Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)

Djork-Arné Clevert, Thomas Unterthiner & Sepp Hochreiter
Institute of Bioinformatics
Johannes Kepler University, Linz, Austria
`{okko,unterthiner,hochreit}@bioinf.jku.at`

arXiv:1511.07289v1 [cs.LG] 23 Nov 2015

Abstract

We introduce the "exponential linear unit" (ELU) which speeds up learning in deep neural networks and leads to higher classification accuracies. Like rectified linear units (ReLUs), leaky ReLUs (LReLUs) and parametrized ReLUs (PReLUs), ELUs also avoid a vanishing gradient via the identity for positive values. However ELUs have improved learning characteristics compared to the units with other activation functions. In contrast to ReLUs, ELUs have negative values which allows them to push mean unit activations closer to zero. Zero means speed up learning because they bring the gradient closer to the unit natural gradient. We show that the unit natural gradient differs from the normal gradient by a bias shift term, which is proportional to the mean activation of incoming units. Like batch normalization, ELUs push the mean towards zero, but with a significantly smaller computational footprint. While other activation functions like LReLUs and PReLUs also have negative values, they do not ensure a noise-robust deactivation state. ELUs saturate to a negative value with smaller inputs and thereby decrease the propagated variation and information. Therefore ELUs code the degree of presence of particular phenomena in the input, while they do not quantitatively model the degree of their absence. Consequently, dependencies between ELUs are much easier to model and distinct concepts are less likely to interfere.

We found that ELUs lead not only to faster learning, but also to better generalization performance once networks have many layers (≥ 5). ELU networks were among top 10 reported CIFAR-10 results and yielded the best published result on CIFAR-100, without resorting to multi-view evaluation or model averaging. On ImageNet, ELU networks considerably speed up learning compared to a ReLU network with the same architecture, obtaining less than 10% classification error for a single crop, single model network.

1 Introduction

Currently the most popular activation function for neural networks is the rectified linear unit (ReLU), which was first proposed for restricted Boltzmann machines Nair & Hinton (2010) and then successfully used for neural networks Glorot et al. (2011); Zeiler et al. (2013); Dahl et al. (2013). The ReLU activation function is linear with slope 1 for positive arguments and zero otherwise, that is, for positive values it is the identity and for negative values the zero function. Besides their sparseness, the main advantage of ReLUs is that they avoid the vanishing gradient Glorot et al. (2011); Maas et al. (2013); Hochreiter & Schmidhuber (1997); Hochreiter (1991; 1998); Hochreiter et al. (2001a) since their derivative is 1 for positive values. However ReLUs are non-negative and, therefore, have a mean activation larger than zero.

Units that have a non-zero mean activation act as bias for the next layer. If such units do not cancel each other out, learning causes a *bias shift* for units in next layer. The more units are correlated, the higher their bias shift. We will see that Fisher optimal learning, i.e., the natural gradient Amari (1997); Yang & Amari (1998); Amari (1998), would correct for the bias shift by adjusting the bias weight update. Thus, less bias shift brings the standard gradient closer to the natural gradient and

speeds up learning. Therefore we aim at activations that have a mean close to zero across the training examples, in order to decrease the bias shift effect.

Centering the activation functions at zero has been proposed in order to keep the off-diagonal entries of the Fisher information matrix small Raiko et al. (2012) but was previously found to speed up learning in neural networks LeCun et al. (1991; 1998); Schraudolph (1998). Centering also improved learning in restricted Boltzmann machines Cho et al. (2011); Montavon & Müller (2012). To counter the internal covariate shift of the network activations, "batch normalization" has been suggested which also centers the hidden activations Ioffe & Szegedy (2015). The Projected Natural Gradient Descent algorithm (PRONG) goes beyond batch normalization by implicitly whitening the activations Desjardins et al. (2015), therefore also centers them at zero. All these methods tweak the network parameters to obtain desired statistics of the activation functions. However distorting the network parameters may take back the last learning update and therefore hamper learning. Instead, natural gradient descent uses efficient update directions without restricting the network dynamics.

An alternative way to push the mean activation towards zero is using an appropriate activation function. Therefore $tanh$ was previously preferred over logistic functions LeCun et al. (1991; 1998). Recently "Leaky ReLUs" (LReLUs) have been proposed, which replace the negative part of the ReLU with a linear function Maas et al. (2013). Leaky ReLUs have been shown to be superior to ReLUs in terms of learning speed and performance Maas et al. (2013); Xu et al. (2015). They have later been generalized to Parametric Rectified Linear Units (PReLUs) He et al. (2015), which adjust the slope of the negative part during learning. PReLUs performed very good on large image benchmark data sets, where they showed better learning behavior than ReLUs He et al. (2015). Another version of leaky ReLUs are Randomized Leaky Rectified Linear Units (RReLUs), where the slope of the negative part is randomly sampled Xu et al. (2015). Evaluations on image benchmark datasets and convolutional networks demonstrated that non-zero slopes for the negative part in ReLUs improve results Xu et al. (2015).

In contrast to ReLUs, LReLUs, PReLUs, and RReLUs do no longer ensure a noise-robust deactivation state of the units. We propose an activation function that has negative values to allow for mean activations close to zero, but at the same time saturate to a negative value with decreasing input. The saturation effect decreases the variation of the units if deactivated, so the precise deactivation value is irrelevant. Such an activation function can code the degree of presence of particular phenomena in the input, but does not quantitatively model the degree of their absence. Therefore, our new activation function is more robust to noise. Consequently, dependencies between coding units are much easier to model and much easier to interpret. Furthermore, distinct concepts are much less likely to interfere with such activation functions since the deactivation state is non-informative, i.e. variance decreasing.

2 ZERO MEAN ACTIVATIONS SPEED UP LEARNING

To derive and analyze the bias shift effect mentioned in Section1 we use the natural gradient. The natural gradient is an update direction which corrects the gradient direction with the Fisher information matrix. The natural gradient method is founded on information geometry Amari (1985); Amari & Murata (1993) and became very popular for independent component analysis Amari et al. (1996); Choi et al. (1998); Zhang et al. (1999). Thereafter, the natural gradient has also been used for stochastic models Park et al. (2000) and for neural networks Amari (1997); Yang & Amari (1998); Amari (1998). The recently introduced Hessian-Free Optimization techniques Martens (2010); Martens & Sutskever (2011); Chapelle & Erhan (2011); Kiros (2013) and the Krylov Subspace Descent methods Mizutani & Demmel (2003); Vinyals & Povey (2012) use an extended Gauss-Newton approximation of the Hessian, therefore they can be interpreted as versions of natural gradient descent Pascanu & Bengio (2014). The natural gradient was also applied to deep Boltzmann machines Desjardins et al. (2014), reinforcement learning Kakade (2002); Peters et al. (2005); Bhatnagar et al. (2008), and natural evolution strategies Sun et al. (2009); Wierstra et al. (2014).

Since for neural networks and deep learning the computation of the Fisher information matrix is too expensive, different approximations of the natural gradient have been proposed. For example the Fisher matrix can be approximated in a block-diagonal form, where unit-wise or quasi-diagonal natural gradients are used Olivier (2013). Unit-wise natural gradients have been previously used Kurita (1993) and go back to natural gradients for perceptrons Amari (1998); Yang & Amari (1998). Top-

moumoute Online natural Gradient Algorithm (TONGA) LeRoux et al. (2008) uses a low-rank approximation of natural gradient descent. FActorized Natural Gradient (FANG) Grosse & Salakhudinov (2015) estimates the natural gradient via an approximation of the Fisher matrix by a Gaussian graphical model. The natural gradient has been combined with the Newton method for an efficient second order method LeRoux & Fitzgibbon (2010).

We assume a parametrized probabilistic model $p(\boldsymbol{x}; \boldsymbol{w})$ with parameter vector $\boldsymbol{w}$ and data $\boldsymbol{x}$. The training data are $\boldsymbol{X} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N) \in \mathbb{R}^{(d+1) \times N}$ with $\boldsymbol{x}_n = (\boldsymbol{z}_n^T, y_n)^T \in \mathbb{R}^{d+1}$, where $\boldsymbol{z}_n$ is the input for example n and y_n is its label. The loss of one example $\boldsymbol{x} = (\boldsymbol{z}^T, y)^T$ using model $p(.; \boldsymbol{w})$ is $L(y, p(\boldsymbol{z}; \boldsymbol{w}))$ and the average loss on the training data $\boldsymbol{X}$ is the empirical risk $R_{\text{emp}}(p(.; \boldsymbol{w}), \boldsymbol{X})$. Gradient descent updates the weight vector $\boldsymbol{w}$ by $\boldsymbol{w}^{\text{new}} = \boldsymbol{w}^{\text{old}} - \eta \nabla_{\boldsymbol{w}} R_{\text{emp}}$ where η is the learning rate. The *natural gradient* is the inverse Fisher information matrix $\boldsymbol{F}^{-1}$ multiplied by the gradient of the empirical risk: $\nabla_{\boldsymbol{w}}^{\text{nat}} R_{\text{emp}} = \boldsymbol{F}^{-1} \nabla_{\boldsymbol{w}} R_{\text{emp}}$.

For a multi-layer perceptron (MLP) $\boldsymbol{a}$ is the vector of the unit's activity and $a_0 = 1$ is the bias unit activity. We consider the ingoing weights to unit i, therefore we drop the index i: $w_j = w_{ij}$ for the weight from unit j to unit i, $\text{net} = \text{net}_i$ for the net input to unit i, $a = a_i$ for the activation of unit i, and $\delta = \delta_i$ for the delta error at unit i. w_0 is the bias weight of unit i. The net input is $\text{net} = \sum_j w_j a_j$. The activation function f maps the net input to the activation $a = f(\text{net})$. For example $\boldsymbol{x} = (\boldsymbol{z}^T, y)^T$, the delta error at unit i is $\delta = \frac{\partial}{\partial \text{net}} L(y, p(\boldsymbol{z}; \boldsymbol{w}))$, which gives for the derivative $\frac{\partial}{\partial w_j} L(y, p(\boldsymbol{z}; \boldsymbol{w})) = \delta a_j$.

For the Fisher information matrix, we need the gradient of the log output function $\frac{\partial}{\partial w_j} \ln p(\boldsymbol{x}; \boldsymbol{w})$. Since $p(\boldsymbol{x}; \boldsymbol{w}) = p(y \mid \boldsymbol{z}; \boldsymbol{w}) p(\boldsymbol{z})$, only $p(y \mid \boldsymbol{z}; \boldsymbol{w})$ depends on the weights. Therefore we define the $\hat{\delta}$ at unit i as $\hat{\delta} = \frac{\partial}{\partial \text{net}} \ln p(y \mid \boldsymbol{z}; \boldsymbol{w})$, which can be computed via the backpropagation algorithm, but starting from the log-output probability instead of the loss function. For the derivative we obtain $\frac{\partial}{\partial w_j} \ln p(y \mid \boldsymbol{z}; \boldsymbol{w}) = \hat{\delta} a_j$.

We restrict the Fisher information matrix to weights leading to unit i, which is the unit Fisher information matrix $\boldsymbol{F}$. The unit Fisher information matrix $\boldsymbol{F}$ captures only the interactions of weights to unit i. Consequently, the unit natural gradient only corrects the interactions of weights to unit i. In particular the affect of the bias weight update can be adjusted by the unit natural gradient. Unit natural gradients have been previously used Kurita (1993); Olivier (2013) and go back to natural gradients for perceptrons Amari (1998); Yang & Amari (1998). The unit Fisher information matrix is

$$[\boldsymbol{F}(\boldsymbol{w})]_{kj} = \mathrm{E}_{p(\boldsymbol{x}; \boldsymbol{w})} \left(\frac{\partial \ln p(y \mid \boldsymbol{z}; \boldsymbol{w})}{\partial w_k} \frac{\partial \ln p(y \mid \boldsymbol{z}; \boldsymbol{w})}{\partial w_j} \right) = \mathrm{E}_{p(\boldsymbol{x}; \boldsymbol{w})} \left(\hat{\delta}^2 a_k a_j \right) \qquad (1)$$

$$= \mathrm{E}_{p(\boldsymbol{z})} \left(\mathrm{E}_{p(y \mid \boldsymbol{z}; \boldsymbol{w})} \left(\hat{\delta}^2 \right) a_k a_j \right) = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 a_k a_j \right) \qquad , \qquad \bar{\delta}^2 = \mathrm{E}_{p(y \mid \boldsymbol{z}; \boldsymbol{w})} \left(\hat{\delta}^2 \right) .$$

Weighting the activations by $\bar{\delta}^2$ is equivalent to adjusting the probability of drawing inputs $\boldsymbol{z}$. $\boldsymbol{z}$ with large $\bar{\delta}^2$ a drawn with higher probability and and inputs $\boldsymbol{z}$ with small $\bar{\delta}^2$ with smaller probability. Since $0 \leq \bar{\delta}^2 = \bar{\delta}^2(\boldsymbol{z})$, we can define a distribution $q(\boldsymbol{z})$:

$$q(\boldsymbol{z}) = \bar{\delta}^2(\boldsymbol{z}) \, p(\boldsymbol{z}) \left(\int \bar{\delta}^2(\boldsymbol{z}) \, p(\boldsymbol{z}) \, d\boldsymbol{z} \right)^{-1} = \bar{\delta}^2(\boldsymbol{z}) \, p(\boldsymbol{z}) \, \mathrm{E}_{p(\boldsymbol{z})}^{-1} \left(\bar{\delta}^2 \right) . \qquad (2)$$

Using the distribution $q(\boldsymbol{z})$, the entries of the Fisher information matrix can be expressed as a covariance:

$$[\boldsymbol{F}(\boldsymbol{w})]_{kj} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 a_k a_j \right) = \int \bar{\delta}^2 a_k a_j \, p(\boldsymbol{z}) \, d\boldsymbol{z} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right) \mathrm{E}_{q(\boldsymbol{z})} \left(a_k a_j \right) . \qquad (3)$$

If the bias unit is unit $j = 0$ with $a_0 = 1$ and weight w_0 then we can divide the complete weight vector $\boldsymbol{w}^+$ into a bias part w_0 and the rest $\boldsymbol{w}$: $\boldsymbol{w}^+ = (\boldsymbol{w}^T, w_0)^T$. For the row $\boldsymbol{b}$ of the Fisher information matrix that is associated with the bias, we have for $\boldsymbol{b} = [\boldsymbol{F}(\boldsymbol{w})]_0$:

$$\boldsymbol{b} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \boldsymbol{a} \right) = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right) \mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right) = \mathrm{Cov}_{p(\boldsymbol{z})} \left(\bar{\delta}^2, \boldsymbol{a} \right) + \mathrm{E}_{p(\boldsymbol{z})} \left(\boldsymbol{a} \right) \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right) . \qquad (4)$$

The next Theorem 1 gives the correction of the standard gradient by the unit natural gradient where the bias weight is treated separately. The resulting formulas are similar to Definition 6 of Olivier (2013).

Theorem 1. *The unit natural gradient corrects the weight update to a unit $(\Delta \boldsymbol{w}^T, \Delta w_0)^T$ by following affine transformation of the gradient $\nabla_{(\boldsymbol{w}, w_0)} R_{\text{emp}} = (\boldsymbol{g}^T, g_0)^T$:*

$$\begin{pmatrix} \Delta \boldsymbol{w} \\ \Delta w_0 \end{pmatrix} = \begin{pmatrix} \boldsymbol{A}^{-1} \left(\boldsymbol{g} - \Delta w_0 \, \boldsymbol{b} \right) \\ s \left(g_0 - \boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{g} \right) \end{pmatrix} , \tag{5}$$

where $\boldsymbol{A}$ is the unit Fisher information matrix without the bias weight. The vector $\boldsymbol{b} = [\boldsymbol{F}(\boldsymbol{w})]_0$ is the unit Fisher matrix column corresponding to the bias weight, and the scalar s is

$$s = \mathrm{E}_{p(\boldsymbol{z})}^{-1} \left(\bar{\delta}^2 \right) \left(1 + \mathrm{E}_{q(\boldsymbol{z})}^T \left(\boldsymbol{a} \right) \mathrm{Var}_{q(\boldsymbol{z})}^{-1} \left(\boldsymbol{a} \right) \mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right) \right) , \tag{6}$$

where $\boldsymbol{a}$ is the vector of activations of units with weights to unit i and

$$q(\boldsymbol{z}) = \bar{\delta}^2(\boldsymbol{z}) \, p(\boldsymbol{z}) \, \mathrm{E}_{p(\boldsymbol{z})}^{-1} \left(\bar{\delta}^2 \right) \quad , \quad \bar{\delta}^2(\boldsymbol{z}) = \mathrm{E}_{p(y|\boldsymbol{z};\boldsymbol{w})} \left(\hat{\delta}^2 \right) . \tag{7}$$

Proof. We separate the local Fisher matrix $\boldsymbol{F}$ of unit i with W incoming weights $\boldsymbol{w}^+ = (\boldsymbol{w}^T, w_0)^T$ into a bias part given by w_0 and the rest $\boldsymbol{w}$. Multiplying the inverse Fisher matrix with the separated gradient $\nabla_{\boldsymbol{w}^+} R_{\text{emp}}(\boldsymbol{w}^+, \boldsymbol{X}) = \boldsymbol{g}^+ = (\boldsymbol{g}^T, g_0)^T$ gives the weight update $\Delta \boldsymbol{w}^+ = (\Delta \boldsymbol{w}^T, \Delta w_0)^T$:

$$\begin{pmatrix} \Delta \boldsymbol{w} \\ \Delta w_0 \end{pmatrix} = \begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{b}^T & c \end{pmatrix}^{-1} \begin{pmatrix} \boldsymbol{g} \\ g_0 \end{pmatrix} = \begin{pmatrix} \boldsymbol{A}^{-1} \boldsymbol{g} + \boldsymbol{u} \, s^{-1} \boldsymbol{u}^T \boldsymbol{g} + g_o \, \boldsymbol{u} \\ \boldsymbol{u}^T \boldsymbol{g} + s \, g_0 \end{pmatrix} . \tag{8}$$

where

$$\boldsymbol{b} = [\boldsymbol{F}(\boldsymbol{w})]_0 \quad , \quad c = [\boldsymbol{F}(\boldsymbol{w})]_{00} \quad , \quad \boldsymbol{u} = -s \, \boldsymbol{A}^{-1} \boldsymbol{b} \quad , \quad s = \left(c - \boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} \right)^{-1} . \tag{9}$$

This formula is derived in Lemma 1. When inserting Δw_0 in the update of $\Delta \boldsymbol{w}$, we obtain

$$\begin{pmatrix} \Delta \boldsymbol{w} \\ \Delta w_0 \end{pmatrix} = \begin{pmatrix} \boldsymbol{A}^{-1} \boldsymbol{g} + s^{-1} \boldsymbol{u} \, \Delta w_0 \\ \boldsymbol{u}^T \boldsymbol{g} + s \, g_0 \end{pmatrix} \quad , \quad \begin{pmatrix} \Delta \boldsymbol{w} \\ \Delta w_0 \end{pmatrix} = \begin{pmatrix} \boldsymbol{A}^{-1} \left(\boldsymbol{g} - \Delta w_0 \, \boldsymbol{b} \right) \\ s \left(g_0 - \boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{g} \right) \end{pmatrix} . \tag{10}$$

The right hand side is obtained by inserting $\boldsymbol{u} = -s \boldsymbol{A}^{-1} \boldsymbol{b}$ in the left hand side update. Since $c = F_{00} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right)$, $\boldsymbol{b} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right) \mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$, and $\boldsymbol{A} = \mathrm{E}_{p(\boldsymbol{z})} \left(\bar{\delta}^2 \right) \mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \boldsymbol{a}^T \right)$, we obtain

$$s = \left(c - \boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{b} \right)^{-1} = \mathrm{E}_{p(\boldsymbol{z})}^{-1} \left(\bar{\delta}^2 \right) \left(1 - \mathrm{E}_{q(\boldsymbol{z})}^T \left(\boldsymbol{a} \right) \mathrm{E}_{q(\boldsymbol{z})}^{-1} \left(\boldsymbol{a} \boldsymbol{a}^T \right) \mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right) \right)^{-1} . \tag{11}$$

Applying Lemma 2 gives the formula for s. $\qquad \square$

Corollary 1. *The weight update correction of the unit natural gradient for the bias weight increases monotonically with the length of the vector $\mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$.*

Proof. The weight update correction is given in Theorem 1. The larger the norm of $\mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$ is, the larger is the multiplicative correction s of the bias weight update. The value s is a positive constant plus a quadratic form of the inverse of the positive definite covariance matrix $\mathrm{Var}_{q(\boldsymbol{z})}^{-1}(\boldsymbol{a})$ in Theorem 1. The additive bias correction of the bias weight update is the inner product between $\boldsymbol{w}$'s natural gradient update $\boldsymbol{A}^{-1} \boldsymbol{g}$ and $\boldsymbol{b}$. Since $\boldsymbol{b}$ is a multiple of $\mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$, the inner product also increases with the length of $\mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$. The correction for $\Delta \boldsymbol{w}$ is directly proportional to both $\boldsymbol{b}$, therefore also to $\mathrm{E}_{q(\boldsymbol{z})} \left(\boldsymbol{a} \right)$, and Δw_0. $\qquad \square$

We compute the *bias shift* of unit i caused by the weight update, i.e. the change of the mean of unit i. Towards this end we compute the average change of the net input of unit i. Since only the activations depend on the input, the bias shift is given by multiplying the weight update by the mean of the activation vector $\boldsymbol{a}$, where we use $\Delta w_0 = -s \, \boldsymbol{b}^T \boldsymbol{A}^{-1} \boldsymbol{g} + s \, g_0$:

$$\begin{pmatrix} \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a}) \\ 1 \end{pmatrix}^T \begin{pmatrix} \Delta \boldsymbol{w} \\ \Delta w_0 \end{pmatrix} = \begin{pmatrix} \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a}) \\ 1 \end{pmatrix}^T \begin{pmatrix} \boldsymbol{A}^{-1} \boldsymbol{g} - \boldsymbol{A}^{-1} \boldsymbol{b} \, \Delta w_0 \\ \Delta w_0 \end{pmatrix} \tag{12}$$

$$= \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{g} + \left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b} \right) \Delta w_0$$

$$= \left(\mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T - \left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b} \right) s \, \boldsymbol{b}^T \right) \boldsymbol{A}^{-1} \boldsymbol{g} + s \left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b} \right) g_0 .$$

In Eq. (12) the term $\left(\mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T - \left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) s \boldsymbol{b}^T\right) \boldsymbol{A}^{-1} \boldsymbol{g}$ is the *bias shift* due weight updates of connections from the previous layer while $s\left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) g_0$ is the shift of the mean value of unit i due to the regular bias weight update. The term $\left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) s \boldsymbol{b}$ is the bias shift correction due to the unit natural gradient.

Next, Theorem 2 states that the bias shift of unit i is mitigated or even prevented by the bias shift correction. Furthermore, the theorem shows that the bias shift correction is goverened by the size of $\mathrm{Cov}_{p(\boldsymbol{z})}\left(\bar{\delta}^2, \boldsymbol{a}\right)$.

Theorem 2. *The bias shift correction of unit i by the unit natural gradient update is equivalent to following correction of the incoming mean $\mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})$:*

$$- (1 + k)\, \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\,, \quad \text{where} \quad k = s\, \mathrm{E}_{p(\boldsymbol{z})}\left(\bar{\delta}^2\right)\, \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\,. \tag{13}$$

For $\mathrm{Cov}_{p(\boldsymbol{z})}\left(\bar{\delta}^2, \boldsymbol{a}\right) = \boldsymbol{0}$, the incoming mean $\mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})$ is corrected to zero.

Proof. We analyze the bias shift correction term defined in Eq. (12), starting with its first factor:

$$1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b} = 1 - \mathrm{E}_{p(\boldsymbol{z})}\left(\bar{\delta}^2\right) \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a}) \tag{14}$$

$$= 1 - \left(\mathrm{E}_{p(\boldsymbol{z})}\left(\bar{\delta}^2\right) \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})^T - \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right)\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})$$

$$= 1 - \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})^T \mathrm{E}_{q(\boldsymbol{z})}^{-1}(\boldsymbol{a}\boldsymbol{a}^T) \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a}) + \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})$$

$$= \mathrm{E}_{p(\boldsymbol{z})}^{-1}\left(\bar{\delta}^2\right) s^{-1} + \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\,,$$

where we used Eq. (4). Next, we reformulate the full bias shift correction term using previous equation:

$$s\left(1 - \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})^T \boldsymbol{A}^{-1} \boldsymbol{b}\right) \boldsymbol{b} = \left(\mathrm{E}_{p(\boldsymbol{z})}^{-1}\left(\bar{\delta}^2\right) + s\, \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\right) \boldsymbol{b} \tag{15}$$

$$= \left(1 + s\, \mathrm{E}_{p(\boldsymbol{z})}\left(\bar{\delta}^2\right) \mathrm{Cov}_{p(\boldsymbol{z})}^T\left(\bar{\delta}^2, \boldsymbol{a}\right) \boldsymbol{A}^{-1} \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\right) \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a}) = (1 + k)\, \mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a})\,.$$

For $\mathrm{Cov}_{p(\boldsymbol{z})}\left(\bar{\delta}^2, \boldsymbol{a}\right) = \boldsymbol{0}$ we have $k = 0$ and $\mathrm{E}_{q(\boldsymbol{z})}(\boldsymbol{a}) = \mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})$ (see Eq. (4)), therefore the mean incoming activation vector $\mathrm{E}_{p(\boldsymbol{z})}(\boldsymbol{a})$ is corrected to a zero mean vector by the unit natural gradient. $\qquad \square$

If the incoming units have non-zero means of the same sign, then the mean of unit i is likely to be shifted via the weight updates depending on the unit natural gradient $\boldsymbol{A}^{-1}\boldsymbol{g}$. The unit natural gradient corrects this bias shift of unit i by the bias weight to ensure optimal learning. This correction can be viewed as shifting the mean activation of the incoming units towards zero. Without the correction, bias shift of unit i lead to oscillations and impede learning progress. Therefore small absolute means of the activations are desired for efficient learning. Small means may be achieved by (i) unit zero mean normalization LeCun et al. (1991; 1998); Schraudolph (1998); Cho et al. (2011); Raiko et al. (2012); Montavon & Müller (2012); Ioffe & Szegedy (2015); Desjardins et al. (2015) or (ii) activation functions with a negative part Maas et al. (2013); Xu et al. (2015); He et al. (2015). We introduce new activation functions with a negative part.

3 EXPONENTIAL LINEAR UNITS (ELUS)

We introduce the *exponential linear unit* (ELU) with $0 < \alpha$:

$$f(x) = \begin{cases} x & \text{if } x \geq 0 \\ \alpha\left(\exp(x) - 1\right) & \text{if } x < 0 \end{cases}, \tag{16}$$

$$f'(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ f(x) + \alpha & \text{if } x < 0 \end{cases}. \tag{17}$$

The ELU hyperparameter α controls the value to which an ELU saturate for negative net inputs (see Fig. 1). ELUs avoid the vanishing gradient problem as rectified linear units (ReLUs) and leaky ReLUs (LReLUs) do. The gradient does not vanish because the positive part of these functions is the identity, therefore their derivative is one. Thus, these activation functions are well suited for deep neural networks with many layers where a vanishing gradient impedes learning Unterthiner et al. (2014; 2015).

In contrast to ReLUs, ELUs have negative values which pushes the mean of the activations closer to zero. Mean activations that are closer to zero enable faster learning as they bring the gradient closer to the natural gradient (see Corollary 1 and Theorem 2). For the same reasons tanh was previously found to speed up learning compared to the logistic function LeCun et al. (1991; 1998). While LReLUs and PReLUs have also negative values, they do not ensure a noise-robust deactivation state, that is, a noise-robust negative value. ELUs saturate to a negative value for negative net inputs. Saturation means a small derivative of the activation function which decreases the variation and the information that is propagated to the next layer Hochreiter et al. (2001b); Hochreiter (1997; 1998). ELUs code the degree of presence of

Figure 1: Different activation functions. The rectified linear unit (ReLU), the leaky ReLU (LReLU) for $\alpha = 0.1$, and the exponential linear unit (ELU) for $\alpha = 1.0$ are depicted.

input concepts, while they do not quantify the degree of their absence Clevert et al. (2015). Dependencies between ELUs are easier to model than between LReLUs or PReLUs because the causes for deactivations are not distinguished. Furthermore ELUs that code distinct concepts are less likely to interfere because uninformative negative values avoid distributed codes. Positively activated ELUs interact by activating units of the next layer.

4 EXPERIMENTS USING ELUS

In this section, we assess the performance of exponential linear units (ELUs) if used for unsupervised and supervised learning of deep autoencoders and deep convolutional networks, by comparing them with (i) Rectified Linear Units (ReLUs) and (ii) Leaky Linear Units (LReLUs). For comparisons we use the following benchmark datasets to compare ELUs with other units: (i) *MNIST* (gray images in 10 classes, 60k training and 10k testing images), (ii) *CIFAR-10* (color images in 10 categories, 50k training and 10k testing images), (iii) *CIFAR-100* (color images in 100 categories, 50k training and 10k testing images), and (iv) *ImageNet* (color images in 1000 categories, 1.3M training and 100k testing images).

4.1 MNIST

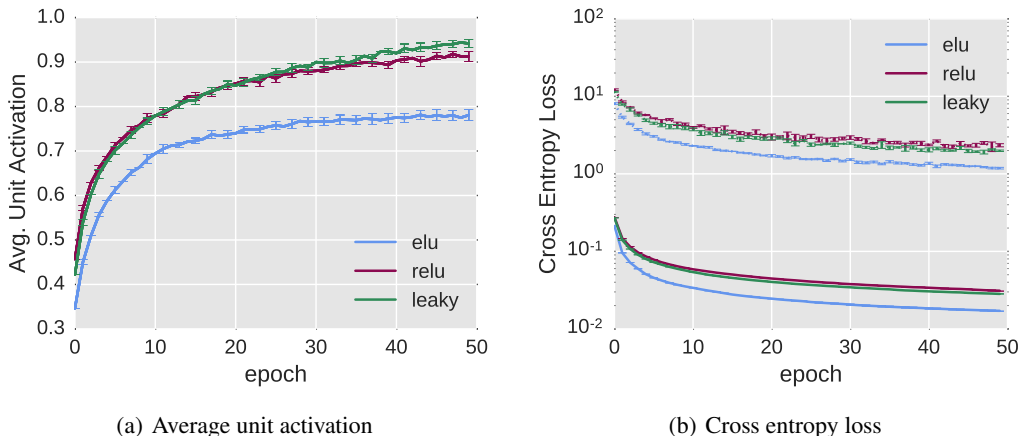(a) Average unit activation

(b) Cross entropy loss

Figure 2: ELU networks evaluated at MNIST. Lines are the average over five runs with different random initializations, error bars show standard deviation. Panel (a): median of the average unit activation for different activation functions. Panel (b): Training set (straight line) and validation set (dotted line) cross entropy loss.

4.1.1 Learning Behavior

We first established that ELUs keep the mean activation closer to zero by training fully connected deep neural networks on the MNIST digit classification dataset and tracking each hidden unit's activation during training. We did this for ELUs, ReLUs as well as Leaky ReLU units with a negative slope of 0.1. Each network had five hidden layers of 256 units each, and was trained by stochastic gradient descent using a learning rate of 0.01 with mini-batches of size 64 was trained for 50 epochs. The ELU hyperparameter α was set to 1.0 for all experiments. The weights have been initialized according to He et al. (2015), which has been specifically designed to speed up learning with ReLU units. After each epoch we calculated the units' average activations on a fixed subset of the training data. We plotted the median over all units in Fig. 2, which shows that ELUs stay closer to zero throughout the training process, while their training error decreases much more rapidly than it does for the other activation functions.

4.1.2 AutoEncoder Learning

To evaluate how well ELUs work in an unsupervised setting. Therefore, we replicated the experiments from Martens (2010) and Desjardins et al. (2015), where a deep autoencoder was trained on the MNIST dataset, with the objective to minimize the reconstruction error. The encoder part of the network consisted of four fully connected hidden layers with sizes 1000, 500, 250 and 30, respectively. The decoder part was symmetrical to the encoder. We performed the experiments using four different learning rates $(10^{-2}, 10^{-3}, 10^{-4}, 10^{-5})$. Learning was based on stochastic gradient descent with mini-batches of 64 samples for 500 epochs. Fig. 3 shows, that ELUs outperform the competing activation functions in terms of training set reconstruction error for all learning rates. We also note that similar to Desjardins et al. (2015), higher learning rates always seem to perform better.
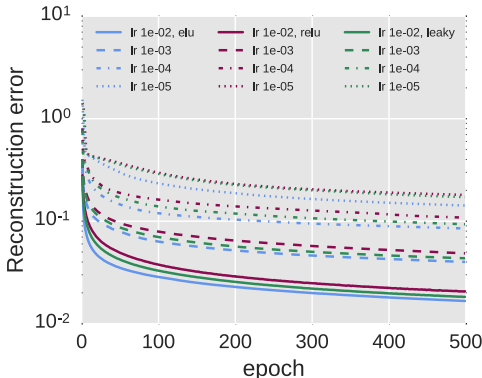

Figure 3: Autoencoder training on MNIST: training set reconstruction error over epochs, using different activation functions and learning rates. The results are medians over several runs with different random initializations.

4.2 Comparison of Activation Functions

In this subsection, we conducted a benchmark experiment using a relatively simple deep convolutional network architecture in order to compare ELUs to ReLUs and LReLUs based on their learning behavior.

The model used for our baseline CIFAR-100 experiment consists of 11 convolutional layers arranged in stacks of $\{1 \times 192 \times 5, 2 \times 240 \times 3, 2 \times 260 \times 2, 2 \times 280 \times 2, 2 \times 300 \times 2, 1 \times 300 \times 1, 1 \times 100 \times 1\}$ layers $\times$ units $\times$ receptive fields. 2×2 max-pooling with a stride of 2 was applied after each stack. For network regularization we used the following drop-out rate for the last layer of each stack $\{0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.0\}$. The $L2$-weight decay regularization term was set to 0.0005. The following learning rate schedule was applied $\{0 - 35k(0.03), 35k - 85k(0.015), 85k - 135k(0.0015), 135k - 165k(0.00015)\}$ (iterations (learning rate)). The momentum term was fixed at 0.9. The dataset was preprocessed as described in Goodfellow et al. (2013) with global contrast normalization and ZCA whitening. Additionally, the images were padded with four zero pixels at all borders. The model was trained on 32×32 random crops with random horizontal flipping. Besides that, we no further augmented the dataset during training.

Fig. 4 shows, that ELUs achieve better training loss and test error than ReLUs or LReLUs. Notice that for the sake of comparison, we used the same step-wise learning rate schedule for all activation functions, which was optimized for ReLU networks during previous experiments. Fig. 4 suggests that ELUs converge faster than ReLU and LReLU networks and therefore would benefit from a better designed schedule adjusted to ELUs. The current schedule was optimized for ReLUs and
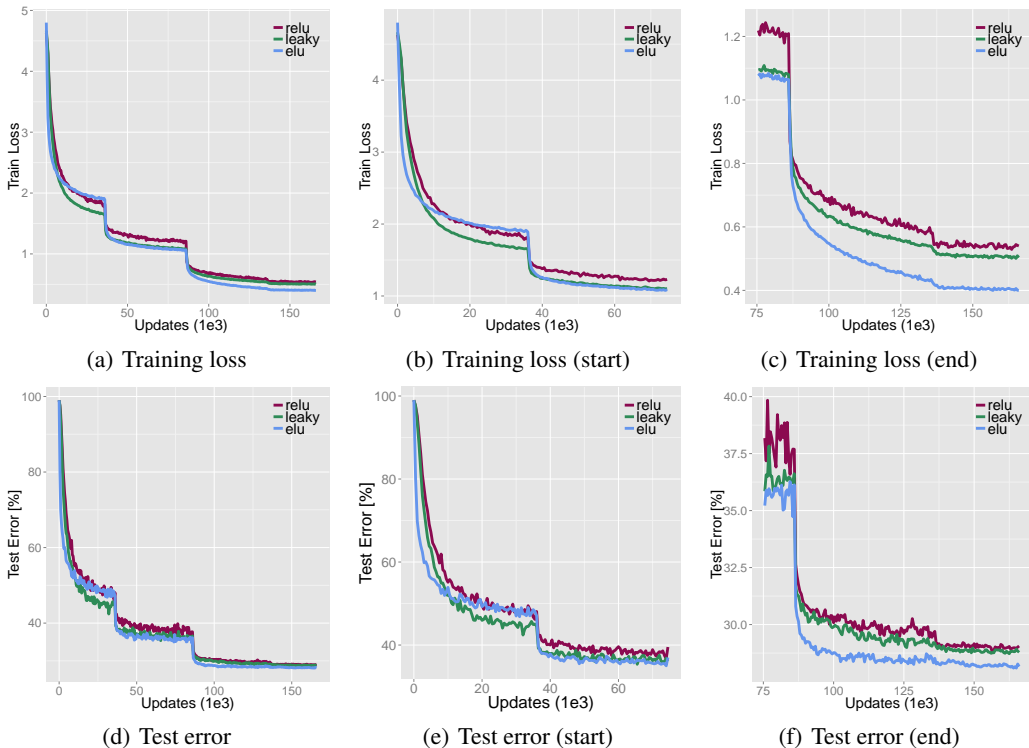
7

Figure 4: The panels (a-c) show the training loss while the lower panels (d-f) show the classification error on CIFAR-100. On CIFAR-100, ELUs achieve better test error and lower training loss. For the sake of comparison, we used the same step-wise learning rate schedule for this experiment.

LReLUs. ELUs yield a test error of 28.16%, while ReLUs and LReLUs yield 29.06% and 28.78%, respectively. In summary, ELUs outperform ReLUs and LReLUs in terms of training loss and test error.

4.3 CLASSIFICATION PERFORMANCE ON CIFAR-100 AND CIFAR-10

After showing that ELUs indeed posses the superior learning behavior as postulated in Section 3, the following experiments were meant to highlight their generalization capabilities. Therefore, the networks in this section have a more sophisticated architecture than the networks in the previous subsection. The architecture of the deep convolution network consists of 18 convolutional layers arranged in stacks of $\{1 \times 384 \times 3, 3 \times 640 \times 3, 4 \times 768 \times 3, 3 \times 896 \times 3, 3 \times 1024 \times 3, 3 \times 1152 \times 3, 1 \times 100 \times 1\}$ layers $\times$ units $\times$ receptive fields. 2×2 max-pooling with a stride of 2 was applied after each stack. For network regularization we used the following drop-out rate for all layers in a stack $\{0.0, 0.15, 0.3, 0.45, 0.6, 0.8, 0.0\}$. The $L2$-weight decay regularization term was set to 0.0005. The initial learning rate was set to 0.01 and decreased by a factor of 10 after 35k iterations. The momentum term was fixed as 0.9 and the mini-batch size is fixed as 50. For both datasets we applied the same preprocessing, padding and cropping as described in the previous section.

A comparison of ELU networks and other convolutional networks on CIFAR-10 and CIFAR-100 is given in Tab.1. On both datasets, ELU-networks performed best and reached a test error of 6.55% and 24.28% for CIFAR-10 and CIFAR-100. ELU networks were among top 10 reported CIFAR-10 results and yielded the best published result on CIFAR-100, without resorting to multi-view evaluation or model averaging.

Table 1: Comparison of ELU networks and other convolutional networks on CIFAR-10 and CIFAR-100. Reported is the test error in percent misclassification for ELU networks and recent convolutional architectures like AlexNet Krizhevsky et al. (2012), DSN Lee et al. (2015), NiN Lin et al. (2013), Maxout Goodfellow et al. (2013), All-CNN Springenberg et al. (2014), Highway Network Srivastava et al. (2015). Best results are in bold.

Network	CIFAR-10 (test error %)	CIFAR-100 (test error %)	augmented
AlexNet	18.04	45.80	
DSN	7.97	34.57	√
NiN	8.81	35.68	√
Maxout	9.38	38.57	√
All-CNN	7.25	33.71	√
Highway Network	7.60	32.24	√
ELU-Network	**6.55**	**24.28**	

4.3.1 IMAGENET CHALLENGE DATASET

In the final experiment, we evaluated ELU-networks on the 1000-class ImageNet dataset, which contains about 1.3M training color images as well as additional 50k images and 100k images for validation and testing, respectively. For this task, we designed a 15 layer deep convolutional network, which was arranged in stacks of $\{1 \times 96 \times 6, 3 \times 512 \times 3, 5 \times 768 \times 3, 3 \times 1024 \times 3, 2 \times 4096 \times FC, 1 \times 1000 \times FC\}$ layers × units × receptive fields or fully-connected (FC). 2×2 max-pooling with a stride of 2 was applied after each stack and spatial pyramid pooling (SPP) with 3 levels before the first FC layer He et al. (2015). For network regularization we set the $L2$-weight decay term to 0.0005 and used 50% drop-out in the two penultimate FC layers. Before training we re-sized the images to 256×256 pixels, subtracted the per-pixel mean. Training was done on 224×224 random crops with random horizontal flipping. Besides that, we didn't augment the dataset during training. The single-model performance was evaluated on the single center crop with no further

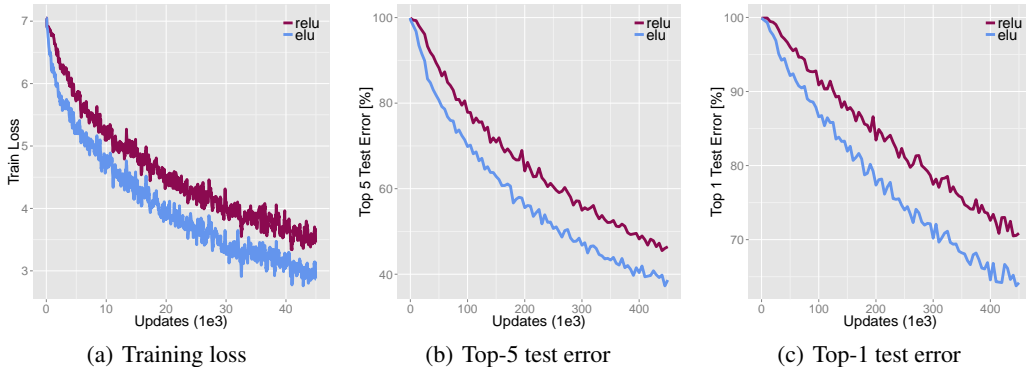(a) Training loss   (b) Top-5 test error   (c) Top-1 test error

Figure 5: ELUs used for the ImageNet classification task. The x-axis gives the number of iterations. The y-axis shows the training loss (a), top-5 error (b) and the top-1 error (c) of 5,000 random validation samples, evaluated on the center crop. Both activation functions ELU (purple) and ReLU (blue) lead for convergence, but ELUs start reducing the error earlier and reach the 50% top-5 error after 26k iterations, while ReLUs need 38k iterations to reach the same error rate.

augmentation and yielded a top-5 validation error below 10%. Fig. 5 shows the learning behavior of our model for ELUs and ReLUs. Panel (b) suggests that ELUs start reducing the error earlier. The ELU-network already reaches the 50% top-5 error after 26k iterations, while the ReLU network needs 38k iterations to reach the same error rate.

5 CONCLUSION

In this paper we introduced the novel *exponential linear units*, units with a nonlinear activation function for deep learning. These units have negative values, which allows the network to push the mean activation of its units closer to zero. This helps to bridge the gap between the normal gradient and the unit natural gradient, thereby speeding up learning. We believe that this property is also the reason for the success of activation functions like LReLUs and PReLUs. In contrast to these methods, ELUs have a clear saturation plateau in its negative regime, allowing them to learn a more robust and stable representation. Experimental results show that ELUs outperform other activation functions on many vision datasets, achieving one of the top 10 best reported results on CIFAR-10 and setting a new state of the art in CIFAR-100 without the need for multi-view test evaluation or model averaging. Furthermore, ELU networks produced competitive results on the ImageNet in much fewer epochs than a corresponding ReLU network.

In our experience ELUs are most effective once the number of layers in a network is larger than 4. For such networks, ELUs consistently outperform ReLUs and its variants with negative slopes. On ImageNet we observed that ELUs are able to converge to a state of the art solution in much less time it takes comparable ReLU networks. Given their outstanding performance, we expect ELU networks to become a real time saver in convolutional networks, which are notably time-intensive to train from scratch otherwise.

REFERENCES

Amari, S. *Differential-Geometrical Methods in Statistic*. Springer, New York, 1985.

Amari, S. and Murata, N. Statistical theory of learning curves under entropic loss criterion. *Neural Computation*, 5(1):140–153, 1993.

Amari, S., Cichocki, A., and Yang, H. H. A new learning algorithm for blind signal separation. In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E. (eds.), *Advances in Neural Information Processing Systems 8 (NIPS8)*, pp. 757–763. The MIT Press, Cambridge, MA, 1996.

Amari, S.-I. Neural learning in structured parameter spaces — natural Riemannian gradient. In Mozer, M. C., Jordan, M. I., and Petsche, T. (eds.), *Advances in Neural Information Processing Systems 9 (NIPS9)*, volume 9, pp. 127–133. The MIT Press, 1997.

Amari, S.-I. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.

Bhatnagar, S., Ghavamzadeh, M., Lee, M., and Sutton, R. S. Incremental natural actor-critic algorithms. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T. (eds.), *Advances in Neural Information Processing Systems 20 (NIPS20)*, pp. 105–112. MIT Press, 2008.

Chapelle, O. and Erhan, D. Improved preconditioner for Hessian free optimization. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. URL `olivier.chapelle.cc/pub/precond.pdf`.

Cho, K., Raiko, T., and Ilin, A. Enhanced gradient and adaptive learning rate for training restricted Boltzmann machines. In Getoor, L. and Scheffer, T. (eds.), *Proceedings of the 28th International Conference on Machine Learning (ICML11)*, pp. 105–112, 2011.

Choi, S., Amari, S., Cichocki, A., and Liu, R. Natural gradient learning with a nonholonomic constraints for blind deconvolution of multiple channels. In *Proceeding of IEEE*, 1998.

Clevert, D.-A., Unterthiner, T., Mayr, A., and Hochreiter, S. Rectified factor networks. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc., 2015.

Dahl, G. E., Sainath, T. N., and Hinton, G. E. Improving deep neural networks for LVCSR using rectified linear units and dropout. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8609–8613. IEEE, 2013.

Desjardins, G., Pascanu, R., Courville, A., and Bengio, Y. Metric-free natural gradient for joint-training of Boltzmann machines. In *International Conference on Learning Representations 2013*, 2014. URL `http://arxiv.org/abs/1301.3545`. arXiv:1301.3545.

Desjardins, G., Simonyan, K., Pascanu, R., and Kavukcuoglu, K. Natural neural networks. *CoRR*, abs/1507.00210, 2015. URL `http://arxiv.org/abs/1507.00210`.

Glorot, X., Bordes, A., and Bengio, Y. Deep sparse rectifier neural networks. In Gordon, G., Dunson, D., and Dudk, M. (eds.), *JMLR W&CP: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS 2011)*, volume 15, pp. 315–323, 2011.

Goodfellow, I. J., Warde-Farley, D., Mirza, M., Courville, A., and Bengio, Y. Maxout networks. *ArXiv e-prints*, 2013.

Grosse, R. and Salakhudinov, R. Scaling up natural gradient by sparsely factorizing the inverse Fisher matrix. *Journal of Machine Learning Research*, 37:2304–2313, 2015. URL `http://jmlr.org/proceedings/papers/v37/grosse15.pdf`. Proceedings of the 32nd International Conference on Machine Learning (ICML15).

He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.

Hochreiter, S. Untersuchungen zu dynamischen neuronalen Netzen. Diploma thesis, Institut für Informatik, Lehrstuhl Prof. Brauer, Technische Universität München, 1991. See www7.informatik.tu-muenchen.de/˜hochreit.

Hochreiter, S. Recurrent neural net learning and vanishing gradient. In Freksa, C. (ed.), *Proceedings in Artificial Intelligence — Fuzzy-Neuro-Systeme 97*, pp. 130–137. INFIX, Sankt Augustin, Germany, 1997.

Hochreiter, S. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(2):107–116, 1998.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. Extended version available in WWW homepages of Hochreiter and Schmidhuber.

Hochreiter, S., Bengio, Y., Frasconi, P., and Schmidhuber, J. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In Kremer and Kolen (eds.), *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press, 2001a.

Hochreiter, S., Younger, A. S., and Conwell, P. R. Learning to learn using gradient descent. In Dorffner, G., Bischof, H., and Hornik, K. (eds.), *International Conference on Artificial Neural Networks (ICANN)*, volume 2130 of *Lecture Notes in Computer Science*, pp. 87–94. Springer, 2001b.

Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Journal of Machine Learning Research*, 37:448–456, 2015. URL `http://jmlr.org/proceedings/papers/v37/ioffe15.pdf`. Proceedings of the 32nd International Conference on Machine Learning (ICML15).

Kakade, S. M. A natural policy gradient. In Dietterich, T. G., Becker, S., and Ghahramani, Z. (eds.), *Advances in Neural Information Processing Systems 14 (NIPS14)*, pp. 1531–1538. MIT Press, 2002.

Kiros, R. Training neural networks with stochastic Hessian-free optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2013. URL `http://arxiv.org/pdf/1301.3641v3`. arXiv:1301.3641.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 25*, pp. 1097–1105. Curran Associates, Inc., 2012.

Kurita, T. Iterative weighted least squares algorithms for neural networks classifiers. In *Proceedings of the Third Workshop on Algorithmic Learning Theory (ALT92)*, volume 743 of *Lecture Notes in Computer Science*, pp. 77–86. Springer, 1993.

LeCun, Y., Kanter, I., and Solla, S. A. Eigenvalues of covariance matrices: Application to neural-network learning. *Physical Review Letters*, 66(18):2396–2399, 1991.

LeCun, Y., Bottou, L., Orr, G. B., and Müller, K.-R. Efficient backprop. In Orr, G. B. and Müller, K.-R. (eds.), *Neural Networks: Tricks of the Trade*, volume 1524 of *Lecture Notes in Computer Science*, pp. 9–50. Springer, 1998.

Lee, Chen-Yu, Xie, Saining, Gallagher, Patrick W., Zhang, Zhengyou, and Tu, Zhuowen. Deeply-supervised nets. In *AISTATS*, 2015.

LeRoux, N. and Fitzgibbon, A. W. A fast natural Newton method. In Fürnkranz, J. and Joachims, T. (eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML10)*, pp. 623–630, 2010.

LeRoux, N., Manzagol, P.-A., and Bengio, Y. Topmoumoute online natural gradient algorithm. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T. (eds.), *Advances in Neural Information Processing Systems 20 (NIPS)*, pp. 849–856, 2008.

Lin, Min, Chen, Qiang, and Yan, Shuicheng. Network in network. *CoRR*, abs/1312.4400, 2013. URL `http://arxiv.org/abs/1312.4400`.

Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the 30th International Conference on Machine Learning (ICML13)*, 2013.

Martens, J. Deep learning via Hessian-free optimization. In Fürnkranz, J. and Joachims, T. (eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML10)*, pp. 735–742, 2010.

Martens, J. and Sutskever, I. Learning recurrent neural networks with Hessian-free optimization. In Getoor, L. and Scheffer, T. (eds.), *Proceedings of the 28th International Conference on Machine Learning (ICML11)*,

pp. 1033–1040, 2011.

Mizutani, E. and Demmel, J. Iterative scaled trust-region learning in Krylov subspaces via Pearlmutter's implicit sparse Hessian-vector multiply. In Thrun, S., Saul, L. K., and Schölkopf, B. (eds.), *Advances in Neural Information Processing Systems 16 (NIPS)*, pp. 209–216. MIT Press, 2003.

Montavon, G. and Müller, K.-R. Deep Boltzmann machines and the centering trick. In Montavon, G., Orr, G. B., and Müller, K.-R. (eds.), *Neural Networks: Tricks of the Trade*, volume 7700 of *Lecture Notes in Computer Science*, pp. 621–637. Springer, 2012.

Nair, V. and Hinton, G. E. Rectified linear units improve restricted Boltzmann machines. In Fürnkranz, J. and Joachims, T. (eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML10)*, pp. 807–814, 2010.

Olivier, Y. Riemannian metrics for neural networks i: feedforward networks. *CoRR*, abs/1303.0818, 2013. URL http://arxiv.org/abs/1303.0818.

Park, H., Amari, S.-I., and Fukumizu, K. Adaptive natural gradient learning algorithms for various stochastic models. *Neural Networks*, 13(7):755–764, 2000.

Pascanu, R. and Bengio, Y. Revisiting natural gradient for deep networks. In *International Conference on Learning Representations 2014*, 2014. URL http://arxiv.org/abs/1301.3584. arXiv:1301.3584.

Peters, L., Vijayakumar, S., and Schaal, S. Natural actor-critic. In *Proceedings of the Sixteenth European Conference on Machine Learning*, pp. 280–291, 2005.

Raiko, T., Valpola, H., and LeCun, Y. Deep learning made easier by linear transformations in perceptrons. In Lawrence, N. D. and Girolami, M. A. (eds.), *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS12)*, volume 22, pp. 924–932, 2012.

Schraudolph, N. N. Centering neural network gradient factor. In Orr, G. B. and Müller, K.-R. (eds.), *Neural Networks: Tricks of the Trade*, volume 1524 of *Lecture Notes in Computer Science*, pp. 207–226. Springer, 1998.

Springenberg, Jost Tobias, Dosovitskiy, Alexey, Brox, Thomas, and Riedmiller, Martin A. Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806, 2014. URL http://arxiv.org/abs/1412.6806.

Srivastava, Rupesh Kumar, Greff, Klaus, and Schmidhuber, Jürgen. Training very deep networks. *CoRR*, abs/1507.06228, 2015. URL http://arxiv.org/abs/1507.06228.

Sun, Y., Wierstra, D., Schaul, T., and Schmidhuber, J. Stochastic search using the natural gradient. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML09*, pp. 1161–1168, 2009.

Unterthiner, T., Mayr, A., Klambauer, G., Steijaert, M., Wegner, J. K., Ceulemans, H., and Hochreiter, S. Deep learning as an opportunity in virtual screening. In *Advances in Neural Information Processing Systems 27 (NIPS 2014), Workshop on Deep Learning and Representation Learning*. Inst. of Bioinformatics, 2014. URL www.bioinf.jku.at/publications/2014/NIPS2014a.pdf.

Unterthiner, T., Mayr, A., Klambauer, G., and Hochreiter, S. Toxicity prediction using deep learning. *CoRR*, abs/1503.01445, 2015. URL http://arxiv.org/abs/1503.01445.

Vinyals, O. and Povey, D. Krylov subspace descent for deep learning. In *AISTATS*, 2012. URL http://arxiv.org/pdf/1111.4259v1. arXiv:1111.4259.

Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J., and Schmidhuber, J. Natural evolution strategies. *Journal of Machine Learning Research*, 15:949–980, 2014.

Xu, B., Wang, N., Chen, T., and Li, M. Empirical evaluation of rectified activations in convolutional network. *CoRR*, abs/1505.00853, 2015. URL http://arxiv.org/abs/1505.00853.

Yang, H. H. and Amari, S.-I. Complexity issues in natural gradient descent method for training multilayer perceptrons. *Neural Computation*, 10(8), 1998.

Zeiler, M. D., Ranzato, M., Monga, R., Mao, M., Yang, K., Le, Q. V., Nguyen, P., Senior, A., Vanhoucke, V., Dean, J., and Hinton, G. E. On rectified linear units for speech processing. In *38th International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3517–3521. IEEE, 2013.

Zhang, L., Amari, S.-I., and Cichocki, A. Natural gradient approach to blind separation of over- and under-complete mixtures. In *Proc. First International Conference on Independent Component Analysis and Blind Source Separation ICA99*, pp. 455–460, 1999.

A INVERSE OF BLOCK MATRICES

Lemma 1. *The positive definite matrix $\boldsymbol{M}$ is in block format with matrix $\boldsymbol{A}$, vector $\boldsymbol{b}$, and scalar c. The inverse of $\boldsymbol{M}$ is*

$$\boldsymbol{M}^{-1} = \begin{pmatrix} \boldsymbol{A} & \boldsymbol{b} \\ \boldsymbol{b}^T & c \end{pmatrix}^{-1} = \begin{pmatrix} \boldsymbol{K} & \boldsymbol{u} \\ \boldsymbol{u}^T & s \end{pmatrix}, \tag{18}$$

where

$$K = A^{-1} + u \, s^{-1} u^T \tag{19}$$

$$u = -s \, A^{-1} b \tag{20}$$

$$s = \left(c - b^T A^{-1} b \right)^{-1} . \tag{21}$$

Proof. For block matrices the inverse is

$$\begin{pmatrix} A & B \\ B^T & C \end{pmatrix}^{-1} = \begin{pmatrix} K & U \\ U^T & S \end{pmatrix} , \tag{22}$$

where the matrices on the right hand side are:

$$K = A^{-1} + A^{-1} B \left(C - B^T A^{-1} B \right)^{-1} B^T A^{-1} \tag{23}$$

$$U = -A^{-1} B \left(C - B^T A^{-1} B \right)^{-1} \tag{24}$$

$$U^T = -\left(C - B^T A^{-1} B \right)^{-1} B^T A^{-1} \tag{25}$$

$$S = \left(C - B^T A^{-1} B \right)^{-1} . \tag{26}$$

Further if follows that

$$K = A^{-1} + U \, S^{-1} U^T . \tag{27}$$

We now use this formula for $B = b$ being a vector and $C = c$ a scalar. We obtain

$$\begin{pmatrix} A & b \\ b^T & c \end{pmatrix}^{-1} = \begin{pmatrix} K & u \\ u^T & s \end{pmatrix} , \tag{28}$$

where the right hand side matrices, vectors, and the scalar s are:

$$K = A^{-1} + A^{-1} b \left(c - b^T A^{-1} b \right)^{-1} b^T A^{-1} \tag{29}$$

$$u = -A^{-1} b \left(c - b^T A^{-1} b \right)^{-1} \tag{30}$$

$$u^T = -\left(c - b^T A^{-1} b \right)^{-1} b^T A^{-1} \tag{31}$$

$$s = \left(c - b^T A^{-1} b \right)^{-1} . \tag{32}$$

Again it follows that

$$K = A^{-1} + u \, s^{-1} u^T . \tag{33}$$

A reformulation using u gives

$$K = A^{-1} + u \, s^{-1} u^T \tag{34}$$

$$u = -s \, A^{-1} b \tag{35}$$

$$u^T = -s \, b^T A^{-1} \tag{36}$$

$$s = \left(c - b^T A^{-1} b \right)^{-1} . \tag{37}$$

$\square$

B   QUADRATIC FORM OF MEAN AND INVERSE SECOND MOMENT

Lemma 2. *For a random variable a holds*

$$\mathrm{E}(a)^T \mathrm{E}^{-1}(a \, a^T) \, \mathrm{E}(a) \leq 1 \tag{38}$$

and

$$\left(1 - \mathrm{E}(a)^T \mathrm{E}^{-1}(a \, a^T) \, \mathrm{E}(a) \right)^{-1} = 1 + \mathrm{E}(a)^T \mathrm{Var}^{-1}(a) \, \mathrm{E}(a) . \tag{39}$$

Proof. The Sherman-Morrison Theorem states

$$\left(\boldsymbol{A} + \boldsymbol{b}\,\boldsymbol{c}^T\right)^{-1} = \boldsymbol{A}^{-1} - \frac{\boldsymbol{A}^{-1}\,\boldsymbol{b}\,\boldsymbol{c}^T\,\boldsymbol{A}^{-1}}{1 + \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}}\,. \tag{40}$$

Therefore we have

$$\boldsymbol{c}^T\left(\boldsymbol{A} + \boldsymbol{b}\,\boldsymbol{c}^T\right)^{-1}\boldsymbol{b} = \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b} - \frac{\boldsymbol{c}^T\boldsymbol{A}^{-1}\,\boldsymbol{b}\,\boldsymbol{c}^T\,\boldsymbol{A}^{-1}\boldsymbol{b}}{1 + \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}} \tag{41}$$

$$= \frac{\boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}\left(1 + \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}\right) - \left(\boldsymbol{c}^T\boldsymbol{A}^{-1}\,\boldsymbol{b}\right)^2}{1 + \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}}$$

$$= \frac{\boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}}{1 + \boldsymbol{c}^T\boldsymbol{A}^{-1}\boldsymbol{b}}\,.$$

Using the identity

$$\mathrm{E}(\boldsymbol{a}\,\boldsymbol{a}^T) = \mathrm{Var}(\boldsymbol{a}) + \mathrm{E}(\boldsymbol{a})\,\mathrm{E}(\boldsymbol{a})^T \tag{42}$$

for the second moment and previous formula, we get

$$\mathrm{E}(\boldsymbol{a})^T\mathrm{E}^{-1}(\boldsymbol{a}\,\boldsymbol{a}^T)\,\mathrm{E}(\boldsymbol{a}) = \mathrm{E}(\boldsymbol{a})^T\left(\mathrm{Var}(\boldsymbol{a}) + \mathrm{E}(\boldsymbol{a})\,\mathrm{E}(\boldsymbol{a})^T\right)^{-1}\mathrm{E}(\boldsymbol{a}) \tag{43}$$

$$= \frac{\mathrm{E}(\boldsymbol{a})^T\mathrm{Var}^{-1}(\boldsymbol{a})\,\mathrm{E}(\boldsymbol{a})}{1 + \mathrm{E}(\boldsymbol{a})^T\mathrm{Var}^{-1}(\boldsymbol{a})\,\mathrm{E}(\boldsymbol{a})} \le 1\,.$$

The last inequality follows from the fact that $\mathrm{Var}(\boldsymbol{a})$ is positive definite.

We obtain further

$$\left(1 - \mathrm{E}(\boldsymbol{a})^T\mathrm{E}^{-1}(\boldsymbol{a}\,\boldsymbol{a}^T)\,\mathrm{E}(\boldsymbol{a})\right)^{-1} = 1 + \mathrm{E}(\boldsymbol{a})^T\mathrm{Var}^{-1}(\boldsymbol{a})\,\mathrm{E}(\boldsymbol{a})\,. \tag{44}$$

$\square$