Why Rectified Linear Neurons Are Efficient: Symmetry-Based, Complexity-Based, and Fuzzy-Based Explanations

L. Olac Fuentes[1], Justin Parra[1], Elizabeth Anthony[2], and
Vladik Kreinovich[1]

Departments of [1]Computer Science and [2]Geological Sciences
University of Texas at El Paso
El Paso, Texas 79968, USA
`ofuentes@utep.edu, jrparra2@miners.utep.edu, eanthony@utep.edu,`
`vladik@utep.edu`

Abstract. Traditionally, neural networks used a sigmoid activation function. Recently, it turned out that piecewise linear activation functions are much more efficient – especially in deep learning applications. However, so far, there have been no convincing theoretical explanation for this empirical efficiency. In this paper, we show that, by using different uncertainty techniques, we can come up with several explanations for the efficiency of piecewise linear neural networks. The existence of several different explanations makes us even more confident in our results – and thus, in the efficiency of piecewise linear activation functions.

1 Rectified Linear Neurons: Formulation of the Problem

Why neural networks: a brief reminder. One of the main objectives of designing computers is that they would solve *intelligent* tasks, tasks that we normally solve by using our brains. It is therefore reasonable, when designing computational devices, to emulate how our brain works.

In the brain, signals come from the special sensor cells in the eyes, ears, etc., and are processed by other cells called *neurons*. The signals from the sensors come as series of electric spikes. The intensity of the corresponding signal is reflected by the frequency of the spikes.

Signal processing cells – neurons – usually:

- take inputs from several cells (sensor cells or other data processing neurons),
- process the summary input signal, and
- send the resulting signal to other neurons – or to the cells that perform some activities (e.g., move a finger, close an eye, slow down the heart rate, etc.).

To be more precise, when a neuron gets signals $x_1, \ldots, x_n$ from different inputs:

- these signals are first aggregated into a linear combination

$$x = w_1 \cdot x_1 + \ldots + w_n \cdot x_n + w_0,$$

and then

– an appropriate transformation $y = s_0(x)$ is applied to the aggregated signal x.

As a result, we get the output

$$y = s_0(w_1 \cdot x_1 + \ldots + w_n \cdot x_n + w_0). \qquad (1)$$

The corresponding function $s_0(x)$ is known as the *activation function*; see, e.g., [2].

This is exactly how the standard artificial neural networks – that emulate biological neural networks – work:

– we feed the inputs x_i into one or more neurons, then
– we feed these neuron's outputs into other neurons, etc.

We can have simple networks, in which inputs go into the intermediate layer, and the outputs of the intermediate layer are collected by neurons from the final layer. We can have neural networks with more layers. Interestingly, it turns out that *deep learning* neural networks – i.e., networks with a large number of layers – are the most efficient ones; see, e.g., [3].

Which activation functions are most effective. In the past, most neural networks used the *sigmoid* activation functions $s_0(x) = \dfrac{1}{1 + \exp(-k \cdot x)}$, the activation function which provides the most adequate description of data processing in biological neurons.

However, recently, it was shows that we can make neural networks more efficient if instead, we use *rectified linear* neurons, with piecewise linear activation function $s_0(x) = \max(x, 0)$, i.e.:

- $s_0(x) = x$ when $x \geq 0$, and
- $s_0(x) = 0$ for $x < 0$.

Such neurons are especially efficient in *deep learning* [3].

In particular, we successfully used rectified linear neurons to predict volcanic eruptions based on preceding seismic activity; see, e.g., [9, 10].

Comment. It is easy to prove that 3-layer neural networks with rectified linear neurons are universal approximators for continuous functions on a bounded domain. Indeed:

– each function can be represented as a difference of two convex functions (see, e.g., [13]), and
– each convex function is a maximum of all tangent linear functions – and thus, can be well approximate if we take finitely many tangent linear functions [13].

Why are rectified linear neurons efficient: an open question. While empirical evidence shows that rectified linear neurons work best, there seems to be no convincing theoretical explanation for this empirical success. Without such an explanation, it is not clear whether these neurons are indeed the best

– or maybe some other activation function would lead to even more efficient computations?

What we do in this paper. In this paper, we show that several different uncertainty-related techniques can lead to a theoretical explanation of why rectified linear activation functions are empirically successful.

The fact that we get several different theoretical explanations makes us very confident that the empirical success of rectified linear neurons is not an accident – it is a well-justified phenomenon.

2   First Symmetry-Based Explanation

Numerical values depend on the measuring unit. When we measure a physical signal, the resulting numerical value depends on what measuring unit we use in this measurement. For example, when we measure the height in meters, the person's height is 1.7. However, if we measure the same height in centimeters, we get a different numerical value: 170.

In general, if instead of the original measuring unit, we use a different unit which us λ times smaller than the previous one, then all the numerical values get multiplied by λ; e.g., if we replace meters by centimeters, all numerical values get multiplied by $\lambda = 100$.

When is an activation function independent of the measuring unit? The choice of a measuring unit is rather arbitrary, it does not change the physical situation. It is reasonable to require that the results of applying the corresponding non-linear activation function not change is we simply change the measuring unit.

In precise terms, this means that if we have $y = s_0(x)$, then for any $\lambda > 0$, we should have $y' = s_0(x')$, where we denoted $x' = \lambda \cdot x$ and $y' = \lambda \cdot y$. Let us see what we can derive based on this requirement.

Definition 1. *We say that a function $s_0(x)$ is a scale-invariant if, for every x, y, and $\lambda > 0$, $y = s_0(x)$ implies that $\lambda \cdot y = s_0(\lambda \cdot x)$.*

Proposition 1. *A function $s_0(x)$ is scale-invariant if and only if it has the following form:*

- *$s_0(x) = c_+ \cdot x$ for $x \geq 0$ and*
- *$s_0(x) = c_- \cdot x$ for $x < 0$,*

for some constants c_+ and c_-.

Comment. One can easily check that each such function has the form

$$s_0(x) = c_- \cdot x + (c_+ - c_-) \cdot \max(x, 0).$$

Thus, if $c_+ \neq c_-$, i.e., if the corresponding activation function is not linear, then the class of functions represented by s_0-neural networks coincides with the class of functions represented by rectified linear neural networks.

Proof of Proposition 1.

$1°$. Every input x is either equal to 0, or positive, or negative. Let us consider these three cases one by one.

$2°$. Let us first consider the case of $x = 0$.

For $x = 0$ and $\lambda = 2$, scale invariance means that if $y = s_0(0)$, then $2y = s_0(0)$. Thus, $2y = y$, hence $y = s_0(0) = 0$.

$3°$. Let us now consider the case of positive values x.

Let us denote $c_+ \stackrel{\text{def}}{=} s_0(1)$. Then, by using scale-invariant with:

- x instead of λ,
- 1 instead of x, and
- c_+ instead of y,

we conclude that for all $x > 0$, $c_+ = s_0(1)$ implies that $c_+ \cdot x = s_0(x)$.

For positive values x, the desired equality is proven.

$4°$. To complete the proof of this result, we need to prove it for negative inputs x.

Let us denote $c_- \stackrel{\text{def}}{=} -s_0(-1)$. In this case, $s_0(-1) = -c$. Thus, for every $x < 0$, by using scale-invariance with:

- $\lambda = |x|$,
- $x = -1$, and
- $y = s_0(-1) = -c_-$,

we conclude that

$$s_0(x) = s_0(|x| \cdot (-1)) = |x| \cdot s_0(-1) = |x| \cdot (-c_-) = c_- \cdot x.$$

The proposition is proven.

3   Second Symmetry-Based Explanation

Idea. We are interested in *optimal* activation functions, i.e., functions which are the best according to some optimality criterion. When we say "an optimality criterion",

- we do not necessarily need to have numerical values attached to each activation function, but
- we do not need to able to compare different activation functions.

Thus, we need to have an order (at least a partial order) $\preceq$ on the set of all possible activation functions.

A function $s_0(x)$ is optimal if is better (or of the same quality) than all other possible activation functions, i.e., if $s \preceq s_0$ for all possible activation functions $s(x)$.

We want an optimality criterion to be useful, i.e., we want to use it to select an activation function. Thus, there should be at least one activation function which is optimal according to this criterion.

What if several different functions are optimal according to the given criterion? In this case, we can use this non-uniqueness to optimize something else. For example, if on a given class of benchmarks, neurons that use several different activation functions have the same average approximation error, we can select, among the, the function with the smallest computational complexity. This way, instead of the original optimality criterion, we, in effect, use a new criterion according to which s_0 is better than s_0 if:

- either it has the smaller average approximation error
- or it has the same average approximation error and smaller computational complexity.

If, based on this modified criterion, we still have several different activation functions which are equally good, we can use this non-uniqueness to optimize something else: e.g., worse-case approximation accuracy, etc.

Thus, every time the optimality criterion selects several equally good activation functions, we, in effect, replace it with a modified criterion, and keep modifying it until finally we get a criterion for which only one activation function is optimal. So, we arrive at the following definition.

Definition 2.

- *By an* optimality criterion*, we mean a partial order $\preceq$ on the set of all continuous functions of one variable.*
- *We say that a function s_0 is* optimal *with respect to the optimality criterion $\preceq$ if $s \preceq s_0$ for all functions s.*
- *We say that an optimality criterion is* final *if there exists exactly one function which is optimal with respect to this optimality criterion.*

Scale-invariance. It is reasonable to require that the quality of an activation function does not depend on the choice of the measuring unit. Let us describe this requirement in precise terms.

Suppose that in some selected units, the activation function has the form $s(x)$. If we replace the original measuring unit by a new unit which is λ times larger that the original one, then the value x in the new units is equivalent to $\lambda \cdot x$ in the old units. If we apply the old-unit activation function to this amount, we get the output of $s(\lambda \cdot x)$ of old units – which is equivalent to $\lambda^{-1} \cdot s(\lambda \cdot x)$ new units.

Thus, after the change in units, the transformation described, in the original units, by an activation function $s(x)$ is described, in the new units, by a modified activation function $\lambda^{-1} \cdot s(\lambda \cdot x)$. So, the above requirement takes the following form:

Definition 3. *We say that an optimality criterion $\preceq$ is* scale-invariant *if for every two functions s and s' and for every $\lambda > 0$, the relation $s \preceq s'$ is equivalent to $T_\lambda(s) \preceq T_\lambda(s')$, where we denoted $(T_\lambda(s))(x) \stackrel{\text{def}}{=} \lambda^{-1} \cdot s(x)$.*

Now, we are ready to formulate our result.

Proposition 2. *A function $s_0(x)$ is optimal with respect to some final scale-invariant optimality criterion if and only if it has the following form:*

- $s_0(x) = c_+ \cdot x$ *for $x \geq 0$ and*
- $s_0(x) = c_- \cdot x$ *for $x < 0$.*

Comment. Thus, we have yet another justification for the success of rectified linear activation functions.

Proof of Proposition 2.

$1°$. For every function $s_0(x)$ of the above type, we can easily find a final scale-invariant optimality criterion for which this function is optimal: namely, we can take the order $\preceq$ in which $s \preceq s_0$ for all continuous functions $s(x)$.

One can easily check:

- that this relation is final and scale-invariant, and
- that the given function $s_0(x)$ is the only function which is optimal with respect to this criterion.

$2°$. Vice versa, let us assume that a function $s_0(x)$ is optimal with respect to some final scale-invariant optimality criterion. Under this assumption, we need to prove that the function $s_0(x)$ has the desired form. To prove this, let us prove that this function is scale-invariant in the sense of Definition 1; if we prove this scale-invariance, then the desired result will follow from Proposition 1.

In terms of the transformation T_λ, scale-invariance means that $s_0 = T_\lambda(s_0)$ for all s. To prove that $T_\lambda(s_0) = s_0$, let us prove that the function $T_\lambda(s_0)$ is optimal. Then, the desired equality will follow from the fact that the optimality criterion is final – and thus, there is only one optimal function.

To prove that the function $T_\lambda(s_0)$ is optimal, we need to prove that $s \preceq T_\lambda(s_0)$ for all s. Due to scale-invariance of the optimality criterion, this condition is equivalent to $T_{\lambda^{-1}}(s) \preceq s_0$ – which is, of course, always true, since s_0 is optimal. Thus, $T_\lambda(s_0)$ is also optimal, hence $T_\lambda(s_0) = s_0$, and by Proposition 1, the function $s_0(x)$ has the desired form.

The proposition is proven.

4   Complexity-Based Explanation

Idea. To speed up computations, we need to make sure that the activation function is as fast to compute as possible.

This idea. Inside the computer, every numerical operation is implemented as a composition of the basic hardware-supported operations. These operations include the basic arithmetic operations:

- addition $a + b$,

- subtraction $a - b$,
- multiplication $a \cdot b$,
- division a/b,

and the operations $\min(a, b)$ and $\max(a, b)$.

Of these operations:

- the functions min and max are the fastest,
- addition + and subtraction − are next fastest,
- followed by multiplication (which involves several additions) and
- division (which involves several multiplications);

see, e.g., [11].

The fastest-to-compute activation function is the one that uses only one hardware supported basic operation.

We are interested in non-linear activation functions (since linear transformation are already taken care in the aggregation procedure, before we invoke the activation function). Out of the above operations, the corresponding functions $s_0(a) = a + a_0$, $s_0(a) = a - a_0$, $s_0(a) = a_0 - a$, $s_0(a) = a \cdot a_0$, and $s_0(a) = a/a_0$ are linear. The only non-linear operations are $\max(a, a_0)$, $\min(a, a_0)$, and a_0/a. Of these three operations, the fastest are piecewise linear operations min and max.

Thus, the computational complexity-based analysis indeed leads to yet another justification for the use of piecewise linear activation functions.

Comment. A similar justification can be made if we are thinking about a hardware implementation of artificial neural networks. Indeed, in this case, a linear combination is straightforward: just place several currents together.

The simplest nonlinear element of an electric circuit is a *diode* that transmits current only in one direction. For the diode, the output is equal to x if $x \geq 0$ and to 0 otherwise, i.e., it is exactly the rectified linear activation function – which is thus the easiest to implement in hardware.

5   Fuzzy-Based Justification

Need to use fuzzy techniques. When we use neural network technique to learn a phenomenon, we generate a neural network that provides a good approximation to this phenomenon. In particular, when we use the neural network technique to provide a solution to a problem – e.g., to provide an appropriate control – we thus produce a neural network that generates the corresponding solution.

In human reasoning, we try our best not only to provide good solutions to real-life problems, but also to provide a clear justification for these solutions.

It is therefore reasonable to look for activation functions for which the corresponding solution makes direct sense, i.e., for which this solution can be interpretable in human-understandable natural-language terms.

The need for translating imprecise ("fuzzy") expert knowledge into precise (and thus, computer-understandable) form has been well recognized since the early 1960s. Techniques that provide such a translation are known as *fuzzy techniques*; see, e.g., [1, 4, 6, 7, 14].

In terms of these techniques, the above idea can be reformulated as follows: we want to select an activation function for which all the functions representing the corresponding neural networks are directly interpretable in fuzzy terms.

Which functions can be interpretable in fuzzy terms. It is known that if we use $1 - a$ as negation, $\min(a + b, 1)$ as an "or"-operation and $\max(a + b - 1, 0)$ as an "and"-operation, then functions that can be represented as compositions of logical operations are exactly piece-wise linear functions with integer coefficients [5, 8, 12].

To these operations, we can add more subtle operations. For example, it is natural to interpret "somewhat A" as $A \vee A$ – which, in the above logic, leads to $2a$ (or, to be more precise, to $\min(2a, 1)$). It is therefore reasonable to define an inverse hedge "very A" as the statement B for which "somewhat B" is equivalent to A. In the above logic, this would mean defining our degree of confidence in "very A" as $a/2$, where A is our degree of confidence in the original statement A.

We can iterate this "very" hedge, thus getting values $a/4$, $a/8$, etc. By combining these hedges and logical operations, we can get any piecewise linear functions with binary-rational coefficients.

This leads to a new justification for piecewise linear activation functions. We want a neural network to be interpretable. For the neural network to be interpretable, we need to make sure that all the data processing algorithms performed by a neural network can be described in fuzzy terms. Since implies that all such algorithms must be piecewise-linear.

This conclusion means, in particular, that the activation function should be piecewise linear. Thus, we indeed get one more justification for using piecewise linear activation functions in neural networks.

Acknowledgments

References

1. R. Belohlavek, J. W. Dauben, and G. J. Klir, *Fuzzy Logic and Mathematics: A Historical Perspective*, Oxford University Press, New York, 2017.
2. C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, New York, 2006.
3. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Leaning*, MIT Press, Cambridge, Massachusetts, 2016.
4. G. Klir and B. Yuan, *Fuzzy Sets and Fuzzy Logic*, Prentice Hall, Upper Saddle River, New Jersey, 1995.

5. R. McNaughton, "A theorem about infinite-valued sentential logic", *Journal of Symbolic Logic*, 1951, Vol. 16, pp. 1–13.
6. J. M. Mendel, *Uncertain Rule-Based Fuzzy Systems: Introduction and New Directions*, Springer, Cham, Switzerland, 2017.
7. H. T. Nguyen and E. A. Walker, *A First Course in Fuzzy Logic*, Chapman and Hall/CRC, Boca Raton, Florida, 2006.
8. V. Novák, I. Perfilieva, and J.. Močkoř, *Mathematical Principles of Fuzzy Logic*, Kluwer, Boston, Dordrecht, 1999.
9. J. Parra, O. Fuentes, E. Anthony, and V. Kreinovich, "Prediction of volcanic eruptions: case study of rare events in chaotic systems with delay", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics SMC'2017*, Banff, Canada, October 5–8, 2017, to appear.
10. J. Parra, O. Fuentes, E. Anthony, and V. Kreinovich, "Use of machine learning to analyze and – hopefully – predict volcano activity", *Acta Politechnica Hungarica*, to appear.
11. D. Patterson and J. L. Hennessy, *Computer Architecture: A Quantitative Approach*, Morgan Kaufman, Waltham, Massachusetts, 2012.
12. I. Perfilieva and A. Tonis, "Functional system in fuzzy logic formal theory", *BUSEFAL*, 1995, Vol. 64, pp. 42–50.
13. R. T. Rockafeller, *Convex Analysis*, Princeton University Press, Princeton, New Jersey, 1970.
14. L. A. Zadeh, "Fuzzy sets", *Information and Control*, 1965, Vol. 8, pp. 338–353.