

Improving Classification Accuracy of Feedforward Neural Networks for Spiking Neuromorphic Chips

Antonio Jimeno Yepes

IBM Research, VIC, Australia
ayepes@au1.ibm.com

Jianbin Tang

IBM Research, VIC, Australia
jbtang@au1.ibm.com

Benjamin Scott Mashford

IBM Research, VIC, Australia
bmashfor@au1.ibm.com

Abstract

Deep Neural Networks (DNN) achieve human level performance in many image analytics tasks but DNNs are mostly deployed to GPU platforms that consume a considerable amount of power. New hardware platforms using lower precision arithmetic achieve drastic reductions in power consumption. More recently, brain-inspired spiking neuromorphic chips have achieved even lower power consumption, on the order of milliwatts, while still offering real-time processing.

However, for deploying DNNs to energy efficient neuromorphic chips the incompatibility between continuous neurons and synaptic weights of traditional DNNs, discrete spiking neurons and synapses of neuromorphic chips need to be overcome. Previous work has achieved this by training a network to learn continuous probabilities, before it is deployed to a neuromorphic architecture, such as IBM TrueNorth Neurosynaptic System, by random sampling these probabilities.

The main contribution of this paper is a new learning algorithm that learns a TrueNorth configuration ready for deployment. We achieve this by training directly a binary hardware crossbar that accommodates the TrueNorth axon configuration constrains and we propose a different neuron model.

Results of our approach trained on electroencephalogram (EEG) data show a significant improvement with previous work (76% vs 86% accuracy) while maintaining state of the art performance on the MNIST handwritten data set.

1 Introduction

Neural networks achieve human level performance in many image analytics tasks but neural networks are mostly deployed to GPU platforms that consume a considerable amount of power. New hardware platforms using lower numerical precision achieve drastic reductions in power consumption. Brain-inspired spiking neuromorphic chips such as the IBM TrueNorth Neurosynaptic System [Merolla *et al.*,

2014] have achieved power consumption on the order of milliwatts while still offering high accuracy and real-time processing. Recent work [Esser *et al.*, 2015; 2016] has used these ideas to implement constrained neural networks to enable high performance image analytics on TrueNorth [Merolla *et al.*, 2014]. With recent advances in hardware platforms already available, there is an urge to provide software that fully exploits the potential of this new type of hardware.

Recent approaches in machine learning have explored training models constrained to binary weights [Courbariaux *et al.*, 2015] or low precision arithmetic [Courbariaux *et al.*, 2014; Cheng *et al.*, 2015; Hwang and Sung, 2014] and spiking deep belief neural networks [Stromatias *et al.*, 2015], achieving state of the art performance in image analytics tasks. Lower power consumption is achieved by using only additions (avoiding expensive multiplication modules), which can benefit from implementing these models in configurable architectures such as FPGA or brain-inspired systems.

Recent methods for training adaptations of fully connected style layers to the TrueNorth core-to-core connectivity [Esser *et al.*, 2015] applied to EEG data have shown that results using neuromorphic platforms exhibit performance levels that are below current work using unconstrained hardware [Nurse *et al.*, 2016]. In this work, we revise the training of deep neural networks for low energy neuromorphic architectures and demonstrate high performance levels comparable to non-constrained neural networks on MNIST and EEG data sets.

The selection of datasets in this study is guided by emerging areas of application for neuromorphic computing platforms. It has been proposed that the development of high-performance classifiers that run within a constrained power consumption limit will enable a new generation of mobile sensing applications [Lane *et al.*, 2015]. Early examples of such applications include healthcare-related wearables that involve machine vision and audio-signal processing [Poggi and Mattocchia, 2016], as well as real-time analysis of imagery for visual monitoring and surveillance [Antoniou and Angelov, 2016].

This paper is organized as follows. First feedforward neural networks and their training are introduced, followed by an overview of the TrueNorth chip architecture. Then a revised constrained model of the feedforward neural network training and deployment for TrueNorth are introduced. Finally results on EEG data and the MNIST handwritten data set are




Figure 1: Neural network

presented with discussion and future work.

2 Feedforward Neural Networks

In feedforward neural networks, inputs x_i are integrated in units I_j using a set of weights w_{ij} and a bias term b_j (which can be shared among units) as shown in equation 1. During training, the weights w_{ij} and bias b_j are learnt.

$$I_j = \sum_i x_i w_{ij} + b_j \quad (1)$$

Neurons output are a combination of the integration function I_j followed by a non-linearity, which can be a sigmoid function. The neurons can be grouped into layers and these layers can be stacked as shown in figure 1.

Prediction is done using the output layer outcome. In classification, typically each neuron is associated to one prediction category. Information is integrated in the output layer neurons, which are then be turned into probabilities using, for instance, softmax (eq. 2). p_l is the prediction probability for category c_l from all the possible K categories.

$$p_l = \frac{e^{c_l}}{\sum_{k=1}^K e^{c_k}} \quad (2)$$

During training, weight and bias values are typically learnt using backpropagation [Duda *et al.*, 2012], which minimizes a loss function that represents classification error of the network using gradient descent. An example of such a loss function is the logarithmic loss, which is used in this work (eq. 3), even though other functions could be explored.

$$E = - \sum_k [y_k \log(p_k) + (1 - y_k) \log(1 - p_k)] \quad (3)$$

The next sections describe the neuromorphic hardware and an explanation about how the neural networks can be constrained, learnt and deployed.

2.1 Deployment Hardware

IBM TrueNorth [Merolla *et al.*, 2014], a low power and highly parallelized brain-inspired chip is used here as an example deployment system, though our method could be generalized to other neuromorphic hardware [Benjamin *et al.*, 2014; Painkras *et al.*, 2013; Pfeil *et al.*, 2012]. In its current implementation it is composed of 4,096 neurosynaptic cores.

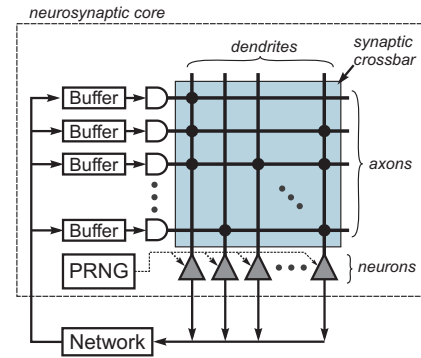


Figure 2: TrueNorth neurosynaptic-core

Figure 2 shows the layout of one of these cores. Each core has highly configurable 256 input axons and 256 output neurons. Input spikes enter the core through axons and the information is passed via the configurable binary synaptic crossbar to the leaky integrate-and-fire neurons, which act as computation units. The neuron assigns weights to the input spikes, integrates and updates its membrane potential, compares it to a threshold and may spike and reset, based on the selected threshold, reset, and stochastic modes [Cassidy *et al.*, 2013].

Programming for TrueNorth means writing core configurations and connections, which in a way could be seen as building a neural network. A programming language has been developed to abstract the chip configuration [Amir *et al.*, 2013]. In our case, the learned neural networks need to be translated into a valid TrueNorth configuration.

In previous work [Esser *et al.*, 2015], a constrained network is trained, which learns the crossbar connection probability and biases. To deploy the model into TrueNorth these probabilities are sampled into discrete synaptic crossbar connections.

3 Constrained Training Model

Neuromorphic chips such as TrueNorth provide a reduction of power consumption, which is achieved by a constrained hardware architecture. There are several constraints compared to more traditional hardware platforms that are explored in this section to define the training of neuromorphic chips. These constraints are: use of spiking neurons, weights constrained to a limited set of integer values and reduced connectivity. In the following sections, we describe these constraints in more detail and provide a way to do learning based on backpropagation. We also describe how to use output layer predictions integrated into backpropagation.

3.1 Spiking Neurons

Spiking neurons have two possible outputs (0, no output, or 1, spike) as shown in the revised neuron output shown in eq. 4.

$$n_j = \begin{cases} 1, & \text{if } I_j > 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

Spiking neurons behavior could be seen as a Heaviside step function, which does not have a gradient for backpropagation since it is not differentiable. This step function can have many

different smooth approximations using, for instance, the logistic function that is used to represent the neuron output \hat{n}_j as shown in eq. 5. The parameter k is the steepness of the function, which in our work has been set empirically to $1/2$. Other values of k could be explored or, may be, approximate its value to a logistic distribution using large data sets.

$$\hat{n}_j \approx \frac{1}{1 + e^{-2kI_j}} \quad (5)$$

The following gradient formulation can then be used in backpropagation.

$$\frac{\partial \hat{n}_j}{\partial I_j} = \frac{1}{1 + e^{-2kI_j}} * \left(1 - \frac{1}{1 + e^{-2kI_j}}\right) \quad (6)$$

In this approach, each neuron is stateless, so its status is reset after each time cycle, or tick. This is implemented in TrueNorth by setting parameters as specified and each neurons leak equal to the rounded bias term, with its threshold to 0, its membrane potential floor to 0 [Esser *et al.*, 2015].

3.2 Constrained Weights

In TrueNorth, effective weights w_{ij} are divided in two components: the axon weight s_{ij} and the crossbar connection c_{ij} , so $w_{ij} = c_{ij}s_{ij}$. Taking this into account, the integration function is defined as:

$$\hat{I}_j = \sum_i x_i c_{ij} s_{ij} + b_j \quad (7)$$

Instead of learning the weights w_{ij} , the axon weights are proportionally distributed as explained below the binary crossbar connection c_{ij} is learnt during training.

The hardware accepts two possible values for the crossbar connections (0,1). The effective weights are dependent on the configuration of the axon types and crossbar connections. In this setup, the total number of effective weights that can be used is a maximum of five, i.e. four axon type weights and no connection (which means zero weight). The distribution of weights in s_{ij} depends on the axon type configuration in a TrueNorth core (we follow [Esser *et al.*, 2015]).

The backpropagation gradients are specified as below and a similar formulation is used for bias values, which are round to the closer integer value during deployment. In the next section, we explain how the binary crossbar connections are learnt.

$$\frac{\partial \hat{n}_j}{\partial c_{ij}} = \frac{\partial \hat{n}_j}{\partial \hat{I}_j} \frac{\partial \hat{I}_j}{\partial c_{ij}}; \frac{\partial \hat{I}_j}{\partial c_{ij}} = x_i s_{ij} \quad (8)$$

3.3 Binary Crossbar Connection Learning

In the previous section, we have shown that binary crossbar connections need to be learnt. We propose adapting previous work [Esser *et al.*, 2016; Courbariaux *et al.*, 2015] to learn binary crossbar connections that are used to obtain the effective network weights as shown in algorithm 1.

This is in contrast to previous work in which connection probabilities are learnt [Esser *et al.*, 2015] that need to be sampled prior to hardware deployment. As well, in contrast to previous work [Esser *et al.*, 2016; Courbariaux *et al.*, 2015],

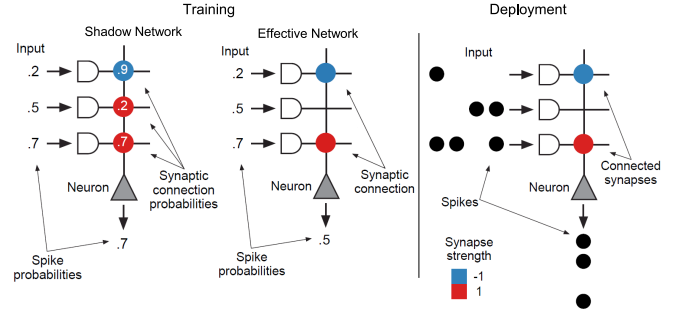


Figure 3: Crossbar connections in training and deployment

the proposed approach allows learning a larger number of weights that are defined by the hardware configuration.

During training, in the forward propagation step the high precision c^{sh} matrix from the shadow network is used to calculate the c binary connection values (eq. 9), which are used to calculate the effective network weights using ($w_{ij} = c_{ij}s_{ij}$). Effective network weights are used to calculate the performance of the trained model (c.f. figure 3).

$$c_{ij} = \begin{cases} 1, & \text{if } c_{ij}^{sh} > 0.5 \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

During backpropagation, gradients are estimated using c connection values that are used to update the c^{sh} high precision connection probabilities.

Algorithm 1 Forward and backward propagation

```

//Forward propagation
for l = 1 : layers do
    for c = 1 : layer(l).cores do
        Calculate core binary crossbar connections  $c_{ij}$  (c.f.
        eq. 9) and effective weights  $w_{ij}$ 
        Calculate core neuron outputs  $\hat{I}_j$  using  $c_{ij}$ 
    end for
end for
Estimate log loss in the last layer
//Backpropagation
for l = 1 : layers do
    for c = 1 : layer(l).cores do
        Calculate gradients using  $c_{ij}$ 
        Update  $c_{ij}^{sh}$  using estimated gradients
    end for
end for
    
```

3.4 Neural Network Connections

The described method allows defining fully connected layer networks, but TrueNorth cores have 256 inputs (axons) and 256 outputs (neurons), which constrains the size of fully connected layers. When designing a network, the first layer, which processes the input data, is defined by blocks that tile to cover the input data. Block size and stride are defined by the number of pixels in the image classification case. Since each of these blocks maps to a TrueNorth core, the number of

input axons required to map the input data per block is calculated as $(blockSize^2 * numberChannels)$ and the maximum value is 256, which is the number of axons per core. In the upper layers, block size and stride are defined by the number of cores (instead of image size as in the input layer). Output neurons of the previous layer are mapped to the input axons of the cores in the upper layer.

3.5 Output Layer

The last layer of the of the network is used for prediction. Spiking neurons in the last layer can be assigned randomly to one category [Esser *et al.*, 2015]. The category with the highest number of neurons firing will be the predicted one.

These counts need to be used in backpropagation. In our work, we have used the number of spiking neurons per category cs_k divided by the number of neurons per category cn_k and used *softmax* as shown in equation 10 to estimate the per category probability that is then used with the logarithmic loss (eq. 3).

$$\hat{p}_l = \frac{e^{\frac{cs_l}{cn_l}}}{\sum_{k=1}^K e^{\frac{cs_k}{cn_k}}} \quad (10)$$

We have described neural networks, the constraints added by neuromorphic chips and how to adapt the training of neural networks to these constraints. In the following sections, we present the results obtained using the MNIST data set and EEG data.

4 Results

We have used two data sets in our experiments. An EEG data set, described in more detail below, and the MNIST handwritten data set.

EEG data¹ was recorded from a participants wearing electrodes placed on the scalp while they performed a self-paced hand squeeze task. Full experimental details are described elsewhere [Nurse *et al.*, 2015]. Briefly, the participant wore surface electromyography (EMG) electrodes attached to the forearms that were used to detect the onset of each hand squeeze. Data was sampled at 1 kHz in four minute sessions, with a total recording time of approximately 30 minutes. Detection of hand squeeze onset was achieved by extracting the root mean squared (RMS) power from the EMG signal, and the corresponding EEG data was labeled with either left-hand squeeze or right-hand squeeze (2 classes). Each EEG instance has been translated to 46x46 images with one 8-bit channel. In this study, EEG signals from participant D are used, which allows comparing the results to previous work. This data set contains 480/468 examples per class in the training set and 66/95 examples per class in the testing set. The data set consisted of a randomly allocated unseen test set (20% of the data) and training set (80% of the data), allowing a direct comparison with previous results [Nurse *et al.*, 2015].

MNIST data set [LeCun *et al.*, 1998] contains handwritten digits from 0 to 9 with 60,000 examples for training and 10,000 samples for testing. Images are 28x28 pixels, with

¹<https://github.com/EwanNurse/A-Generalizable-BCI-using-Machine-Learning-for-Feature-Discovery>

Name	Layer	Definition
Small network	1	Block size 16 Stride 12
	2	Block size 1 Stride 1
	3	Block size 2 Stride 1
Large network	1	Block size 16 Stride 4
	2	Block size 2 Stride 1
	3	Block size 2 Stride 1
	4	Block size 2 Stride 1

Table 1: Networks used in the experiments. MNIST small has 3 layers and requires 9 TrueNorth cores. MNIST large has 4 layers and requires 30 TrueNorth cores.

one 8-bit gray scale channel. Gray scale values have been normalized dividing them by 255.

During training EEG and MNIST data are represented using floating point values. During deployment, the data values a single input axon over n ticks, which means that floating point values need to be represented as a set of spikes. Rate code [Jimeno Yepes *et al.*, 2016] has been used to encode the data. In rate code, the number of spikes in the time window represents the encoded value, k. Spikes are evenly distributed within the time window.

Experiments were run using a MATLAB implementation of the proposed method, this includes training and evaluation of the trained models in a simulated MATLAB environment. We have used two configurations for the axon type weights in the neuron configuration $s=[-1,1]$, used in previous work, and $s=[-2,-1,1,2]$.

Table 1 lists the networks used in the experiments. Network configurations are derived from previous work [Esser *et al.*, 2015]. There is a small network with 3 layers that requires 9 TrueNorth cores and a larger network with 4 layers that requires 30 TrueNorth cores.

Performance of the deployed trained models has been measured using classification accuracy, which is calculated by dividing the number of correctly predicted instances by all instances in the test set. McNemar’s test has been used to look if the outcome of the systems is statistically significant [Haralick, 1993]. During test, data is converted into spikes using a rate code scheme. Since rate code is used and there is no need to sample from the trained network crossbar probabilities, results are stable (e.g. repetitions of the experiments provide the same result) compared to previous methods.

We have found that the initial optimal learning rate lr was provided by equation 11, which takes into account the number of neurons assigned to each category and the batch size.

$$lr = (neurons_per_category * batch_size)^{-1} \quad (11)$$

4.1 EEG Results

For EEG experiments, the initial learning rate lr was multiplied by 0.1. We have updated the learning rate every 100 iterations using a rate of 0.1. Batch size is 25 and we have run a total of 500 epochs.

Ranges of values vary a lot for EEG data across features. We have normalized each individual feature x_i but instead of using values between -1 and 1 as it is typically done with

T	s=[-2,-1,1,2]		s=[-1,1]	
	Original	Normalized	Original	Normalized
1	61.43±3.72	76.58±2.17	62.36±5.27	79.19±2.45
2	64.66±4.82	78.01±1.58	66.96±4.22	81.80±2.43
4	68.94±2.40	80.93±1.78	70.06±3.98	84.22±1.50
8	71.12±2.22	81.30±2.46	72.36±2.92	84.66±1.58
16	73.73±3.21	82.55±1.48	75.96±2.17	84.07±1.40
32	76.27±3.14	83.17±0.99	80.00±2.34	85.90±0.66
64	76.52±2.23	83.60±0.52	80.62±1.89	86.09±0.73

Table 2: EEG accuracy results using the large network and two axon weights s configuration and with (Normalized) and without normalization (Original). T stands for the number of ticks used to encode test instances with rate code.

unconstrained neural networks [Ioffe and Szegedy, 2015], the values have been normalized to be between 0 and 1 as shown in the formula below. A small ϵ value has been inserted to avoid zero values.

$$\hat{x}_i = \frac{x_i - E[x_i^{train}]}{\sqrt{\text{var}(x_i^{train}) + \epsilon}} + 0.5 \quad (12)$$

Results are shown in Table 2 training the network with two different axon type weights and the combination of both trained models. Previously reported results using TrueNorth showed a 76% accuracy for participant D on this 2-class task [Nurse *et al.*, 2016]. With the proposed method we obtain over 86%, which is the same accuracy as previously reported for the same data set using an unconstrained neural network [Nurse *et al.*, 2015]. Better results are obtained when a more restricted set of axon type weights ($s = [-1, 1]$) is used. Normalization of the set at the feature level improves the performance of the trained models, which indicates that feature preprocessing can benefit predictive performance. Both results are statistically significant ($p < 0.005$). Using more than 32 ticks does not show a statistically significant improvement.

4.2 MNIST Results

Training has been carried out using batches of 100 images for each iteration and with and without data augmentation. To perform data augmentation, random changes to the images have been done in each epoch. We used a maximum rotation of 7.5 degrees, a maximum shift of 2.5 and a maximum rescale of 7.5%. Augmentation changes are randomly performed. For experiments without augmentation, the initial learning rate lr is multiplied by 0.1 and 500 epochs were used. For augmentation experiments, the initial learning rate lr is multiplied by 0.1, which is further multiplied by 0.1 after 2,000 epochs. 3,000 epochs were used for training. More iterations were required for the training set using augmentation to converge compared to the number of iterations used without augmentation.

Tables 3 and 4 show the results for several configurations. The large network shows better performance compared to the smaller one. The configuration with axon types ($s = [-1, 1]$) provide better performance. Augmentation increases the performance, which is similar to previous published work [Esser *et al.*, 2015]. Improvements in performance by these three

T	s=[-2,-1,1,2]		s=[-1,1]	
	No aug	Aug	No aug	Aug
1	95.05±0.14	95.99±0.13	93.73±0.21	95.20±0.16
2	96.27±0.09	97.46±0.13	96.07±0.20	97.17±0.11
4	97.00±0.09	98.20±0.08	97.28±0.11	98.22±0.11
8	97.30±0.09	98.44±0.03	97.75±0.06	98.58±0.09
16	97.44±0.06	98.56±0.07	97.95±0.08	98.76±0.05
32	97.56±0.07	98.64±0.05	98.03±0.05	98.84±0.02
64	97.65±0.04	98.66±0.03	98.08±0.04	98.86±0.03

Table 3: MNIST small network results using data augmentation (Aug) and not using it (No aug). T stands for the number of ticks used to encode test instances with rate code.

T	s=[-2,-1,1,2]		s=[-1,1]	
	No aug	Aug	No aug	Aug
1	95.45±0.13	96.97±0.13	95.36±0.17	97.26±0.13
2	96.50±0.11	98.07±0.11	96.45±0.12	98.32±0.10
4	97.16±0.06	98.60±0.07	97.38±0.10	98.88±0.07
8	97.53±0.09	98.78±0.04	97.70±0.04	99.16±0.05
16	97.68±0.05	98.91±0.04	97.88±0.05	99.29±0.04
32	97.80±0.05	98.97±0.04	97.94±0.04	99.29±0.03
64	97.82±0.05	98.99±0.03	97.99±0.03	99.32±0.02

Table 4: MNIST large network results using data augmentation (Aug) and not using it (No aug). T stands for the number of ticks used to encode test instances with rate code.

configurations are statistically significant ($p < 0.005$). As with the EEG results, using more than 32 ticks does not show a statistically significant improvement.

5 Discussion

Results on EEG data show that the performance of the trained network increased from 76% [Nurse *et al.*, 2016] to 86%, which is close to state-of-the-art performance obtained using unconstrained fully connected layers [Nurse *et al.*, 2015]. Results on MNIST show that the performance of the proposed method is similar compared to previously reported work depending on the configuration.

The use of several ticks for encoding input data shows that a plateau is reached after 8 ticks for MNIST, additional ticks to encode data does not significantly increase performance.

We have examined the effective weights and bias values. Figures 4 and 5 show the values for MNIST and figure 6 shows the values for the EEG data set.

In all combinations, approximately half of the weights have a value of zero for all the layers, which implies that half of the crossbar connections are disconnected. The other weights are evenly distributed. Furthermore, bias values tend to be around zero for all layers in many cases. On the other hand, when using augmentation in MNIST or in the EEG experiments, bias values tend to be distributed mostly on the negative values for the first layers. Negative values might indicate that there is a tendency to prevent neurons from spiking.

The values for weights and biases imply that in a deployment scenario, hardware configuration could be further constrained to a limited set of bias values and further research is required to understand implications in hardware design to improve crossbar usage.

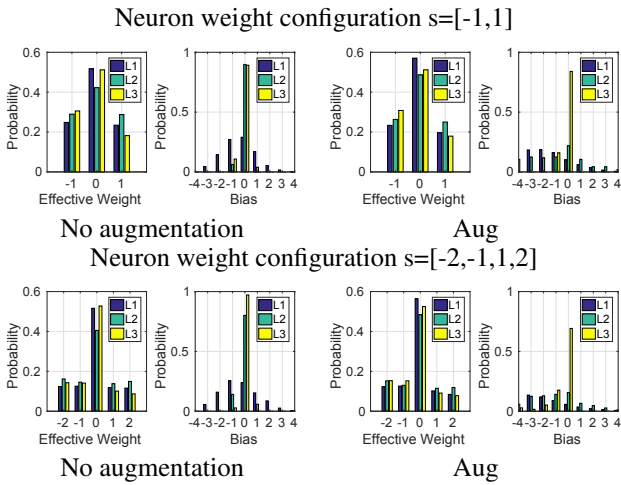


Figure 4: Effective weights and bias values for the small network (3 layers, 9 cores) for MNIST data set. Two weight configurations have been used $s=[-1,1]$ and $s=[-2,-1,1,2]$. The networks have been trained without augmentation (No augmentation) and with augmentation (Aug).

6 Conclusions and Future Work

We have shown that it is possible to train constrained fully connected layers on neuromorphic chips for two sets of problems (image classification with MNIST and EEG data analytics), which provides better or similar performance compared to previous work depending on the data set.

Results on EEG data sets which were previously analysed through the original method yield over 86% classification accuracy using the proposed approach, which is comparable to unconstrained neural networks. The trained models use a small portion of the TrueNorth chip (30 cores vs. 4096 available in the current version of the chip), thus requiring a much less than 70mW to work, which makes these models suitable for portable autonomous devices with large autonomy.

The outcome of this research provides insights into the possibility of learning network parameters that translate into deployment networks for brain-inspired chips. We will explore feasibility of expanding the introduced new methodology to train other kinds of existing deep learning algorithms including convolutional neural networks and recurrent neural networks. Furthermore, analysis of the learnt parameters provide insights that might complement hardware design, thus providing a more efficient deployment of the trained models.

Acknowledgments

The authors would like to thank the SyNAPSE team at the IBM Almaden Research Center for their support.

References

[Amir *et al.*, 2013] Arnon Amir, Pallab Datta, William P Risk, Andrew S Cassidy, Jeffrey Kusnitz, Steven K Esser, Alexander Andreopoulos, Theodore M Wong, Myron Flickner, Rodrigo Alvarez-Icaza, et al. Cognitive computing programming paradigm: a corelet language for

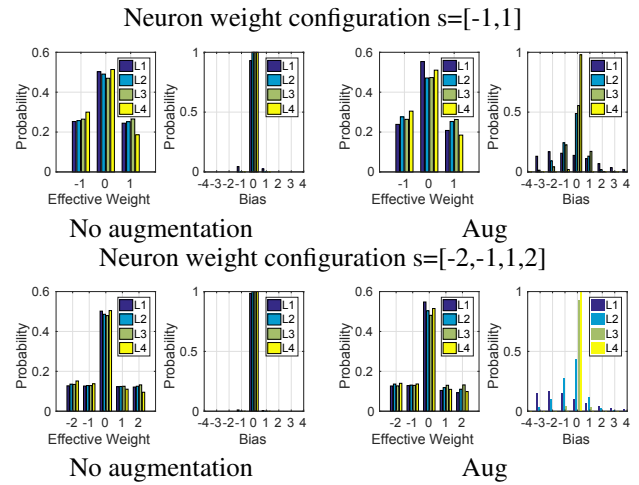


Figure 5: Effective weights and bias values for the large network (4 layers, 30 cores) for MNIST data set. Two weight configurations have been used $s=[-1,1]$ and $s=[-2,-1,1,2]$. The networks have been trained without augmentation (No augmentation) and with augmentation (Aug).

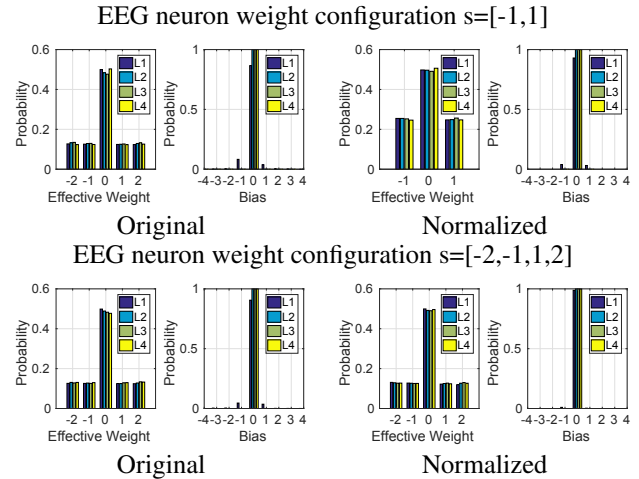


Figure 6: Effective weights and bias values for the large network (4 layers, 30 cores) for EEG data set. Two weight configurations have been used $s=[-1,1]$ and $s=[-2,-1,1,2]$. The networks have been trained using the original set or the normalized one.

composing networks of neurosynaptic cores. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–10. IEEE, 2013.

[Antoniou and Angelov, 2016] Antreas Antoniou and Plamen Angelov. A general purpose intelligent surveillance system for mobile devices using deep learning. In *Neural Networks (IJCNN), 2016 International Joint Conference on*, pages 2879–2886. IEEE, 2016.

[Benjamin *et al.*, 2014] Ben Varkey Benjamin, Peiran Gao, Emmett McQuinn, Swadesh Choudhary, Anand R Chandrasekaran, Jean-Marie Bussat, Rodrigo Alvarez-Icaza, John V Arthur, Paul A Merolla, and Kwabena Boahen. Neurogrid: A mixed-analog-digital multichip system for

- large-scale neural simulations. *Proceedings of the IEEE*, 102(5):699–716, 2014.
- [Cassidy *et al.*, 2013] Andrew Cassidy, Paul Merolla, John Arthur, Steve Esser, Bryan Jackson, Rodrigo Alvarez-Icaza, Pallab Datta, Jun Sawada, Theodore Wong, Vitaly Feldman, Arnon Amir, Daniel Ben-Dayana Rubin, et al. Cognitive computing building block: A versatile and efficient digital neuron model for neurosynaptic cores. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–10. IEEE, 2013.
- [Cheng *et al.*, 2015] Zhiyong Cheng, Daniel Soudry, Zexi Mao, and Zhenzhong Lan. Training binary multilayer neural networks for image classification using expectation backpropagation. *arXiv preprint arXiv:1503.03562*, 2015.
- [Courbariaux *et al.*, 2014] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Low precision arithmetic for deep learning. *arXiv preprint arXiv:1412.7024*, 2014.
- [Courbariaux *et al.*, 2015] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in Neural Information Processing Systems*, pages 3105–3113, 2015.
- [Duda *et al.*, 2012] Richard O Duda, Peter E Hart, and David G Stork. *Pattern classification*. John Wiley & Sons, 2012.
- [Esser *et al.*, 2015] Steve K Esser, Rathinakumar Appuswamy, Paul Merolla, John V Arthur, and Dharmendra S Modha. Backpropagation for energy-efficient neuromorphic computing. In *Advances in Neural Information Processing Systems*, pages 1117–1125, 2015.
- [Esser *et al.*, 2016] Steven K Esser, Paul A Merolla, John V Arthur, Andrew S Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J Berg, Jeffrey L McKinstry, Timothy Melano, Davis R Barch, et al. Convolutional networks for fast, energy-efficient neuromorphic computing. *arXiv preprint arXiv:1603.08270*, 2016.
- [Haralick, 1993] Robert M Haralick. Performance characterization in computer vision. In *International Conference on Computer Analysis of Images and Patterns*, pages 1–9. Springer, 1993.
- [Hwang and Sung, 2014] Kyuyeon Hwang and Wonyong Sung. Fixed-point feedforward deep neural network design using weights+ 1, 0, and- 1. In *Signal Processing Systems (SiPS), 2014 IEEE Workshop on*, pages 1–6. IEEE, 2014.
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [Jimeno Yepes *et al.*, 2016] Antonio Jimeno Yepes, Jianbin Tang, Shreya Saxena, Tobias Brosch, and Arnon Amir. Weighted population code for low power neuromorphic image classification. In *Neural Networks (IJCNN), 2016 International Joint Conference on*, pages 4294–4301. IEEE, 2016.
- [Lane *et al.*, 2015] Nicholas D Lane, Sourav Bhattacharya, Petko Georgiev, Claudio Forlivesi, and Fahim Kawsar. An early resource characterization of deep learning on wearables, smartphones and internet-of-things devices. In *Proceedings of the 2015 International Workshop on Internet of Things towards Applications*, pages 7–12. ACM, 2015.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [Merolla *et al.*, 2014] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Filipp Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197):668–673, 2014.
- [Nurse *et al.*, 2015] Ewan S Nurse, Philippa J Karoly, David B Grayden, and Dean R Freestone. A generalizable brain-computer interface (bci) using machine learning for feature discovery. *PLoS one*, 10(6):e0131328, 2015.
- [Nurse *et al.*, 2016] Ewan Nurse, Benjamin S Mashford, Antonio Jimeno Yepes, Isabell Kiral-Kornek, Stefan Harrer, and Dean R Freestone. Decoding eeg and lfp signals using deep learning: heading truenorth. In *Proceedings of the ACM International Conference on Computing Frontiers*, pages 259–266. ACM, 2016.
- [Painkras *et al.*, 2013] Eustace Painkras, Luis A Plana, Jim Garside, Steve Temple, Francesco Galluppi, Cameron Patterson, David R Lester, Andrew D Brown, and Steve B Furber. Spinnaker: A 1-w 18-core system-on-chip for massively-parallel neural network simulation. *IEEE Journal of Solid-State Circuits*, 48(8):1943–1953, 2013.
- [Pfeil *et al.*, 2012] Thomas Pfeil, Andreas Grübl, Sebastian Jeltsch, Eric Müller, Paul Müller, Mihai A Petrovici, Michael Schmuker, Daniel Brüderle, Johannes Schemmel, and Karlheinz Meier. Six networks on a universal neuromorphic computing substrate. *arXiv preprint arXiv:1210.7083*, 2012.
- [Poggi and Mattocchia, 2016] Matteo Poggi and Stefano Mattocchia. A wearable mobility aid for the visually impaired based on embedded 3d vision and deep learning. In *Computers and Communication (ISCC), 2016 IEEE Symposium on*, pages 208–213. IEEE, 2016.
- [Stromatias *et al.*, 2015] E Stromatias, D Neil, F Galluppi, M Pfeiffer, SC Liu, and S Furber. Scalable energy-efficient, low-latency implementations of spiking deep belief networks on spinnaker. In *2015 International Joint Conference on Neural Networks (IJCNN)(Killarney)*, 2015.