



HAL
open science

Max-plus Operators Applied to Filter Selection and Model Pruning in Neural Networks

Yunxiang Zhang, Samy Blusseau, Santiago Velasco-Forero, Isabelle Bloch,
Jesus Angulo

► **To cite this version:**

Yunxiang Zhang, Samy Blusseau, Santiago Velasco-Forero, Isabelle Bloch, Jesus Angulo. Max-plus Operators Applied to Filter Selection and Model Pruning in Neural Networks. International Symposium on Mathematical Morphology, Jul 2019, Saarbrücken, Germany. hal-02071450v1

HAL Id: hal-02071450

<https://hal.science/hal-02071450v1>

Submitted on 18 Mar 2019 (v1), last revised 1 Apr 2019 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Max-plus Operators Applied to Filter Selection and Model Pruning in Neural Networks

Yunxiang Zhang^{1,2,3}, Samy Blusseau³, Santiago Velasco-Forero³, Isabelle Bloch², and Jesús Angulo³

¹ Ecole Polytechnique, 91128 Palaiseau, France

² LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France

³ Centre for Mathematical Morphology, Mines ParisTech, PSL Research University, France

Abstract. Following recent advances in morphological neural networks, we propose to study in more depth how Max-plus operators can be exploited to define morphological units and how they behave when incorporated in layers of conventional neural networks. Besides showing that they can be easily implemented with modern machine learning frameworks, we confirm and extend the observation that a Max-plus layer can be used to select important filters and reduce redundancy in its previous layer, without incurring performance loss. Experimental results demonstrate that the filter selection strategy enabled by a Max-plus is highly efficient and robust, through which we successfully performed model pruning on different neural network architectures. We also point out that there is a close connection between Maxout networks and our pruned Max-plus networks by comparing their respective characteristics. The code for reproducing our experiments is available online⁴.

Keywords: Mathematical morphology · Morphological neural networks · Max-plus operator · Deep learning · Filter Selection · Model Pruning.

1 Introduction

During the previous era of high interest in artificial neural networks, in the late 1980s, morphological networks [4,25] were introduced merely as an alternative to their linear counterparts. In a nutshell, they consist in changing the elementary neural operation from the usual scalar product to dilations, erosions and more general rank filters. Besides the practical achievements in this research line, which reached state-of-the-art results at its time [18], it questioned the main practices in the field of artificial neural networks on crucial topics such as architectures and optimization methods, improving their understanding.

After the huge technological leap taken by deep neural networks in the last decade, understanding them remains challenging and insights can still be brought by testing alternatives to the most popular practices. However, since the emergence of highly efficient deep learning frameworks (Caffe, Torch, TensorFlow,

⁴ For code release, please visit <https://github.com/yunxiangzhang>.

etc.), morphological neural networks have been very little re-investigated [2,16]. Motivated by the promising results in [2], this paper proposes to extend and deepen the study on morphological neural networks with Max-plus layers. More precisely, we aim to validate and exploit a specific property, namely that Max-plus layer (i.e. a dilation layer) following a conventional linear layer tends to select a reduced number of filters from the latter, making the others useless.

Our findings and contributions are three-fold. First, we propose an efficient training framework for morphological neural networks with Max-plus blocks, and theoretically demonstrate that they are universal function approximators under mild conditions. Secondly, we perform extensive experiments and visualization analysis to validate the robustness and effectiveness of the filter selection property provided by Max-plus blocks. Thirdly, we successfully applied the resulting Max-plus blocks to the task of model pruning on different neural network architectures. Related work is briefly summarized in Section 2. Then we show how to introduce Max-plus operators in neural networks in Section 3. Experimental results are discussed in Section 4.

2 Related Work

Morphological neural networks were defined almost simultaneously in two different ways at the end of the 1980s [4,25]. Davidson [4] introduced neural units that can be seen as pure dilations or erosions, whereas Wilson [25] focused on a more general formulation based on rank filters, in which max and min operators are two particular cases. Davidson’s definition interprets morphological neurons as bounding boxes in the feature space [19,20,23]. In the latter studies, networks were trained to perform perfectly on training sets after few iterations, but little attention was drawn to generalization. Only recently, a backpropagation-based algorithm was adopted and improved constructive ones [26]. Still, the “bounding-box” approach does not seem to generalize well to test set when faced with high-dimensional problems like image analysis.

Wilson’s idea, on the other hand, inspired hybrid linear/rank filter architectures, which were trained by gradient descent [17] and backpropagation [18]. In this case, the geometrical interpretation of decision surfaces becomes much richer, and the resulting framework was successfully applied to an image classification problem. The previously mentioned study [2] is one of the latest in this area, introducing a hybrid architecture that experimentally shows an interesting property on network pruning. In this classification experiment [2], each Max-plus unit shows, after training, one or two large weight values compared to the others. At the inference stage, this non-uniform distribution of weight values induces a selection of important filters in the previous layer, whereas the other filters (the majority) are no longer used in the subsequent classification task. Therefore, after removal of the redundant filters, the network behaves exactly as it did before pruning. In this paper, we focus on the exploration of this property and show that it becomes more stable and effective provided that a proper regularization is applied.

Model pruning stands for a family of algorithms that explore the redundancy in model parameters and remove unimportant ones while not compromising the performance. Rapidly increased computational and storage requirements for deep neural networks in recent years have largely motivated research interests in this domain. Early works on model pruning dates back to [9,12], which prune model parameters based on the Hessian of the loss function. More recently, the model pruning problem was addressed by first dropping the neuronal connections with insignificant weight value and then fine-tuning the resulting sparse network [8]. This technique was later integrated into the Deep Compression framework [7] to achieve even higher compression rate. The HashNets model [3] randomly groups parameters into hash buckets using a low-cost hash function and performs model compression via parameter sharing. While previous methods achieved impressive performance in terms of compression rate, one notable disadvantage of these non-structured pruning algorithms lies in the sparsity of the resulting weight matrices, which cannot lead to speedup without dedicated hardwares or libraries.

Various structured pruning methods were proposed to overcome this subtlety in practical applications by pruning at the level of channels or even layers. Filters with smaller L_1 norm were pruned based on a pre-defined pruning ratio for each layer in [13]. Model pruning was transformed into an optimization problem in [15] and the channels to remove were determined by minimizing next layer reconstruction error. A branch of algorithms in this category employs L_p regularization to induce shrinkage and sparsity in model parameters. Sparsity constraints were imposed in [14] on channel-wise scaling factors and pruning was based on their magnitude, while in [24] group-sparsity was leverage to learn compact CNNs via a combination of L_1 and L_2 regularization. One minor drawback of these regularization-based methods is that the training phase generally requires more iterations to converge. Our approach also falls into the structured pruning category and thus no dedicated hardware or libraries are required to achieve speedup, yet no regularization is imposed during model training. Moreover, in contrast to most existing pruning algorithms, our method does not need fine-tuning to regain performance.

3 Max-plus Operator as a Morphological Unit

3.1 Morphological Perceptron

In traditional literature on machine learning and neural networks, a perceptron [21] is defined as a linear computational unit, possibly followed by a non-linear activation function. Among all popular choices of activation functions, such as logistic function, hyperbolic tangent function and rectified linear unit (ReLU) function, ReLU [5] generally achieves better performance due to its simple formulation and non-saturating property. Instead of multiplication and addition, the morphological perceptron employs addition and maximum, which results in a non-linear computational unit. A simplified version [2] of the initial formulation [4,19] is defined as follows.

Definition 1. (Morphological Perceptron). *Given an input vector $\mathbf{x} \in \mathbb{R}_{max}^n$ (with $\mathbb{R}_{max} = \mathbb{R} \cup \{-\infty\}$), a weight vector $\mathbf{w} \in \mathbb{R}_{max}^n$, and a bias $\mathbf{b} \in \mathbb{R}_{max}$, the morphological perceptron computes its activation as:*

$$a(\mathbf{x}) = \max \left\{ \mathbf{b}, \max_{i \in \{1, \dots, n\}} \{ \mathbf{x}_i + \mathbf{w}_i \} \right\} \quad (1)$$

where \mathbf{x}_i (resp. \mathbf{w}_i) denotes the i -th component of \mathbf{x} (resp. \mathbf{w}).

This model may also be referred to as $(\max, +)$ perceptron since it relies on the $(\max, +)$ semi-ring with underlying set \mathbb{R}_{max} . It is a dilation on the complete lattice $(\mathbb{R} \cup \{\pm\infty\})^n, \leq_n$ with \leq_n the Pareto ordering.

3.2 Max-plus Block

Based on the formulation of the morphological perceptron, we define the Max-plus block as a standalone module that combines a fully-connected layer (or convolutional layer) with a Max-plus layer [2]. Let us denote the input vector of the fully-connected layer⁵, the input and output vectors of the Max-plus layer respectively by \mathbf{x} , \mathbf{y} and \mathbf{z} , whose components are indexed by $i \in \{1, \dots, I\}$, $j \in \{1, \dots, J\}$ and $k \in \{1, \dots, K\}$, respectively. The corresponding weight matrices are denoted by $\mathbf{w}^f \in \mathbb{R}_{max}^{I \times J}$ and $\mathbf{w}^m \in \mathbb{R}_{max}^{J \times K}$. Then the operation performed in this Max-plus block is (see Figure 1):

$$\begin{aligned} \mathbf{y}_j &= \sum_{i \in \{1, \dots, I\}} \mathbf{x}_i \cdot \mathbf{w}_{ij}^f \\ \mathbf{z}_k &= \max_{j \in \{1, \dots, J\}} \{ \mathbf{y}_j + \mathbf{w}_{jk}^m \} \end{aligned} \quad (2)$$

Note that the bias vector of the fully-connected layer (convolutional layer) is removed in our formulation, since its effect overlaps with that of the weight matrix \mathbf{w}^m . In addition, the bias vector of the Max-plus layer is shown to be ineffective in practice and is therefore not used here.

3.3 Universal Function Approximator Property

As shown in [6], Maxout networks with enough affine components in each Maxout unit are universal function approximators. Recall that a model is called a universal function approximator if it can approximate arbitrarily well any continuous function provided enough capacity. Similarly, provided that the input vector (or input feature maps) $\mathbf{y} \in \mathbb{R}_{max}^J$ of the Max-plus layer may have arbitrarily many affine components (or affine feature maps), we show that a Max-plus model with just two output units in its Max-plus block can approximate arbitrarily well any continuous function of the input vector (or input feature maps) $\mathbf{x} \in \mathbb{R}^I$ of the block on a compact domain. A diagram illustrating the basic idea of the proof is shown in Figure 1.

⁵ This formulation can be easily generalized to the case of convolutional layers.

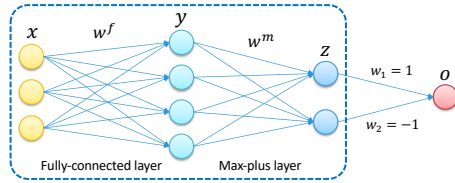


Fig. 1. A Max-plus block with two output units, applied to an input vector x .

Theorem 1. (Universal function approximator). *A Max-plus model with two output units in its Max-plus block can approximate arbitrarily well any continuous function of the input of the block on a compact domain.*

Proof. We provide here a sketch of the proof. By Proposition 4.2 in [6], any continuous function defined on a compact domain $C \subset \mathbb{R}^I$ can be approximated arbitrarily well by a piecewise linear (PWL) continuous function g , composed of k affine regions. By Proposition 4.1 in [6], there exist two matrices $W_1, W_2 \in \mathbb{R}^{k \times I}$ and two vectors $b_1, b_2 \in \mathbb{R}^k$ such that

$$\forall x \in C, g(x) = \max_{1 \leq j \leq k} \{W_{1j}x + b_{1j}\} - \max_{1 \leq j \leq k} \{W_{2j}x + b_{2j}\}, \quad (3)$$

where W_{ij} is the j -th row of matrix W_i and b_{ij} the j -th coefficient of b_i , $i = 1, 2$. Now, Equation 3 can be exactly the output of the Max-plus block of Figure 1, by setting $\mathbf{w}^f = [W_1; W_2] \in \mathbb{R}^{2k \times I}$ as the matrix of the fully connected layer, $\mathbf{w}_1^m = [b_1^T, -\infty, \dots, -\infty] \in \mathbb{R}_{max}^{2k}$ and $\mathbf{w}_2^m = [-\infty, \dots, -\infty, b_2^T] \in \mathbb{R}_{max}^{2k}$ the two rows of \mathbf{w}^m . This concludes the proof.

4 Experiments

In this section, we present experimental results for our Max-plus blocks by integrating them in different types of neural networks. All neural network models shown in this section are implemented with the open-source machine learning library TensorFlow [1] and trained on benchmark datasets MNIST [11] or CIFAR-10 [10] depending on the model complexity and capability.

4.1 Filter Selection Property

In an attempt to reproduce and confirm the experimental results reported in [2], we first implemented a simple Max-plus model composed of a fully-connected layer with J units followed by a Max-plus layer with 10 units, namely a Max-plus block in our terminology, to perform image classification on MNIST dataset. In contrast to the original formulation in [2], the two bias vectors are removed for practical concerns explained in Section 3.2.

Table 1 summarizes the classification accuracy achieved by this simple model on the validation set of MNIST dataset for different values of J in columns 3

to 5. Note that all the experiments contained in these 3 columns are conducted under the same training setting (initial learning rate, learning rate decay steps, batch size, optimizer, etc.) except for parameter initialization (each column corresponds to a different random seed). The performance of the original model in [2] is included in column 2 for comparison. As shown in Table 1, provided a proper initialization, our simple Max-plus model generally achieves an improved performance compared to the original model. More interestingly, through horizontal comparison across different runs, we find that the performance of this naive Max-plus model is highly sensitive to parameter initialization.

Table 1. Classification accuracy on the validation set of MNIST dataset.

J units	[2]	Max-plus			Max-plus + dropout		
24	84.3%	76.7%	84.4%	76.0%	93.7%	93.7%	94.2%
32	84.8%	84.5%	76.5%	94.0%	94.6%	94.2%	93.7%
48	84.6%	94.0%	93.8%	75.8%	94.6%	94.5%	94.5%
64	92.1%	85.4%	94.7%	94.9%	94.8%	95.1%	94.8%
100	–	94.8%	85.3%	93.7%	94.8%	95.5%	95.2%
144	–	95.1%	85.8%	85.4%	95.3%	95.7%	95.9%

In order to gain more insight into this instability problem, we follow the approach of [2] and visualize the weight matrix of the Max-plus layer $\mathbf{w}^m \in \mathbb{R}_{max}^{J \times 10}$ by splitting and reshaping it into 10 gray-scale images $\mathbf{w}_1^m \dots \mathbf{w}_i^m \dots \mathbf{w}_{10}^m$, each corresponding to a specific class. In addition, we also visualized the 10 linear filters from the fully-connected layer that correspond to the maximum value of each weight vector \mathbf{w}_i^m , defined as:

$$\text{Filter}_i = \mathbf{w}_{j_{\max}^{(i)}}^f \in \mathbb{R}_{max}^I \quad \text{where } j_{\max}^{(i)} = \arg \max_{j \in \{1, \dots, J\}} \{\mathbf{w}_{ji}^m\} \quad (4)$$

Figure 2 shows the visualization results for the Max-plus model that achieves an accuracy of 85.8% with $J = 144$ units. We notice that there exists a severe filter-collision problem in the Max-plus block, namely different output units select the same linear filter in the fully-connected layer to compute their outputs. More specifically, class 3 and class 8 share the same *argmax* (highlighted by red circles) in their weight vectors and consequently select the same linear filter (visualized filters for class 3 and class 8 are the same). This collision between output units directly leads to classification confusion (because the Max-plus layer is the last layer in this simple model) since the classes that employ the same linear filter are completely indistinguishable. Furthermore, we only observed this filter-collision problem in the experiments that achieve relatively poor performance compared to the others (same model with different initialization). This observation consistently verifies our hypothesis that filter-collision is at the root of lower performance.

In order to separate the output units that got stuck with the same linear filter, we applied a dropout regularization [22] to randomly switch off the neu-

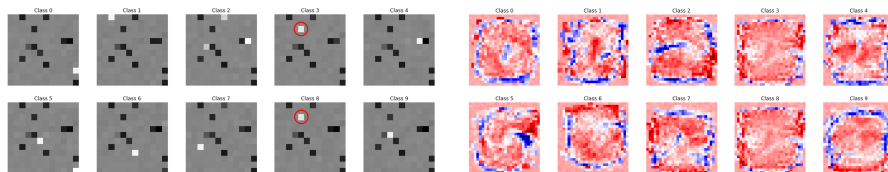


Fig. 2. Visualization of the weight matrix \mathbf{w}^m (left) and the 10 linear filters that correspond to the maximum value of each weight vector \mathbf{w}_i^m (right).

ronal connections between the fully-connected layer and the Max-plus layer during training. Empirically, we found this approach highly effective, the performance of dropout-regularized Max-plus model becomes much less sensitive to parameter initialization. Now if we perform the same visualization as before, we no longer observe the filter-collision problem. The performance of dropout-regularized Max-plus model is shown in columns 6 to 8 (each corresponds to a non-regularized column) in Table 1. There exist more significant values (bright points) in each weight vector \mathbf{w}_i^m of the dropout-regularized Max-plus model. Visualization of the linear filters that correspond to these large values shows similar images to those in Figure 2 (right), i.e. digit-like shape with high contrast. Now if we visualize the linear filters that correspond to smaller values in the weight matrix \mathbf{w}^m , we can only get noisy images or low-contrast images showing the shape of a specific digit, which means that these filters were activated by few training examples. Figure 3 (left) shows several linear filters of this kind.



Fig. 3. Visualization of several linear filters corresponding to small values in \mathbf{w}^m (left) and comparison between the activations of the fully-connected layer for 2 training examples and their corresponding weight vectors in \mathbf{w}^m (right).

We further validate the effect provided by dropout regularization through an additional experiment. With $J = 144$, we carried out 25 runs with different initializations for different dropout ratios. As shown in Figure 4, dropout reduces dramatically the variability across experiments. Interestingly, we observe a slight accuracy drop when dropout ratio surpasses a certain level, which indicates that a trade-off between stability and performance needs to be found.

The visualization above suggests that large values in \mathbf{w}^m correspond to linear filters of the fully-connected layer that contain rich information for the subsequent classification task, and hence that strongly respond to the samples of a specific class. If that holds, then it is likely that only these filters shall achieve the maximum in Equation 2, and contribute to the classification output. There-

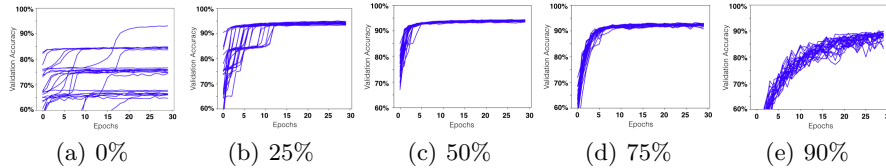


Fig. 4. Classification accuracy per epochs for 25 runs with different random initializations on the MNIST validation set with different dropout ratios.

fore, we might be able to reduce the complexity of a model while not degrading its performance by exploiting this filter selection property of Max-plus layers.

In order to further validate the stability of this connection before taking advantage of it, we also visualized the activations of the fully-connected layer as gray-scale images for several training examples and compare⁶ them to the visualization of \mathbf{w}^m . As shown in Figure 3 (right), there is a clear correspondence between the maximum activation value of a training example with label i and the largest two or three values of the weight vector $\mathbf{w}_{\cdot(i+1)}^m$, which indicates that the linear filters corresponding to large values in the weight matrix \mathbf{w}^m are effectively used for the subsequent classification task.

4.2 Application to Model Pruning

Now that our approach to filter selection via Max-plus layers is proved to be quite effective and stable, we formalize our model pruning strategy as follows: given a fixed threshold $s \in [0, 1]$, for each weight vector $\mathbf{w}_{\cdot(i+1)}^m$, we only keep the values that are larger than $s \times \max_{j \in \{1, \dots, J\}} \{ \mathbf{w}_{j(i+1)}^m \} + (1 - s) \times \min_{j \in \{1, \dots, J\}} \{ \mathbf{w}_{j(i+1)}^m \}$ and the linear filters that correspond to these retained values. Therefore, if in total J_r linear filters are kept in the pruned model, then the remaining parameters in the Max-plus layer no longer form a weight matrix but a weight vector of size J_r , where each entry corresponds to a linear filter in the fully-connected layer. The pruned fully-connected layer and the pruned Max-plus layer combined together perform a standard linear transformation followed by a maximum operation over uneven groups. Note that the pruning process is conducted independently for each output unit \mathbf{z}_k of the Max-plus block ($1 \leq k \leq 10$), thus the number of retained linear filters in the pruned model may vary from one class to another. From now on, we shall call these selected linear filters by *active filters* and the others by *non-active filters*. Figure 5 shows a graphical illustration of the comparison between the original model and the pruned model.

We tried different pruning levels on this simple Max-plus model by varying the threshold s and tested the pruned models on the validation set and test set

⁶ For example, if a training example has label i , then we compare its activation vector $\mathbf{y} \in \mathbb{R}_{max}^J$ with the weight vector $\mathbf{w}_{\cdot(i+1)}^m \in \mathbb{R}_{max}^J$.

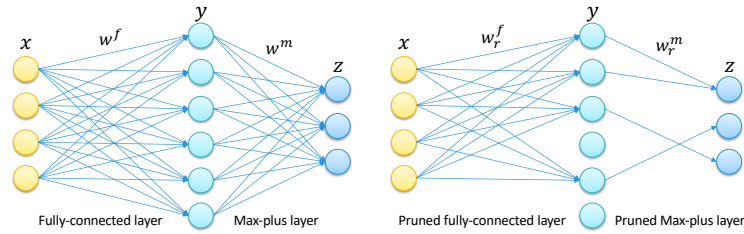


Fig. 5. Illustration of the comparison between the original Max-plus model (left) and the pruned Max-plus model (right). Here we have $\mathbf{w}_r^f \in \mathbb{R}_{max}^{I \times J_r}$ and $\mathbf{w}_r^m \in \mathbb{R}_{max}^{J_r \times J}$.

of MNIST dataset. We plotted the resulting classification accuracy in function of the number of active filters in Figure 6. The performance of a single-layer softmax model and a single-layer Maxout model (number of affine components in each Maxout unit is 2) is also provided for comparison.

J_r	Accuracy	J_r	Accuracy
10	82.7%	18	90.4%
11	86.4%	19	94.6%
12	86.4%	20	94.7%
13	89.9%	21	94.7%
14	90.0%	22	94.8%
15	90.2%	23	94.8%
16	90.2%	24	95.7%
17	90.3%	25	95.7%

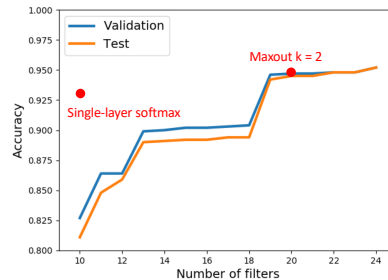


Fig. 6. Classification accuracy of pruned Max-plus model in function of the number of retained active filters.

As we can see on the diagram, the performance of the pruned Max-plus model is quite inferior to that of a single-layer softmax model when only one active filter is allowed to be selected for each class, i.e. the threshold is fixed to 1.0. However, as we relaxed the constraint on the number of total active filters by decreasing the threshold, the accuracy recovers rapidly and approaches that of the unpruned Max-plus model in a monotonic way. With exactly 24 active filters retained in the pruned model, we achieve a full-recovery of the original Max-plus model performance, which means that the other 120 linear filters do not contribute to the classification task. Moreover, we can achieve comparable performance as the 2-degree Maxout model with roughly the same amount of parameters, which again validates the effectiveness of our Max-plus models.

With the same method, we successfully performed model pruning on a much more challenging CNN model by replacing the last fully-connected layer with a Max-plus layer. In order to facilitate the training of deep Max-plus model, we resort to transfer learning by initializing the two convolutional layers with

pre-trained weights. The pruned Max-plus model achieves slightly better performance than the CNN model while reducing 94.8% of parameters of the second last fully-connected layer and eliminating the last fully-connected layer compared to the CNN model. Note that we could achieve a full-recovery of the unpruned Max-plus model performance with only 10 active filters in this case, namely one linear filter for each output unit. Table 2 summarizes the architectures of the CNN model, the unpruned Max-plus model and the pruned Max-plus model, along with their classification accuracy on the test set of CIFAR-10 dataset.

Table 2. The architectural specifications of the CNN model, unpruned and pruned Max-plus model, along with their performance on the test set of CIFAR-10 dataset.

CNN	Max-plus	Pruned Max-plus
conv(5*5)	conv(5*5)	conv(5*5)
maxpool(2*2)	maxpool(2*2)	maxpool(2*2)
conv(5*5)	conv(5*5)	conv(5*5)
maxpool(2*2)	maxpool(2*2)	maxpool(2*2)
fc(384)	fc(384)	fc(384)
fc(192)	fc(192)	fc(10)
fc(10)	maxplus(10)	maxplus(10)
83.5%	83.9%	83.9%

4.3 Comparison to Maxout Networks

It is noticeable that the pruned Max-plus network differs from Maxout networks only in their grouping strategy for maximum operations. Maxout networks impose this rigid constraint in a way that each Maxout unit has an equal number of affine components while Max-plus networks are more tolerant in this respect. Figure 7 shows a graphical illustration of this comparison.

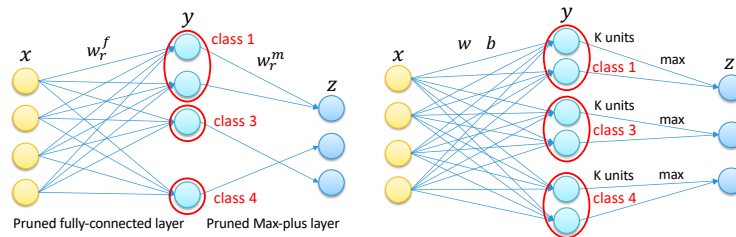


Fig. 7. Illustration of the comparison between the pruned Max-plus model (left) and Maxout model (right).

This higher flexibility hence endows Max-plus blocks with the capability of adapting the number of active filters used for each output unit accordingly. If a latent concept (say digit 1, which can be easily confused with digit 7) is considerably tougher to capture than some other concepts (say digits 3, 4 and 8, each of which has a relatively unique shape among the 10 Arabic numbers), then the Max-plus layer will select more linear filters to abstract it than for the others. For example, the partition of the 24 active filters for the 10 digit classes in Section 4.2 is $[2, 3, 3, 1, 2, 3, 3, 3, 2, 2]$, which is consistent with our point. This adaptive behavior of the filter selection property of Max-plus blocks makes the pruned Max-plus network more computationally efficient (fewer model parameters, smaller run-time memory footprint and faster inference) and is highly desirable in real-life applications.

5 Conclusions and Future Work

In this work we went a step further on a very new and promising topic, namely the reduction of deep neural networks with Max-plus blocks. Our experiments show strong evidence that model pruning via this method is compatible with high performance when a proper dropout regularization is applied during training. This was tested on data and architectures of variable complexity. Just as interesting as the obtained results are the many questions raised by these new insights. In particular, training these architectures is a challenging task which requires a better understanding on them, both theoretical and practical. We observed that training a deep model containing a Max-plus block is not straightforward, as we needed to resort to transfer learning. New optimization tricks will be needed to train deep architectures with *several* Max-plus blocks. The extension to a convolutional version of Max-plus blocks is also an open question, which we hope can be addressed based on the elements provided by this work.

Acknowledgements. This work was partially funded by a grant from Institut Mines Telecom.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., et al.: Tensorflow: a system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation. vol. 16, pp. 265–283 (2016)
2. Charisopoulos, V., Maragos, P.: Morphological perceptrons: geometry and training algorithms. In: International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing. pp. 3–15 (2017)
3. Chen, W., Wilson, J., Tyree, S., Weinberger, K., Chen, Y.: Compressing neural networks with the hashing trick. In: International Conference on Machine Learning. pp. 2285–2294 (2015)
4. Davidson, J.L., Ritter, G.X.: Theory of morphological neural networks. In: Digital Optical Computing II. vol. 1215, pp. 378–389 (1990)
5. Glorot, X., Bordes, A., Bengio, Y.: Deep sparse rectifier neural networks. In: 14th International Conference on Art. Intell. and Statistics. pp. 315–323 (2011)

6. Goodfellow, I.J., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. arXiv preprint arXiv:1302.4389 (2013)
7. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In: International Conference on Learning Representations (2015)
8. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: Advances in Neural Information Processing Systems. pp. 1135–1143 (2015)
9. Hassibi, B., Stork, D.G.: Second order derivatives for network pruning: Optimal brain surgeon. In: Advances in Neural Information Processing Systems. pp. 164–171 (1993)
10. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
11. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
12. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: Advances in Neural Information Processing Systems. pp. 598–605 (1990)
13. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. In: ICLR (2016)
14. Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: ICCV. pp. 2755–2763 (2017)
15. Luo, J., Wu, J., Lin, W.: ThiNet: A filter level pruning method for deep neural network compression. In: ICCV. pp. 5068–5076 (2017)
16. Mondal, R., Santra, S., Chanda, B.: Dense morphological network: An universal function approximator. arXiv preprint arXiv:1901.00109 (2019)
17. Pessoa, L.F., Maragos, P.: MRL-filters: A general class of nonlinear systems and their optimal design for image processing. *IEEE Transactions on Image Processing* **7**(7), 966–978 (1998)
18. Pessoa, L.F., Maragos, P.: Neural networks with hybrid morphological/rank/linear nodes: a unifying framework with applications to handwritten character recognition. *Pattern Recognition* **33**(6), 945–960 (2000)
19. Ritter, G.X., Sussner, P.: An introduction to morphological neural networks. In: 13th International Conference on Pattern Recognition. vol. 4, pp. 709–717 (1996)
20. Ritter, G., Urcid, G.: Lattice algebra approach to single-neuron computation. *IEEE Transactions on Neural Networks* **14**(2), 282–295 (2003)
21. Rosenblatt, F.: Principles of neurodynamics: perceptrons and the theory of brain mechanisms. Tech. rep., Cornell Aeronautical Lab Inc Buffalo NY (1961)
22. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* **15**(1), 1929–1958 (2014)
23. Sussner, P., Esmi, E.L.: Morphological perceptrons with competitive learning: Lattice-theoretical framework and constructive learning algorithm. *Information Sciences* **181**(10), 1929–1950 (2011)
24. Wen, W., Wu, C., et al.: Learning structured sparsity in deep neural networks. In: Advances in Neural Information Processing Systems. pp. 2074–2082 (2016)
25. Wilson, S.S.: Morphological networks. In: Visual Communications and Image Processing IV. vol. 1199, pp. 483–496 (1989)
26. Zamora, E., Sossa, H.: Dendrite morphological neurons trained by stochastic gradient descent. *Neurocomputing* **260**, 420–431 (2017)