Projectionnet: Learning efficient on-device deep networks using neural projections. *arXiv preprint arXiv:1708.00630*.
- Sopena, J. M.; Romero, E.; and Alquezar, R. 1999. Neural networks with periodic and monotonic activation functions: a comparative study in classification problems.
- Tan, D.; Chen, W.; and Wang, H. 2017. On the use of monte-carlo simulation and deep fourier neural network in lane departure warning. In *IEEE Intelligent Transportation Systems Magazine*, volume 9, 76–90.
- Velichko, A.; Belyaev, M.; and Boriskov, P. 2019. A model of an oscillatory neural network with multilevel neurons for pattern recognition and computing. In *Electronics*, volume 8, 75.
- Xu, B.; Wang, N.; Chen, T.; and Li, M. 2015. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*.
- Zahedi, S., and Tornberg, A.-K. 2010. Delta function approximations in level set methods by distance function extension. *Journal of Computational Physics* 229(6):2199–2219.