

Semiring Activation in Neural Networks

Bart M.N. Smets

Peter D. Donker

Jim W. Portegies

*Department of Mathematics and Computer Science
Eindhoven University of Technology, the Netherlands*

B.M.N.SMETS@TUE.NL

P.D.DONKER@STUDENT.TUE.NL

J.W.PORTEGIES@TUE.NL

Abstract

We introduce a class of trainable nonlinear operators based on semirings that are suitable for use in neural networks. These operators generalize the traditional alternation of linear operators with activation functions in neural networks. Semirings are algebraic structures that describe a generalised notation of linearity, greatly expanding the range of trainable operators that can be included in neural networks. In fact, max- or min-pooling operations are convolutions in the tropical semiring with a fixed kernel.

We perform experiments where we replace the activation functions for trainable semiring-based operators to show that these are viable operations to include in fully connected as well as convolutional neural networks (ConvNeXt). We discuss some of the challenges of replacing traditional activation functions with trainable semiring activations and the trade-offs of doing so.

Keywords: neural network, activation function, semiring, trainable nonlinearity

1. Introduction

Neural networks come in a large variety of types and architectures, one common characteristic the majority of them share is the alternation between trainable linear operations (or affine operations if one includes bias) on the one hand and non-trainable nonlinear operators in the form of a scalar activation function on the other hand. Sometimes a multi-variate nonlinear function like max/min-pooling is also used, but this is again a fixed, non-trainable function similar in nature to a scalar activation function. Even in the case of transformers (Vaswani et al., 2017) one trains linear maps that are then composed in a fixed manner through the inner product and a soft-max function. While this binary setup is the de facto standard in machine learning, exceptions exist. These exceptions can be roughly classified into three classes: trainable activation functions, non-standard neurons and morphological. We will briefly discuss these three approaches. For a comprehensive overview of trainable activation functions and non-standard neurons see Apicella et al. (2021).

Trainable activation functions. The most straightforward way of obtaining a trainable nonlinearity is using an activation function that has one or more parameters affecting it and train these parameters as part of the neural network. This is typically done either by adding some shape parameter to an existing fixed activation function or building the activation function as an ensemble of some fixed basis functions and have the ensemble coefficients be trainable parameters. An example of the first kind is the *wish* function

(Ramachandran et al., 2017) defined as

$$\text{swish}_\alpha(x) := x \cdot \text{sigmoid}(\alpha \cdot x),$$

where $\alpha \in \mathbb{R}$ is a parameter that can either be trained or constant. Note that for $\alpha = 1$ the swish function reduces to the SiLU (Elfwing et al., 2018) and for $\alpha \rightarrow \infty$ reduces to the familiar ReLU.

An example of the second kind is the *adaptive blending unit* (Sütfield et al., 2018) given by

$$\text{ABU}(x) := \sum_{i=1}^k \alpha_i \cdot f_i(x),$$

where the α_i 's are the trainable parameters and the functions f_i are a selection of fixed activation functions such as tanh, ReLU, id, swish, etc. In the original study the α_i parameters were initialized with $\frac{1}{k}$ and constrained using a normalization scheme.

As Apicella et al. (2021) notes, in most cases these trainable activation functions can also be expressed (or approximated) as a small feed-forward neural network with fixed activation functions. This is not surprising since the fixed activation functions remain the core building blocks. As a consequence, the same benefit of using these trainable activation functions can be achieved by simply making the neural network deeper.

Non-standard neurons. The alternative to making the activation function trainable is doing away with the standard neuronal model (linear map followed by an activation function) altogether. A basic example is a *maxout network* (Goodfellow et al., 2013) where each neuron has a number of linear maps with scalar codomain and the final output is the maximum of the group.

A more elaborate non-standard neuron is used in an *morphological neural network* as introduced by Ritter and Sussner (1996). Here the idea is to replace the inner linear combination in the classic neuron given by

$$y_j = \sigma \left(\sum_{i=1}^n w_{ij} x_i + b_j \right),$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is an activation function, with a *tropical* combination given by:

$$y_j = \sigma \left(\max \{ w_{ij} + x_i \mid i = 1 \dots n \} + b_j \right).$$

and where σ is again a choice of scalar activation function. More examples of non-standard neurons can be found in Apicella et al. (2021).

Morphological. Mathematical morphology is a theory and set of techniques for analyzing and processing geometric features, most commonly in images (Serra, 1982). The general technique is similar to convolution where one move a kernel along an input image and multiply and integrate to generate an output. Instead one calls the kernel the *structuring element* and instead of multiplying one add or subtracts and then takes the maximum or minimum instead of integrating. This operation is also referred to as *morphological convolution*, and can be thought of as a nonlinear version of the familiar linear convolution. The use of morphological techniques in deep learning can be considered a special case of non-standard neurons, but it is specific enough to warrant separate discussion.

The earliest work we are aware of is the *PConv* operator in Masci et al. (2013), here the morphological convolution is not implemented directly but approximated through counter-harmonic means, which is also the approach taken in Mellouli et al. (2017). Later work uses other approximations like soft maximum and minimum in Shih et al. (2019). Both approaches have some trouble when executed in floating point arithmetic, that can however be ameliorated to some degree by a smart choice of bias as in Shen et al. (2019). Direct computation of morphological convolution is also possible as is used in *PDE-based CNNs* by Smets et al. (2023).

Our work looks at linear convolution and morphological convolution as two special cases of a more general family of operators where a particular choice of codomain algebra and domain translation equivariance naturally yields a convolution type operator.

Our approach. Instead of looking at deep neural networks as consisting of layers of neurons, we can also take the view that we are alternating between linear and nonlinear operators. Sure, we usually choose the nonlinear operator to be a scalar activation function and let the linear operators be the trainable part but that is but one possible design choice.

Suppose we could place the nonlinear operator on an equal footing to the linear one, what would this look like? Clearly, the nonlinear operator needs some sort of structure similar to the linear operators to make this work. If $A : V \rightarrow W$ is some map between vector spaces V and W then we say it is linear if

$$A(av_1 + bv_2) = aA(v_1) + bA(v_2)$$

for all $v_1, v_2 \in V$ and scalars a and b . We could require a similar structure of a nonlinear operator $B : X \rightarrow Y$, so that for some binary operations \oplus and \odot we have

$$B(a \odot x \oplus b \odot y) = a \odot B(x) \oplus b \odot B(y) \tag{1}$$

for all $x_1, x_2 \in X$ and scalars a and b . The spaces X and Y would have to be spaces where these operations make sense of course. If this allows for the nonlinear operator B to be written as a matrix, similar to A then we could train B in the same way we train A and effectively have linear and nonlinear operators on the same footing.

This then will be our approach, look at a class of nonlinear operators that are *quasilinear* as in (1) and then instead of building neural networks with trainable linear (or affine) operators A_i and activation functions σ as

$$A_L \circ \sigma \circ A_{L-1} \circ \dots \circ A_2 \circ \sigma \circ A_1,$$

we build it with linear operators A_i and nonlinear operators B_i as

$$A_L \circ B_{L-1} \circ A_{L-1} \circ \dots \circ A_2 \circ B_1 \circ A_1,$$

where both types are trainable. We will refer to this idea as *semiring activation*.

An additional property of this approach is that it is entirely compatible with the notion of *geometric equivariance*: the property of a model where when a certain geometric transformation is applied to the input (say a translation is applied to an image input) then the output of the model should undergo a corresponding transformation (if the output is also

an image, that image should also be translated, if the output is a classification label then the label should not change, i.e. be invariant).

The original impetus of the approach we take in this work came from our previous research into *PDE-based equivariant CNNs* in Smets et al. (2023). In that work we also built neural networks without activation functions and used trainable nonlinear operators based on certain PDE solvers. The PDE solvers we considered – while nonlinear – did have a *quasilinear* property of the same kind as we will discuss in this work. The question we asked ourselves was how much of the performance of the PDE-G-CNNs in Smets et al. (2023) is due to the PDE structure? Or, is much of the performance explained by just the semiring structure of the nonlinearities? In this work we examine how far we can get using just semiring based nonlinearities.

2. Quasilinear operators from semirings

To construct quasilinear operator as in (1) we need our generalized addition \oplus and generalized multiplication \odot to share some behaviour with the conventional “+” and “.”. Specifically, we will require them to form a semiring.

Definition 1 (Semiring) *A semiring is a set R equipped with two binary operations \oplus and \odot , called addition and multiplication, so that*

- (i) *addition and multiplication are associative,*
- (ii) *addition is commutative,*
- (iii) *addition has an identity element 0_R ,*
- (iv) *multiplication has an identity element 1_R ,*
- (v) *multiplication distributes over addition:*

$$a \odot (b \oplus c) = a \odot b \oplus a \odot c \quad \text{and} \quad (a \oplus b) \odot c = a \odot c \oplus b \odot c,$$

- (vi) *multiplication with 0_R annihilates: $0_R \odot a = a \odot 0_R = 0_R$.*

Additionally, if $a \oplus a = a$ we say the semiring is *idempotent* and if $a \odot b = b \odot a$ we say the semiring is *commutative*. As is conventional we let multiplication take precedence over addition: $a \oplus b \odot c = a \oplus (b \odot c)$.

In the same way we construct linear maps through matrix multiplication we can now construct a quasilinear map. Let R be some commutative semiring, $x \in R^n$ and $B \in R^{m \times n}$ and define

$$B \odot x = \begin{bmatrix} b_{11} & \cdots & b_{1n} \\ \vdots & \ddots & \vdots \\ b_{m1} & \cdots & b_{mn} \end{bmatrix} \odot \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} := \begin{bmatrix} b_{11} \odot x_1 \oplus \cdots \oplus b_{1n} \odot x_n \\ \vdots \\ b_{m1} \odot x_1 \oplus \cdots \oplus b_{mn} \odot x_n \end{bmatrix}. \quad (2)$$

Which is nothing more than the usual matrix-vector multiplication operation when we take the linear semiring $R = (\mathbb{R}, +, \cdot)$. When we further overload the symbols \oplus and \odot to component-wise addition resp. scalar multiplication in R^n , i.e.:

$$a \odot \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} := \begin{bmatrix} a \odot x_1 \\ \vdots \\ a \odot x_n \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \oplus \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} := \begin{bmatrix} x_1 \oplus y_1 \\ \vdots \\ x_n \oplus y_n \end{bmatrix},$$

then it can be verified that B satisfies

$$B \odot (a \odot x \oplus b \odot y) = a \odot (B \odot x) \oplus b \odot (B \odot y). \quad (3)$$

Which makes $x \mapsto B \odot x$ satisfy (1), we say it is a *quasilinear operator with respect to R* . Adding a bias can also be done in this setting: $x \mapsto B \odot x \oplus c$ with $c \in R^m$.

Remark 2 (Semimodules and their homomorphisms) *In algebraic terms if we have a semiring R then R^n equipped with the component-wise addition $\oplus : R^n \times R^n \rightarrow R^n$ and scalar multiplication $R \times R^n \rightarrow R^n$ from above forms a left R -semimodule. A semimodule is a generalization of the concept of a vector space, where the underlying set of scalars is a semiring but not necessarily a field like \mathbb{R} . All fields are semirings and all vector space are semimodules but not the other way around. Quasilinearity as in (1) and (3) can then be understood as a homomorphism between semimodules.*

3. Logarithmic and tropical semirings

Semirings, being fairly general, come in a large variety. For our current purpose we are only interested in semirings that have the real numbers, possibly extended with $\pm\infty$, as their set. This way the new quasilinear operators can coexist with the linear operators since each can deal with the output of the other. The new addition and multiplication operations should also be easy enough to compute in practice so as to not incur an unreasonable performance penalty.

Logarithmic semirings. The first semirings we consider is the family of *logarithmic semirings*. Let $\mu \neq 0$ and define

$$a \oplus_{\mu} b := \frac{1}{\mu} \log \left(e^{\mu a} + e^{\mu b} \right) \quad \text{and} \quad a \odot b := a + b.$$

Here we adopt the convention that $e^{\infty} = \infty$, $e^{-\infty} = 0$ and correspondingly $\log(\infty) = \infty$ and $\log(0) = -\infty$. Then for $\mu > 0$ we have that $R_{\log}^{\mu} := (\mathbb{R} \cup \{-\infty\}, \oplus_{\mu}, +)$ forms a commutative semiring. Associativity, commutativity and distributivity are easy to check. The identity element for addition is $0_R = -\infty$ and for multiplication $1_R = 0$. The annihilation axiom is also satisfied since $0_R \odot a = -\infty + a = -\infty = 0_R$.

Similarly for $\mu < 0$ we have that $R_{\log}^{\mu} := (\mathbb{R} \cup \{\infty\}, \oplus_{\mu}, +)$ is a commutative semiring except that now the identity element for the addition is $0_R = \infty$.

Tropical semirings. The second family we consider is that of the *tropical semirings* $R_{\max} := (\mathbb{R} \cup \{-\infty\}, \max, +)$ and $R_{\min} := (\mathbb{R} \cup \{\infty\}, \min, +)$, also called the *max-plus semiring* and *min-plus semiring* respectively. For the max-plus semiring, $-\infty$ is the additive identity since $\max(a, -\infty) = a$ while for the min-plus semiring ∞ is the additive identity, 0 is the multiplicative identity for both. These two semirings are isomorphic under negation $x \mapsto -x$. Both the tropical semirings are commutative and idempotent since $\max(a, a) = \min(a, a) = a$.

Remark 3 *Morphological neural networks as in Ritter and Sussner (1996) replace the linear combination of the neuron with what we could now call a tropical combination, indeed*

$\max_{i=1\dots n}\{w_i + x_i\}$ is a quasilinear function in the max-plus semiring and can be written as

$$y_j = \bigoplus_{i=1}^n w_{ij} \odot x_i,$$

with respect to the tropical \oplus and \odot . Our approach is distinct in that 1) we do not aim to replace the linear combination but rather the activation function while keeping the linear combination, 2) have a variety of semirings to consider.

4. In fully connected networks

To check the viability of our proposed scheme we first run a set of experiments with small fully connected networks on a variety of datasets. The datasets for this series of experiments are:

- the classic iris dataset from Anderson (1936), available at <https://www.kaggle.com/datasets/uciml/iris>,
- the heart disease dataset available at <https://www.kaggle.com/datasets/johnsmith88/heart-disease-dataset>,¹
- the circles and spheres dataset from Naitzat et al. (2020), available at https://github.com/topnn/topnn_framework,
- the FashionMNIST dataset from Xiao et al. (2017), available at <https://github.com/zalandoresearch/fashion-mnist>.

The code of the experiments is available at <https://github.com/bmnsnets/semitorch>.

4.1 Architectures

We will train a set of network architectures that only slightly vary based on the dataset. Every model will be based on a common architecture with a linear stem and head with two layers with residual connections in between, see Figure 1a. The stem and head are there to convert the number of features in the dataset n to a common internal network width w and back to the number of classes c in the dataset. The internal layers we will vary between a traditional ReLU based layer (Figure 1b) and semiring based layers (Figure 1c and 1d). For the iris and heart disease datasets we use the ReLU layer without layer normalization and semiring layer from Figure 1c, for the circles, spheres and FashionMNIST dataset we use the ReLU layer with layer normalization and the semiring layer from Figure 1d.

4.2 Training

Training the networks with semiring based activation proved challenging, the usual setups that are known to work well for conventional neural network do not necessarily carry over.

1. The provenance of this dataset is complicated and its contents are of questionable value for predicting heart disease, see Simmons II (2021) for an investigation into this dataset. For our purpose of comparing network architectures this dataset is still perfectly serviceable to see how well models deal with poor data.

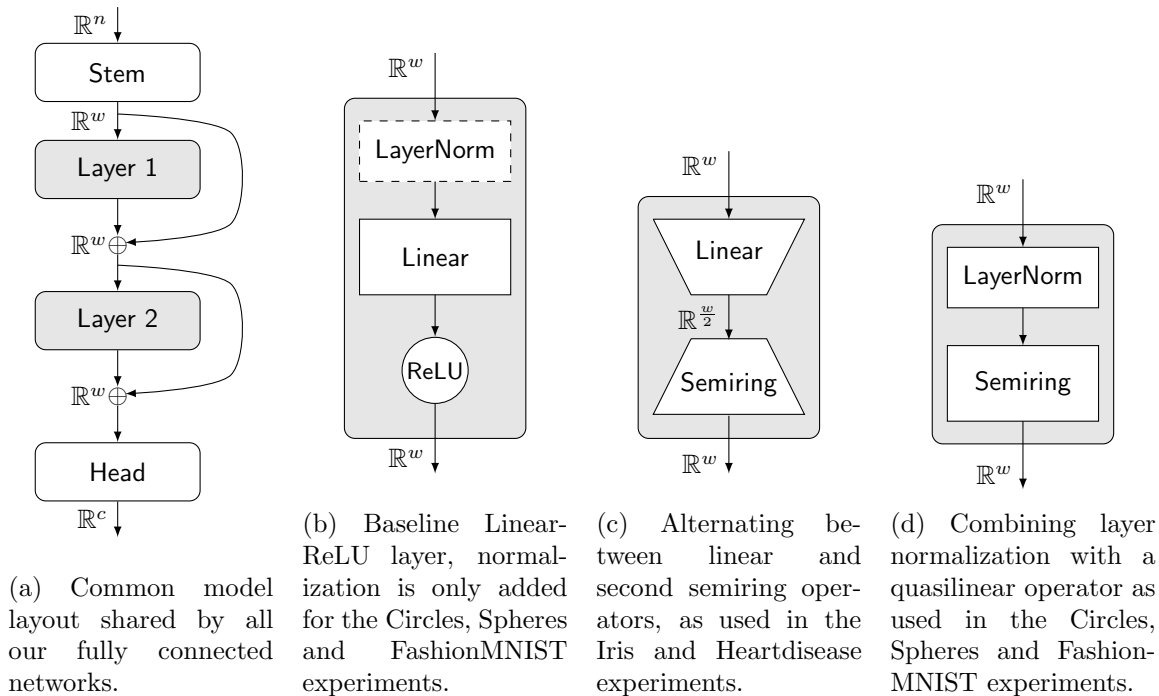


Figure 1: Network architecture for our fully connected experiments. The **Head** and **Stem** modules are linear modules. None of the modules include biases. The layer normalization modules include affine transforms. The number of input features n and number of output classes c are dataset dependent, the internal width parameter w is chosen per experiment. Each network under consideration has the exact same number of parameters per experiment.

Indeed, parameter initialization, normalization layers, optimizers, and schedulers are all areas that have seen much research to obtain the best possible results with conventional neural networks. We spent a lot of time finding a training setup that worked for our modified networks but we expect that this is an area where more gains could be found. We document the training setup that we used for our fully-connected experiments next.

Optimizer & learning rate scheduler. We ended up settling on a combination of the *AdamW* optimizer (Loshchilov and Hutter, 2019) and the *1-cycle* learning rate scheduler (Smith and Topin, 2019). Key in getting comparable performance out of the semiring based networks as compared to the baseline network was assigning a separate optimizer to the linear parameters and the parameters of the semiring module. In general we needed to assign smaller learning rates to the semiring parameters to have stable training. The FashionMNIST experiment is an exception to this, there we obtained best results by increasing the semiring learning rate and decreasing the linear learning rate. The hyperparameters for all the experiments are listed in Appendix 6.

4.3 Initialization

An important aspect of training neural networks is proper initialization of its parameters. The standard parameter initialization schemes like *Xavier* initialization (Glorot and Bengio, 2010) and *Kaiming* initialization (He et al., 2015) are derived specifically around the forward and backward stability of linear maps and so do not apply to semiring-based maps. Consequently, we need to come up with new initialization schemes for the semiring weights, specifically for tropical and logarithmic cases.

Tropical. For tropical operators (i.e. max-plus and min-plus) we propose an initialization scheme based on ‘fair’ backpropagation of gradients. Consider the max-plus operator $R_{\max}^m \rightarrow R_{\max}^n$ given by

$$y_i = \max_{j=1..m} w_{ij} + x_j \quad (4)$$

where $x_1, \dots, x_m \in R_{\max}$ are the inputs, $y_1, \dots, y_n \in R_{\max}$ are the outputs and $[w_{ij}]_{ij} \in R_{\max}^{n \times m}$ are the trainable parameters. Then the partial derivatives are given by

$$\frac{\partial y_i}{\partial x_j} = \begin{cases} 1 & \text{if } j = \arg \max_{k=1..m} w_{ik} + x_k, \\ 0 & \text{else,} \end{cases}$$

for $i = 1 \dots n$ and $j = 1 \dots m$. This could be problematic for the backward pass since if an input x_j never ‘wins’ one of the maxima, i.e. $\frac{\partial y_i}{\partial x_j} = 0$ for all $i = 1, \dots, m$, then its gradient will always be zero. Conversely, if an input x_j happens to be very large on a consistent basis then it will accumulate all the gradients of all the outputs y_1, \dots, y_n to itself. The result would be a very unbalanced gradient distribution during the backward pass, something we know from previous research on parameter initialization (Glorot and Bengio, 2010; He et al., 2015) is undesirable.

The extreme case for the operator (4) would be having an x_1 value (for example) that is consistently much larger than any other $x_j + w_{ij}$ so that for all $i = 1, \dots, n$ we always have $y_i = w_{i1} + x_1$.

Of course, we have limited control over the degree that this effect will manifest during training, but we can at least avoid it at the start by choosing an appropriate initialization scheme. The idea is to make sure that at initialization there is a high probability that there is at least one $i = 1, \dots, n$ so that $\frac{\partial y_i}{\partial x_j} = 1$ for each $j = 1, \dots, m$. This is of course only possible when $n \geq m$, but in all our experiments we have used models for which this is the case (see Figure 1 and 2). Our aim will be to initialize the weights w_{ij} in (4) so that there is a high probability of each of the n inputs ‘winning’ roughly $\frac{n}{m}$ of the m outputs.

Assuming the inputs x_j generally stay in a range $[-\frac{K}{2}, \frac{K}{2}]$ for some $K > 0$, then in the max-plus case we can initialize the weight matrix $W = [w_{ij}]_{ij} \in \mathbb{R}^{m \times n}$ as

$$w_{ij} = \text{Unif}[-\varepsilon, \varepsilon] + \begin{cases} 0 & \text{if } i = j \text{ mod } m \\ -K & \text{else.} \end{cases} \quad (5)$$

The second term applies a penalty of $-K$ to each input unless $i = j \text{ mod } m$ ensuring that the m available ‘wins’ are fairly distributed amount the n inputs (at least with high probability based on our assumptions on the inputs). We additionally add a modest uniform

distribution $\text{Unif}[-\varepsilon, \varepsilon]$ to keep the initialization scheme from being deterministic from run to run, ε is chosen on the order of $\frac{K}{2}$.

A matrix initialized like this, but without the stochastic term, looks like

$$\begin{bmatrix} 0 & -K & \cdots & -K \\ -K & 0 & & -K \\ \vdots & & \ddots & \\ -K & & & 0 \\ 0 & -K & \cdots & -K \\ -K & 0 & & -K \\ \vdots & & \ddots & \end{bmatrix}. \quad (6)$$

We see that in each row we get a single zero coefficient with the others having the value $-K$. Assuming that the inputs are in the range $[-\frac{K}{2}, \frac{K}{2}]$, and we forget the stochastic term for the moment, then the input corresponding to the column with the zero value will achieve the maximum for the output corresponding with that row. This has the effect of $(m \bmod n)$ number of inputs contributing to $\lceil m/n \rceil$ outputs and the remaining inputs contributing to $\lfloor m/n \rfloor$ outputs. This avoids the vanishing and exploding gradient problem that would be caused by a single input dominating.

For the min-plus case we similarly set

$$w_{ij} = \text{Unif}[-\varepsilon, \varepsilon] + \begin{cases} 0 & \text{if } i = j \bmod m \\ K & \text{else.} \end{cases} \quad (7)$$

In our experiments normalization layers and weight decay keep input values fairly small, so our starting assumption of inputs being in a range $[-\frac{K}{2}, \frac{K}{2}]$ generally holds. We found $K = 1$ gave the intended effect of an initial fair distribution during the backward pass without overly biasing the initialization. We call this scheme *fair tropical initialization*.

Logarithmic. For the logarithmic semiring (for some choice of $\mu \in \mathbb{R} \setminus \{0\}$) consider the operator $(R_{\log}^{\mu})^m \rightarrow (R_{\log}^{\mu})^n$ given by

$$y_i = \frac{1}{\mu} \log \left(\sum_{j=1}^m e^{\mu(w_{ij} + x_j)} \right) \quad (8)$$

for all $i = 1 \dots n$, where $x_1, \dots, x_m \in R_{\log}^{\mu}$ are the inputs, $y_1, \dots, y_n \in R_{\log}^{\mu}$ are the outputs and $[w_{ij}]_{ij} \in (R_{\log}^{\mu})^{n \times m}$ are the trainable parameters. Then the partial derivatives are given by

$$\frac{\partial y_i}{\partial x_j} = \frac{e^{\mu(x_j + w_{ij})}}{\sum_{k=1}^m e^{\mu(x_k + w_{ik})}},$$

which is nothing but the *softmax* function with temperature $\frac{1}{\mu}$ over the values $\{x_k + w_{ik}\}_{k=1}^m$.

Let us now take the same approach as Glorot and Bengio (2010) and look at the inputs x_j and outputs y_i as random variables and control forward variance. We assume all the

inputs x_j are i.i.d. with expected value $\mathbb{E}[x_j] = 0$ and some finite variance $\text{Var}(x) < \infty$. Then we can estimate the variance of the outputs using the *delta method* as

$$\text{Var}(y_i) \approx \sum_{j=1}^m \frac{\partial y_i}{\partial x_j} (\mathbb{E}[x_1], \dots, \mathbb{E}[x_m])^2 \text{Var}(x_j),$$

and since we assumed the input distributions to be i.i.d. and centered we have

$$\text{Var}(y_i) \approx \text{Var}(x) \sum_{j=1}^m \frac{\partial y_i}{\partial x_j} (0, \dots, 0)^2 = \text{Var}(x) \sum_{j=1}^m \left(\frac{e^{\mu w_{ij}}}{\sum_{k=1}^m e^{\mu w_{ik}}} \right)^2.$$

Ideally we want $\text{Var}(y_i) \approx \text{Var}(x)$, which is the case if

$$\sum_{j=1}^m \left(\frac{e^{\mu w_{ij}}}{\sum_{k=1}^m e^{\mu w_{ik}}} \right)^2 \approx 1$$

or

$$\sum_{j=1}^m (e^{\mu w_{ij}})^2 \approx \left(\sum_{j=1}^m e^{\mu w_{ij}} \right)^2 \quad (9)$$

for all $i = 1, \dots, n$. We can satisfy this condition exactly by choosing a single $J_i = 1, \dots, m$ for each $i = 1, \dots, n$ and set that weight to $w_{i,J_i} = 0$ and set the other weights to $w_{ij} = -\text{sgn}(\mu)\infty$. In that case $e^{\mu w_{ij}} = 1$ if $j = J_i$ and $e^{\mu w_{ij}} = 0$ if not and consequently both sides of (9) are equal to 1.

We can also make a similar analysis for the backward pass. Let $(\bar{y}_i)_{i=1}^n$ be the loss gradients of the outputs and $(\bar{x}_j)_{j=1}^m$ be the loss gradients of the inputs. We interpret these as random variables, where we assume the output gradients \bar{y}_i are i.i.d. with $\mathbb{E}[\bar{y}_i] = 0$ and finite variance $\text{Var}(\bar{y}_i) = \text{Var}(\bar{y}) < \infty$ for all $i = 1, \dots, n$. Then the loss gradients of the inputs are computed as

$$\bar{x}_j = \sum_{i=1}^n \frac{\partial y_i}{\partial x_j} (x_1, \dots, x_m) \bar{y}_i,$$

or if we are talking about the expected backward pass over the possible inputs we can say

$$\bar{x}_j \approx \sum_{i=1}^n \frac{\partial y_i}{\partial x_j} (0, \dots, 0) \bar{y}_i.$$

The variance of the input gradients can then be approximated as

$$\text{Var}(\bar{x}_j) \approx \text{Var} \left(\sum_{i=1}^n \frac{\partial y_i}{\partial x_j} (0, \dots, 0) \bar{y}_i \right) = \text{Var}(\bar{y}) \sum_{i=1}^n \left(\frac{e^{\mu w_{ij}}}{\sum_{k=1}^n e^{\mu w_{ik}}} \right)^2,$$

which implies that to get $\text{Var}(\bar{x}_j) \approx \text{Var}(\bar{y})$ we need

$$\sum_{i=1}^n \left(\frac{e^{\mu w_{ij}}}{\sum_{k=1}^n e^{\mu w_{ik}}} \right)^2 \approx 1.$$

or

$$\sum_{i=1}^n (e^{\mu w_{ij}})^2 \approx \left(\sum_{i=1}^n e^{\mu w_{ij}} \right)^2 \quad (10)$$

for all $j = 1, \dots, m$. We can satisfy this condition exactly by choosing a single $I_j = 1, \dots, n$ for each $j = 1, \dots, m$ and set that weight to $w_{I_j, j} = 0$ and set the other weights to $w_{ij} = -\text{sgn}(\mu)\infty$. In that case $e^{\mu w_{ij}} = 1$ if $i = I_j$ and $e^{\mu w_{ij}} = 0$ if not and consequently both sides of (10) are equal to 1.

Satisfying both forward (9) and backward condition (10) is only possible in the case that $m = n$ where we can set $I_j = j$ and $J_i = i$. But even then, this scheme is not satisfactory. If we initialize parameters to $\pm\infty$ (depending on the sign of μ) they can not be changed by the addition of finite numbers thus making training impossible. We might then be tempted to substitute $\pm\infty$ by a very large negative or positive value to initialize with. But this still causes a problem for training.

Consider the loss gradient for an individual parameter:

$$\bar{w}_{ij} = \frac{e^{\mu(x_j + w_{ij})}}{\sum_{k=1}^m e^{\mu(x_k + w_{ik})}} \bar{y}_i. \quad (11)$$

If we try to satisfy (9) and (10) then we would set $w_{ik} = 0$ for some $k = 1, \dots, n$ for every $i = 1, \dots, n$. This implies that the denominator in (11) is greater than or equal to 1. If we subsequently set $w_{ij} = -\text{sgn}(\mu)K$ for $j \neq k$ and K a very large number then the numerator of (11) becomes vanishingly small, possibly zero in floating point format. This essentially freezes the parameter's value since any updates applied to it would become practically zero.

Assuming the inputs $(x_j)_{j=1}^m$ are centered, the expected value of the fraction in (11) is given by

$$\frac{e^{\mu w_{ij}}}{\sum_{k=1}^m e^{\mu w_{ik}}}, \quad (12)$$

which gives values in the range $[0, 1]$ for any $i = 1, \dots, n$ and $j = 1, \dots, m$. We can avoid this fraction becoming (effectively) zero from the start by initializing all the parameters to roughly the same value, i.e. for all $i = 1, \dots, n$ set

$$w_{i1} \approx w_{i2} \approx \dots \approx w_{im}. \quad (13)$$

In this case we achieve the *maximum of the minimum* where (12) gives the value $1/m$ for all $i = 1, \dots, n$ and $j = 1, \dots, m$.

Satisfying the forward condition (9), backward condition (10) and parameter trainability condition (13) all at the same time is not possible. However, the fair tropical initialization scheme does present a reasonable trade-off between these three clashing requirements. Indeed if we initialize using this scheme (assuming $n > m$ for the moment) we get a matrix like (6), where we let $K > 0$ if $\mu > 0$ and $K < 0$ if $\mu < 0$. First, most entries have the same $-K$ value and so satisfy (13) to some degree. Second, in every row we have a single element (the zero values) that are relatively dominant with respect to the other elements in the row, thus catering to (9). Third, in every column we have at most $\lceil n/m \rceil$ zero values that are relatively dominant with respect to the other elements in the column, thus catering to (10).

4.4 Results

We perform 10 training runs for each type of network on each dataset using the training setup we described. We measure performance by the accuracy (mean \pm standard deviation over the runs) on the testing dataset that was not seen during training. We use the same training/testing split of the dataset for every type of network. The results are summarized in Table 1.

Model / dataset	Iris	Heartdisease	Circles	Spheres	FashionMNIST
ReLU	97.14 \pm 0.62	83.93 \pm 2.16	84.50 \pm 0.39	80.91 \pm 1.62	83.82 \pm 0.35
maxplus	97.52 \pm 1.01	83.50 \pm 2.27	84.84 \pm 0.86	81.69 \pm 0.42	83.50 \pm 0.34
minplus	97.62 \pm 0.49	82.84 \pm 1.22	84.91 \pm 0.35	81.61 \pm 0.60	83.39 \pm 0.24
logplus $\mu = -10$	97.58 \pm 1.00	81.72 \pm 1.62	85.06 \pm 0.25	81.52 \pm 0.71	83.46 \pm 0.37
logplus $\mu = -1$	97.90 \pm 0.39	83.26 \pm 2.15	73.92 \pm 6.84	69.41 \pm 5.05	83.50 \pm 0.28
logplus $\mu = 1$	97.97 \pm 0.49	82.38 \pm 1.73	75.06 \pm 7.89	67.28 \pm 5.53	83.46 \pm 0.16
logplus $\mu = 10$	97.46 \pm 0.62	81.86 \pm 1.15	85.16 \pm 0.26	81.62 \pm 0.45	83.56 \pm 0.37
Parameters	60	5328	2336	2336	2288

Table 1: Accuracy (mean \pm standard deviation) of the trained fully connected networks on the testing sets of the various classification datasets. The best result for each dataset is indicated in **purple**, the second best result in **green**.

The semiring activation networks manage to modestly outperform the classic network in 3 of the 5 cases and are only slightly behind in the other 2 cases. A standout failure is the logarithmic semiring networks with $\mu \in \{-1, 1\}$ in the *circles* and *spheres* datasets. We can explain this based on the fact that these are low-dimensional problem (2 respectively 3 input dimensions) where there is a sharp boundary between the classes. At the same time the logarithmic maps are fairly gradual for small absolute values of μ and weight decay keeps the parameters from becoming large enough to compensate for that. Consequently, the logarithmic networks have a hard time separating classes that are close together.

On the whole though we can conclude that the semiring approach is at least as viable as the standard approach for these smaller problems. We can now try to see if we can scale up to some larger networks.

5. In convolutional neural networks

We perform a final experiment to see how viable replacing activation functions with semiring activation is in an existing modern and optimized network architecture. We start from the *ConvNeXt* network from Liu et al. (2022), specifically the *atto* variant available from Wightman (2019), and train it on FashionMNIST.

5.1 Architecture modifications

The core *block* at the heart of a ConvNeXt consists of two phases:

1. a *depthwise convolution* that applies a kernel per feature map and

2. a *normalization* followed by a *reverse bottleneck MLP* that is applied *pixel-wise*,

the result of these two operations is added to the input to form a residual connection, see Figure 2a.

Our adaptation keeps the (linear) depthwise convolution and normalization but replaces the MLP with a concatenation of a non-linear semiring operator and a linear operator. The semiring operation does a $4\times$ fan-out while the linear operation reduces again to the original amount of channels, thus retaining the same reverse bottleneck of the original, see Figure 2b.

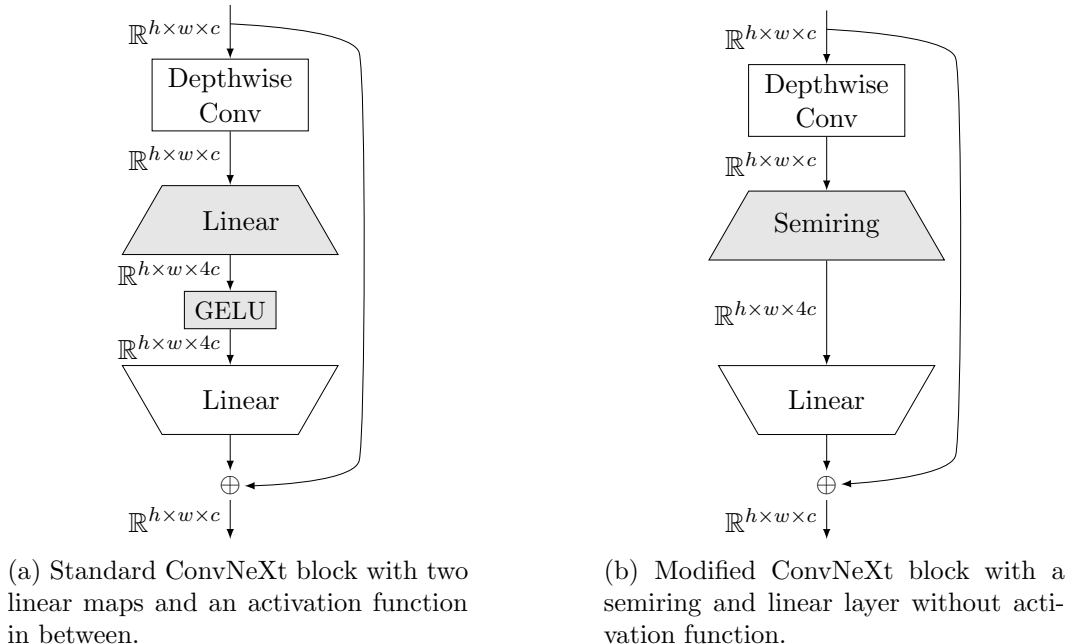


Figure 2: Standard and semiring-based ConvNeXt (Liu et al., 2022) blocks compared. Normalization and dropout modules are omitted.

5.2 Results

We compare the performance of the semiring based networks against the baseline network using an MLP. We perform 10 training runs per model and record accuracy on the test dataset. We train 4 semiring based models: maxplus, minplus, logplus($\mu = -1$) and logplus($\mu = 1$). We train all models for the same 50 epochs and batchsize of 512 and adapt the rest of the training setup to each model. Details for the training setup can be found in Appendix 6. The results are summarized in Table 2.

Training with these (relatively) large scale networks proved more challenging than with the previous toy networks. In particular the logarithmic networks proved challenging and required us to remove the affine transform from the normalization layer to avoid numeric stability issues. The problem being that in single precision floating point, the function $x \mapsto e^x$ already overflows at $x = 89$. This would not happen if the input x is normalized but is likely enough to happen if we let x be an affine transform of a normalized input. The

overflow would only need to happen at one place in the network, after which ‘inf’ values would propagate throughout the network and ruin training. As we already experienced in the previous experiments, both the tropical and logarithmic variants proved to be more sensitive to the training hyperparameters than the baseline network. This sensitivity makes finding a good training setup harder for the semiring networks than for the baseline network.

Feed-forward type	Test accuracy(%)	Train accuracy(%)	Gap(%)
MLP (linear-GELU-linear)	91.24 \pm 0.13	99.97 \pm 0.07	-8.63
maxplus-linear	89.08 \pm 0.29	93.26 \pm 2.32	-4.18
minplus-linear	89.93 \pm 0.24	94.74 \pm 2.79	-4.81
logplus-linear $\mu = -1$	88.15 \pm 0.15	91.72 \pm 2.52	-3.57
logplus-linear $\mu = 1$	88.31 \pm 0.23	91.60 \pm 2.52	-3.29

Table 2: Accuracy(%) (mean \pm standard deviation) of the trained ConvNeXt models on the FashionMNIST test dataset and on the last 100 training batches. The generalization gap is the difference between the two mean accuracies.

As the second column of Table 2 shows, both the tropical and logarithmic network’s accuracy falls significantly short of the baseline network. This shortfall does not bode well for the semiring idea but there are some nuances to be made.

Looking at the third column of Table 2 we see that the baseline network has saturated its performance on the training data. During the last 100 batches the baseline network has virtually 100% training accuracy, consequently it can not benefit much more from further training. At the same time there is a fairly large gap between the training and testing accuracy as can be seen in the last column of Table 2. On the side of the semiring networks we see that after 50 epochs there is still significant room for improvement on the training data and that the gap between the testing and training accuracy is much more modest.

We conjecture that this difference is partly explained by the standard training regime being very well suited for the baseline network but that the modifications we made—chiefly the parameter initialization scheme—are not sufficient to extract maximum performance from the semiring networks. Indeed, the methods for neural network training have evolved much over the last decade with much research into optimizers, schedulers, initialization, regularization, etc., all focused on the linear with activation type networks. It would not be unreasonable to assume similar efforts could yield similar progress in training semiring based networks.

6. Discussion & Concluding remarks

Viability. In this article we have proposed a general framework for constructing trainable nonlinearities for neural networks. We constructed several such trainable nonlinearities based on the tropical and logarithmic semirings and introduced an associated parameter initialization scheme. We did a series of experiments to show the viability of replacing the traditional activation function in neural networks with nonlinearities based on the aforementioned semirings.

Unrealized potential. Our experiments showed that the semiring approach is viable in that we can get very good results for small networks and decent but not state-of-art results for larger networks. With regard to the larger scale experiment we concluded that there is more performance on the table for semiring networks that can potentially be unlocked by designing more suitable optimizers, schedulers, initialization schemes, regularization schemes, etc. Whether such a line of research would be worthwhile is debatable for two reasons.

First, while our experiments show viability, they do not immediately show a clear advantage over traditional neural networks with activation functions.

Second, the current development of deep learning hardware (Dhilleswararao et al., 2022) focuses to a large degree on optimizing linear operations. Most GPUs in general use today already contain hardware dedicated to linear matrix multiplication (Markidis et al., 2018), this makes linear computations much more efficient than doing semiring computations that would have to be executed by general purpose computing units at a higher cost in time and energy.

Conclusion on PDE-G-CNNs. Recall that our interest in the semiring structure was instigated by our previous research into PDE-G-CNNs (Smets et al., 2023). In those networks we also had trainable nonlinear operators based on the tropical semiring, but these operators were further structured to satisfy equivariant and PDE properties. We asked whether the benefits from PDE-G-CNNs were perhaps not largely a consequence of the tropical semiring structure rather than the equivariant and PDE structures.

We can now answer this question in the negative. Without the equivariance and PDE constraints the tropical operator becomes very hard to deal with and a challenge to train. PDE-G-CNNs can be successfully trained without much adaptation of the training regime and are not unusually sensitive to hyperparameters. PDE-G-CNNs also do not require more epochs to saturate their training data than normal CNNs, which was clearly an issue for the tropical semiring networks in our ConvNeXt experiment.

Overall, we conclude that the equivariant and PDE structures play an important role in making PDE-G-CNNs work and the semiring structure is not—by itself—the cause of the benefits of PDE-G-CNNs.

Appendix A. Hyperparameters for the fully connected experiments

Hyperparameter	Iris	Heartdisease	Circles	Spheres	FashionMNIST
Epochs	40	40	100	100	40
Batchsize	8	16	32	16	512
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW
Scheduler	1-cycle cos	1-cycle cos	1-cycle cos	1-cycle cos	1-cycle cos
Learning rate (linear)	0.020	0.010	0.020	0.020	0.008
Learning rate (tropical)	0.004	0.008	0.010	0.010	0.040
Learning rate (logarithmic)	0.040	0.008	0.008	0.008	0.040
Weigh decay	0.01	0.05	0.01	0.01	0.01
Warmup epochs	18	18	45	45	18
Warmup factor	1/10	1/10	1/10	1/10	1/10
Annihilation factor	1/1000	1/1000	1/1000	1/1000	1/1000
Internal width w	4	48	16	32	8
RNG seed	42	42	42	42	42
Parameters	60	5328	2336	2336	2288

Table 3: Dataset dependent hyperparameters for the fully connected experiments.

Appendix B. Hyperparameters for the ConvNeXt experiments

Hyperparameter	MLP	tropical-linear	logplus-linear
Epochs	50	50	50
Batchsize	512	512	512
Affine LayerNorm	Yes	Yes	No
Optimizer	AdamW	AdamW	AdamW
Scheduler	1-cycle cos	1-cycle cos	1-cycle cos
Learning rate (linear)	$4 \cdot 10^{-3}$	$4 \cdot 10^{-3}$	$6 \cdot 10^{-3}$
Learning rate (semiring)	-	$1 \cdot 10^{-3}$	$5 \cdot 10^{-4}$
Weight decay (linear)	$5 \cdot 10^{-3}$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
Weight decay (semiring)	-	$1 \cdot 10^{-2}$	$1 \cdot 10^{-2}$
Warmup epochs	5	5	5
Warmup factor (linear)	1/25	1/50	1/50
Warmup factor (semiring)	-	1/50	1/50
Annihilation factor	1/500	1/250	1/250
RNG seed	42	42	42
Initialization (linear)	Kaiming	Kaiming	Kaiming
Initialization (semiring)	-	Tropical fair	Tropical fair
Parameters	3, 375, 850	3, 375, 850	3, 372, 170

Table 4: Hyper parameters settings for the ConvNeXt experiments.

References

- E. Anderson. The species problem in iris. *Annals of the Missouri Botanical Garden*, 23(3): 457–509, 1936.
- A. Apicella, F. Donnarumma, F. Isgrò, and R. Prevete. A survey on modern trainable activation functions. *Neural Networks*, 138:14–32, 2021.
- P. Dhilleswararao, S. Boppu, M. S. Manikandan, and L. R. Cenkeramaddi. Efficient hardware architectures for accelerating deep neural networks: Survey. *IEEE access*, 10: 131788–131828, 2022.
- S. Elfving, E. Uchibe, and K. Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.
- X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- I. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. In *International conference on machine learning*, pages 1319–1327. PMLR, 2013.
- K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A convnet for the 2020s, 2022.
- I. Loshchilov and F. Hutter. Decoupled weight decay regularization, 2019.
- S. Markidis, S. W. Der Chien, E. Laure, I. B. Peng, and J. S. Vetter. Nvidia tensor core programmability, performance & precision. In *2018 IEEE international parallel and distributed processing symposium workshops (IPDPSW)*, pages 522–531. IEEE, 2018.
- J. Masci, J. Angulo, and J. Schmidhuber. A learning framework for morphological operators using counter-harmonic mean. In *Mathematical Morphology and Its Applications to Signal and Image Processing: 11th International Symposium, ISMM 2013, Uppsala, Sweden, May 27-29, 2013. Proceedings 11*, pages 329–340. Springer, 2013.
- D. Mellouli, T. M. Hamdani, M. B. Ayed, and A. M. Alimi. Morph-cnn: A morphological convolutional neural network for image classification. In *Neural Information Processing: 24th International Conference, ICONIP 2017, Guangzhou, China, November 14-18, 2017, Proceedings, Part II 24*, pages 110–117. Springer, 2017.
- G. Naitzat, A. Zhitnikov, and L.-H. Lim. Topology of deep neural networks. *The Journal of Machine Learning Research*, 21(1):7503–7542, 2020.
- P. Ramachandran, B. Zoph, and Q. V. Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.

- G. X. Ritter and P. Sussner. An introduction to morphological neural networks. In *Proceedings of 13th International Conference on Pattern Recognition*, volume 4, pages 709–717. IEEE, 1996.
- J. Serra. *Image analysis and mathematical morphology*. Academic press, 1982.
- Y. Shen, X. Zhong, and F. Y. Shih. Deep morphological neural networks, 2019.
- F. Y. Shih, Y. Shen, and X. Zhong. Development of deep learning framework for mathematical morphology. *International Journal of Pattern Recognition and Artificial Intelligence*, 33(06):1954024, 2019.
- B. Simmons II. *Investigating Heart Disease Datasets and Building Predictive Models*. PhD thesis, Elizabeth City State University, 2021.
- B. M. Smets, J. Portegies, E. J. Bekkers, and R. Duits. Pde-based group equivariant convolutional neural networks. *Journal of Mathematical Imaging and Vision*, 65(1):209–239, 2023.
- L. N. Smith and N. Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pages 369–386. SPIE, 2019.
- L. R. Sütfeld, F. Brieger, H. Finger, S. Füllhase, and G. Pipa. Adaptive blending units: Trainable activation functions for deep neural networks, 2018.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- R. Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.