



Biologically Plausible Learning of Text Representation with Spiking Neural Networks

Marcin Białas¹(✉) , Marcin Michał Mironczuk¹ , and Jacek Mańdziuk² 

¹ National Information Processing Institute, al. Niepodległości 188 b,
00-608 Warsaw, Poland

{marcin.bialas,marcin.mironczuk}@opi.org.pl

² Faculty of Mathematics and Information Sciences, Warsaw University
of Technology, Koszykowa 75, 00-662 Warsaw, Poland
mandziuk@mini.pw.edu.pl

Abstract. This study proposes a novel biologically plausible mechanism for generating low-dimensional spike-based text representation. First, we demonstrate how to transform documents into series of spikes (*spike trains*) which are subsequently used as input in the training process of a spiking neural network (SNN). The network is composed of biologically plausible elements, and trained according to the unsupervised Hebbian learning rule, Spike-Timing-Dependent Plasticity (STDP). After training, the SNN can be used to generate low-dimensional spike-based text representation suitable for text/document classification. Empirical results demonstrate that the generated text representation may be effectively used in text classification leading to an accuracy of 80.19% on the *bydate* version of the *20 newsgroups* data set, which is a leading result amongst approaches that rely on low-dimensional text representations.

Keywords: Spiking neural network · STDP · Hebbian learning · Text processing · Text representation · Spike-based representation · Representation learning · Feature learning · Text classification · 20 newsgroups bydate

1 Introduction

Spiking neural networks (SNNs) are an example of biologically plausible artificial neural networks (ANNs). SNNs, like their biological counterparts, process sequences of discrete events occurring in time, known as spikes. Traditionally, spiking neurons, due to their biological validity, have been studied mostly by theoretical neuroscientists, and have become a standard tool for modeling brain processes on a micro scale. However, recent years have shown that spiking computation can also successfully address common machine learning challenges [35]. Another interesting aspect of SNNs is the adaptation of such algorithms to neuromorphic hardware which is a brain-inspired alternative to the traditional von

© The Author(s) 2020

T. Bäck et al. (Eds.): PPSN 2020, LNCS 12269, pp. 433–447, 2020.

https://doi.org/10.1007/978-3-030-58112-1_30

Neumann machine. Thanks to mimicking processes observed in brain synaptic connections, neuromorphic hardware is a highly fault-tolerant and energy-efficient substitute for classical computation [27].

Recently we have witnessed significant growth in the volume of research into SNNs. Researchers have successfully adapted SNNs for the processing of images [35], audio signals [9, 39, 40], and time series [18, 31]. However, to the best of the authors knowledge, there is only one work related to text processing with SNNs [37]. This state of affairs is caused by the fact that text, due to its structure and high dimensionality, presents a significant challenge to tackle by the SNN approach. The motivation of this study is to broaden the current knowledge of the application of SNNs to text processing. More specifically, we have developed and evaluated a novel biologically inspired method for generation of spike-based text representation that may be used in text/document classification task [25].

1.1 Objectives and Summary of Approach

This paper proposes an *Spike Encoder for Text* (SET) which generates spike-based text representation suitable for classification task. Text data is highly dimensional (the most common text representation is in the form of a vector with many features) which, due to the *curse of dimensionality* [2, 17, 20, 26], usually leads to overfitted classification models with poor generalisation [4, 13, 30, 32, 38].

Processing highly dimensional data is also computationally expensive. Therefore, researchers have sought text representations which may overcome this drawback [3]. One of possible approaches is based on transformation of high dimensional feature space to low-dimensional representation [5, 6, 36].

In the above context we propose the following two-phase approach to SNN based text classification. Firstly, the text is transformed into *spike trains*. Secondly, spike trains representation is used as the input in the SNN training process performed according to biologically plausible unsupervised learning rule, and generating the spike-based text representation. This representation has significantly lower dimensionality than the spike trains representation and can be used effectively in subsequent SNN text classification. The proposed solution has been empirically evaluated on the publicly available version, *bydate* [21] of the real data set known as *20 newsgroups*, which contains 18 846 text documents from twenty different newsgroups of *Usenet*, a worldwide distributed discussion system.

Both the input and output of the SNN rely on spike representations, though of very different forms. For the sake of clarity, throughout the paper the former representation (SNN input) will be referred to as *spike trains*, and the latter one (SNN output) as *spike-based*, or *spiking encoding*, or *low-dimensional*.

1.2 Contribution

The main contribution of this work can be summarized as follows:

- To propose an original approach to document processing using SNNs and its subsequent classification based on generated spike-based text representation;

- To experimentally evaluate the influence of various parameters on the quality of generated representation, which leads to better understanding of the strengths and limitations of SNN-based text classification approaches;
- To propose an SNN architecture which may potentially contribute to development of other SNN based approaches. We believe that the solution presented may serve as a building block for larger SNN architectures, in particular deep spiking neural networks (DSNNs) [35];

1.3 Related Work

As mentioned above, we are aware of only one paper related to text processing in the context of SNNs context [37] which, nevertheless, differs significantly from our approach. The authors of [37] focus on transforming word embeddings [23, 28] into spike trains, whilst our focus is not only on representation of text in the form of spike trains, but also on training the SNN encoder which generates low-dimensional text representation. In other words, our goal is to generate a low-dimensional text representation with the use of SNN base, whereas in [37] the transformation of an existing text embedding into spike trains is proposed.

This remainder of the paper is structured as follows. Section 2 presents an overview of the proposed method; Sect. 3 describes the evaluation process of the method and experimental results; and Sect. 4 presents the conclusions.

2 Proposed Spiking Neural Method

The proposed method transforms input text to spike code and uses it as training input for the SNN to achieve a meaningful spike-based text representation. The method is schematically presented in Fig. 1. In phase *I*, text is transformed into a vector representation and afterwards each vector is encoded as spike trains. Once the text is encoded in the form of neural activity, it can be used as input to the core element of our method - a *spiking encoder*. The encoder is a two-layered SNN with adaptable synapses. During the learning phase (*II*), the spike trains are propagated through the encoder in a feed-forward manner and synaptic weights are modified simultaneously according to unsupervised learning rule. After the learning process, the output layer of the spiking encoder provides spike-based representation of the text presented to the system.

In the remainder of this section all elements of the system described above are discussed in more detail.

2.1 Input Transformation

2.1.1 Text Vectorization

During a *text to spike transformation* phase like the one illustrated in Fig. 1 text is preprocessed for further spiking computation. Text input data (*data corpus*) is organized as a set D of documents $d_i, i = 1, \dots, K$. In the first step a dictionary T containing all unique words $t_j, j = 1, \dots, |T|$ from the corpus data is built.

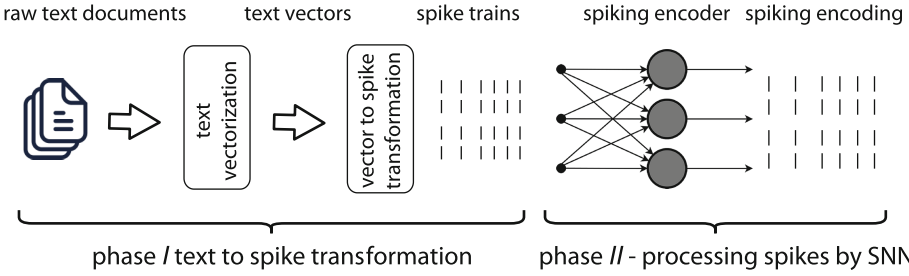


Fig. 1. A schema of the proposed method for generating spike-based low-dimensional text representation.

Next, each document d_i is transformed into an M -dimensional ($M = |T|$) vector W_i , the elements of which, $W_i[j] := w_{ij}, j = 1, \dots, M$ represent the relevance of words t_j to document d_i . In effect, the corpus data is represented by a real-valued matrix $W_{K \times M}$ also called *document-term matrix*.

The typical weighting functions are *term-frequency* (TF), *inverse document frequency* (IDF), or their combination TF-IDF [12, 22]. In TF the weight w_{ij} is equal to the number of times the j -th word appears in d_i with respect to the length of $|d_i|$ (the number of all non-unique words in d_i). IDF takes into account the whole corpus D and sets w_{ij} as the logarithm of a ratio between $|D|$ and the number of documents containing word t_j . Consequently, IDF mitigates the impact of words that occur very frequently in a given corpus and are presumably less informative from the point of view of document classification than the words occurring in a small fraction of the documents. TF-IDF sets w_{ij} as a product of TF and IDF weights. In this paper we use TF-IDF weighting which is the most popular approach in text processing domain.

2.1.2 Vector to Spike Transformation

In order to transform a vector representation to spike trains one, presentation time t_p which establishes for how long each document is presented to the network, and the time gap between two consecutive presentations Δt_p , must be defined. A time gap period, without any input stimuli is necessary to eliminate interference between documents and allow dynamic parameters of the system to decay and “be ready” for the next input.

Technically, for a given document d_i , represented as M dimensional vector of weights w_{ij} , for each weight w_{ij} in every millisecond of document presentation a spike is generated with probability proportional to w_{ij} . Thanks to this procedure, we ultimately derive a spiking representation of the text.

In our experiments each document is presented for $t_p = 600[ms]$ and $\Delta t_p = 300[ms]$, and proportionality coefficient α is set to 1.5.

For a better clarification, let’s consider a simple example and assume that for a word *baseball* the corresponding weight w_{ij} in some document d_i is equal to 0.1. Then for each millisecond of a presentation time a probability of emitting

a spike $P(\text{spike}|\text{baseball})$ equals $\alpha \cdot 0.1 = 0.15$. Hence, 90 spikes during 600[ms] presentation time are expected, on average, to be generated.

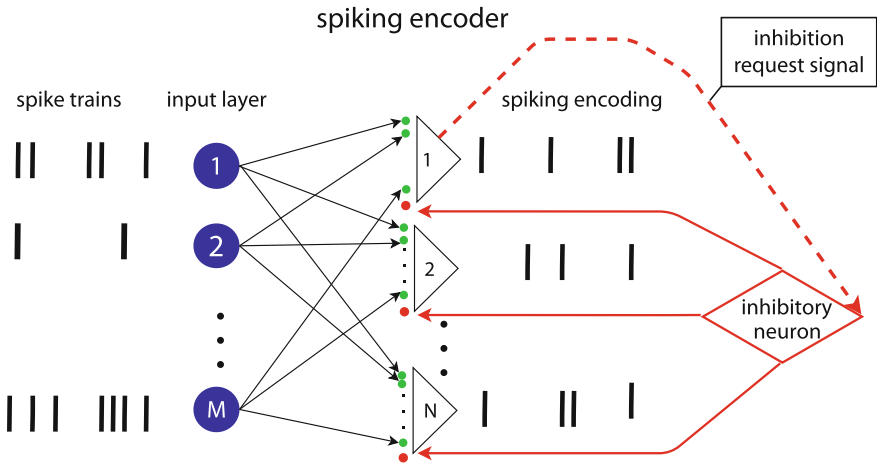


Fig. 2. Spiking encoder architecture.

2.2 Spiking Encoder Architecture and Dynamics

A spiking encoder is the key element of the proposed method. The encoder, presented in Fig. 2, is a two layered SNN equipped with an additional inhibitory neuron. The first layer contains M neurons (denoted by blue circles) and each of them represents one word t_j from the dictionary T . Neuron dynamics is defined by the *spike trains* generated based on weights w_{ij} corresponding to documents $d_i, i = 1, \dots, K$. Higher numbers of spikes are emitted by neurons representing words which are statistically more relevant for a particular document, according to the chosen TF-IDF measure. The *spike trains* for each neuron are presented in Fig. 2 as a row of short vertical lines.

In the brain spikes are transmitted between neurons via synaptic connections. A neuron which generates a spike is called a *presynaptic neuron*, whilst a target neuron (spike receiver) is a *postsynaptic neuron*. In the proposed SNN architecture (cf. Fig. 2) two different types of synaptic connections are utilised: *excitatory* ones and *inhibitory* ones. Spikes transmitted through excitatory connections (denoted by green circles in Fig. 2) leads to firing of postsynaptic neuron, while impulses traveling through inhibitory ones (red circles in Fig. 2) hinder *postsynaptic neuron* activity. Each time an encoder neuron fires its weights are modified according to the proposed learning rule. The neuron simultaneously sends an *inhibition request signal* to the *inhibitory neuron* and activates it. Then the inhibitory neuron suppresses the activity of all encoder output layer neurons using *recursive inhibitory connection* (red circles). The proposed architecture

satisfies the competitive learning paradigm [19] with a winner-takes-all (WTA) strategy.

In this work we consider a biologically plausible neuron model known as leaky integrate and fire (LIF) [11]. The dynamics of such a neuron is described in terms of changes of its membrane potential (MP). If the neuron is not receiving any spikes its potential is close to the value of $u_{rest} = -65[mV]$ known as resting membrane potential. When the neuron receives spikes transmitted through excitatory synapses, the MP moves towards excitatory equilibrium potential, $u_{exc} = 0[mV]$. When many signals are simultaneously transmitted through excitatory synapses the MP rises and at some point can reach a threshold value of $u_{th} = -52[mV]$, in which case the neuron fires. After firing, the neuron resets its MP to u_{rest} and becomes inactive for $t_{ref} = 3[ms]$ (the refractory period). In the opposite scenario, when the neuron receives spikes through the inhibitory synapse, its MP moves towards inhibitory equilibrium potential $u_{inh} = -90[mV]$, i.e. further away from its threshold value, which decreases the chance of firing. The dynamics of the membrane potential u in the LIF model is described by the following equation:

$$\tau \frac{du}{dt} = (u_{rest} - u) + g_e(u_{exc} - u) + g_i(u_{inh} - u) \quad (1)$$

where g_e and g_i denote excitatory and inhibitory conductance, resp. and $\tau = 100[ms]$ is membrane time constant. The values of g_e and g_i depend on presynaptic activity. Each time a signal is transmitted through the synapse the conductance is incremented by the value of weight corresponding to that synapse, and decays with time afterwards according to Eq. (4)

$$\tau_e \frac{dg_e}{dt} = -g_e, \quad \tau_i \frac{dg_i}{dt} = -g_i, \quad (2)$$

where $\tau_e = 2[ms]$, $\tau_i = 2[ms]$ are decay time constants. In summary, if there is no presynaptic activity the MP converges to u_{rest} . Otherwise, its value changes according to the signals transmitted through the neuron synapses.

2.3 Hebbian Synaptic Plasticity

We utilise a modified version of the Spike-Timing-Dependent Plasticity (STDP) learning process [33]. STDP is a biologically plausible unsupervised learning protocol belonging to the family of Hebbian learning (HL) methods [14]. In short, the STDP process results in an increase of the synaptic weight if the postsynaptic spike is observed soon after presynaptic one ('pre-before-post'), and in a decrease of the synaptic weight in the opposite scenario ('post-before-pre'). The above learning scheme increases the relevance of those synaptic connections which contribute to the activation of the postsynaptic neuron, and decreases the importance of the ones which do not. We modify STDP in a manner similar to [8, 29], i.e. by skipping the weight modification in the post-before-pre scenario

and introducing an additional scaling mechanism. The plasticity of the excitatory synapse s_{ij} connecting a presynaptic neuron i from the input layer with postsynaptic neuron j from the encoder layer can be expressed as follows:

$$\Delta s_{ij} = \eta(A(t) - (R(t) + 0.1)s_{ij}) \quad (3)$$

where

$$A(t) = -\tau_A \frac{dA(t)}{dt}, \quad R(t) = -\tau_R \frac{dR(t)}{dt} \quad (4)$$

and $\eta = 0.01$ is a small learning constant. In Eqs. (3)–(4) $A(t)$ represents a presynaptic trace and $R(t)$ is a scaling factor which depends on the history of postsynaptic neuron activity. Every time the presynaptic neuron i fires $A(t)$ is set to 1 and exponentially decays in time ($\tau_A = 5[ms]$). If the postsynaptic neuron fires just after the presynaptic one (‘pre-before-post’) $A(t)$ is close to 1 and the weight increase is high. The other component of Eq. (3) $(R(t) + 0.1)s_{ij}$ is a form of synaptic scaling [24]. Every time the postsynaptic neuron fires $R(t)$ is incremented by 1 and afterwards decays with time ($\tau_R = 70[ms]$). The role of the small constant factor 0.1, is to maintain scaling even when activity is relatively small. The overall purpose of synaptic scaling is to decrease the weights of the synapses which are not involved in firing the postsynaptic neuron. Another benefit of synaptic scaling is to restrain weights from uncontrolled growth which can be observed in HL [1].

2.4 Learning Procedure

For a given *data corpus* (set of documents) D the training procedure is performed as follows. Firstly, we divide D into s subsets $u_i, i = 1, \dots, s$ in the manner described in Sect. 3.1. Secondly, each subset u_i is transformed to spike trains and used as input for a separate SNN encoder $H_i, i = 1, \dots, s$ composed of N neurons. Please note that each encoder is trained with the use of one subset only. Such a training setup allows the processing of the data in parallel manner. Another advantage is that this limits the number of excitatory connections per neuron, which reduces computational complexity (the number of differential equations that need to be evaluated for each spike) as during training encoder H_i is exposed only to the respective subset, T_i of the training set dictionary T and the number of its excitatory connections is limited to $|T_i| < |T|$. Spike trains are presented to the network four times (in four training epochs).

Once the learning process is completed, the *connection pruning* procedure is applied. Please observe that HL combined with competitive learning should lead to highly specialised neurons which are activated only for some subset of the inputs. The specialisation of a given neuron depends on the set of its connection weights. If the probability of firing should be high for some particular subset of the inputs, the weights representing words from those inputs must be high. The other weights should be relatively low due to the synaptic scaling mechanism. Based on this assumption, after training, for each output layer neuron we prune θ per cent of its incoming connections with the lowest weights. θ is a hyper parameter of the method empirically evaluated in the experimental section.

3 Empirical Evaluation and Results Comparison

This section presents experimental evaluation of the method proposed. In Subsect. 3.1 the technical issues related to the setup of experiment and implementation of the training and evaluation procedures are discussed. The final two subsections focus respectively on the experimental results and compare them with the literature.

3.1 Experiment Setup

3.1.1 Data Set and Implementation Details

The *bydate* version¹ of *20 newsgroups* is a well known benchmark set in the text classification domain. The set contains newsgroups post related to different categories (topics) gathered from *Usenet*, in which each category corresponds to one newsgroup. Categories are organised into a hierarchical structure with the main categories being *computers*, *recreation and entertainment*, *science*, *religion*, *politics*, and *forsale*. The corpus consists of 18 846 documents nearly equally distributed among twenty categories and explicitly divided into two subsets: the training one (60%) and the test one (40%).

The dynamics of the spiking neurons (including the plasticity mechanism) was implemented using the *BRIAN 2* simulator [34]. *Scikit-learn Python library*² was used for processing the text and creating the TF-IDF matrix.

3.1.2 Training

As mentioned in Sect. 2.4 the training set was divided into $s = 11$ subsets u_i each of which, except for u_{11} , contained 1 500 documents. Division was performed randomly. Firstly the entire training set was shuffled, and then consecutively assigned to the subsets according to the resulting order with a 500 document redundancy (overlap) between the neighbouring subsets, as described in Table 1.

The overlap between subsequent subsets resulted from preliminary experiments which suggested that such an approach improves classification accuracy. While we found the concept of partial data overlap to be reasonably efficient, it by no means should be regarded as an optimal choice. The optimal division of data into training subsets remains an open question and a subject of our future research.

3.1.3 Evaluation Procedure

The outputs of all SNNs $H_i, i = 1, \dots, s$, i.e. spike-based encodings represented as sums of spikes per document were joined to form a single matrix (a final low-dimensional text representation) which was evaluated in the context of a classification task. The joined matrix of spike rates was used as an input to the Logistic Regression (LR) [15, 17] classifier with *accuracy* as the performance measure.

¹ <http://qwone.com/~jason/20Newsgroups/20news-bydate.tar.gz>.


² <https://scikit-learn.org/>.

Table 1. Division of the training set into subsets u_1 – u_{11} .

Subset	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}	u_{11}
First index	0	1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
Last index	1500	2500	3500	4500	5500	6500	7500	8500	9500	10500	11314
Size	1500	1500	1500	1500	1500	1500	1500	1500	1500	1500	1314

3.2 Experimental Results

In the first experiment we looked more closely at the weights of neurons after training and the relationship between the inhibition mechanism and the quality/efficacy of resulting text representation. We trained eleven SNN encoders with 50 neurons each according to the procedure presented above. After training, 5 neurons from the first encoder (H_1) was randomly sampled and their weights were used for further analysis. Figure 3 illustrates the highest 200 weights sorted in descending order.

**Fig. 3.** Weights extracted from the encoder’s neurons.

The weights of each neuron are presented with a different colour. The plots show that every neuron has a group of dominant connections represented by the weights with the highest values - the first several dozen connections. It means that each neuron will be activated more easily by the inputs that contain words corresponding to these weights. For example neuron 4 will potentially produce more spikes for documents related to religion because its 10 highest weights corresponds to words ‘*jesus*’, ‘*god*’, ‘*paul*’, ‘*faith*’, ‘*law*’, ‘*christians*’, ‘*christ*’, ‘*sabbath*’, ‘*sin*’, ‘*jewish*’. A different behaviour is expected from neuron 2 whose 10 highest weights corresponds to words ‘*drive*’, ‘*scsi*’, ‘*disk*’, ‘*hard*’, ‘*controller*’, ‘*ide*’, ‘*drives*’, ‘*help*’, ‘*mac*’, ‘*edu*’. This one will be more likely activated for computer related documents. On the other hand, not all neurons can be classified so easily. For instance 10 highest weights of neuron 5 are linked to words ‘*cs*’,

'serial', 'ac', 'edu', 'key', 'bit', 'university', 'windows', 'caronni', 'uk', hence a designation of this neuron is less obvious. We have repeated the above sampling and weigh inspection procedure several times and the observations are qualitatively the same. For the sake of space savings we do not report them in detail.

Hence, a question arises as to how well documents can be encoded with the use of neurons trained in the manner described above? Intuitively, in practice the quality of encoding may be related to the level of competition amongst neurons in the evaluation phase. If the inhibition value is kept high enough to satisfy WTA strategy then only a few neurons will be activated and the others will be immediately suppressed. This scenario will lead to highly sparse representations of the input documents, with just a few or (in extreme cases) only one neuron dominating the rest. Since differences between documents belonging to different classes may be subtle, such a sparse representation may not be the optimal setup. In order to check the influence of the inhibition level on the resulting spike-based representation we tested the performance of the trained SNNs H_1-H_{11} for various inhibition levels by adjusting the value of the neurons' inhibitory synapses. The results are illustrated in Fig. 4 (top).

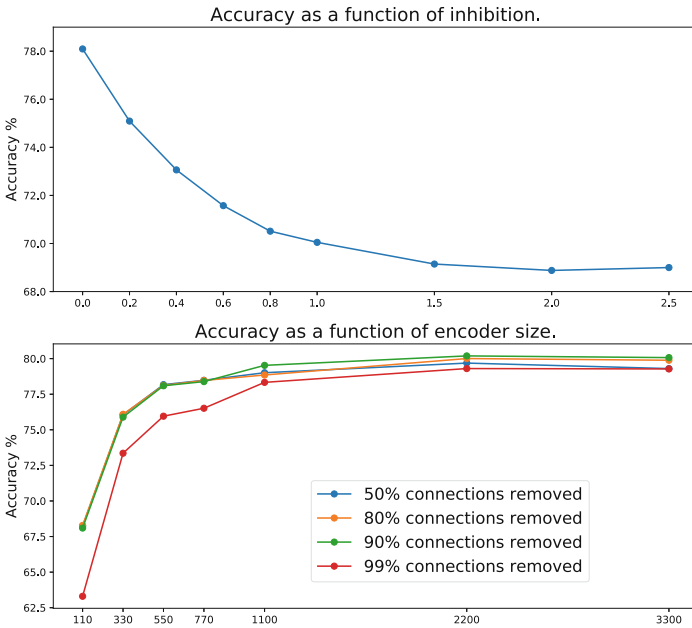


Fig. 4. Accuracy for various inhibition levels (top) and encoder sizes (bottom).

Clearly the accuracy strongly depends on the inhibition level. The best outcomes ($\approx 78\%$) are accomplished with inhibition set to 0 and rapidly decrease along with the inhibition raise. For the inhibition values higher than 1.5 the accuracy plot enters a plateau at the level of approximately 68%. The results

show that the most effective representation of documents is generated with the absence of inhibition during the evaluation phase, i.e. when all neurons have the same chance of being activated and contribute to the document representation.

The second series of experiments aimed at exploring the relationship between the efficacy of document representation and the size of the encoders. Furthermore, the sensitivity of the trained encoders to *connection pruning*, with respect to their efficiency, was verified. The results of both experiments are shown in the bottom plot of Fig. 4. Seven encoders of various sizes (between 110 and 3 300 neurons) were trained, and once the training was completed the before the *connection pruning* procedure took place.

In the plot, four colored curves illustrate particular pruning scenarios and their impact on classification accuracy for various encoder sizes. 99%, 90%, 80%, and 50% of the weakest weights were respectively removed in the four discussed cases. Overall, for smaller SNN encoders (between 110 and 1 100 neurons) the accuracy rises rapidly along with the encoder size increase. For larger SSNs, changes in the accuracy are slower and for all four curves stay within the range [77.5%, 80.19%].

In terms of the *connection pruning* degree the biggest changes in accuracy (between 63% and 79%) are observed when 99% of connections have been deleted (the red curve). In particular, the results of the encoders with fewer than 1 100 neurons demonstrate that this range of pruning heavily affects classification accuracy. In larger networks additional neurons compensate the features removed by the *connection pruning* mechanism and the results are getting closer to other pruning setups.

Interestingly, for the networks smaller than 770 neurons, the differences in accuracy between 50%, 80%, and 90% pruning setups are negligible, which suggests that relatively high redundancy of connections still exist in the networks pruned in the range of 50% to 80%. Apparently, retaining as few as 10% of the weights does not impact the quality of representation and does not cause deterioration of results. This result well correlates with the outcomes of the weight analysis reported above and confirms that a meaningful subset of connections is sufficient for proper encoding the input. The best overall classification result (80.19%) was achieved by the SNN encoder with 2 200 neurons and level of pruning set to 90% (the green curve). It proves that SEM can effectively reduce dimensionality of the text input from initial $\approx 130\,000$ (the size of *20 newsgroups* training vocabulary) to the size of 550 – 2 200, and maintain classification accuracy above 77.5%.

3.3 Results Analysis. Comparison with the Literature

Since to our knowledge this paper presents the first attempt of using SNN architecture to text classification, in order to make some comparisons we selected results reported for other neural networks trained with similar input (*document-term matrix*) and yielding low-dimensional text representation as the output. The results are presented in Table 2. SET achieved 80.19% accuracy and outperformed the remaining shallow approaches. While this result looks promising we

believe that there is still room for improvement with further tuning of the method (in particular a division of samples into training subsets), as well as extension of the SNN encoder by adding more layers. Another interesting direction would be to learn semantic relevance between different words and documents [10, 41].

Table 2. Accuracy [%] comparison for several low-dimensional text representation methods on *bydate* version of *20 newsgroups* data set.

Method	Accuracy
SET (this paper)	80.19
K-competitive Autoencoder for TExt (KATE) [7]	76.14
Class Preserving Restricted Boltzmann Machine (CPr-RBM) [16]	75.39
Variational Autoencoder [7]	74.30

4 Conclusions

This work offers a novel approach to text representation relying on Spiking Neural Networks. Using the proposed low-dimensional text representation the LR classifier accomplished 80.19% accuracy on a standard benchmark set (*20 newsgroups bydate*) which is a leading result among shallow approaches relying on low-dimensional representations.

We have also examined the influence of the inhibition mechanism and synaptic connections sparsity on the quality of the representation showing that (i) it is recommended that inhibition be disabled during the SNN evaluation phase, and (ii) pruning out as many as 90% of connections with lowest weights did not affect the representation quality while heavily reducing the SNN computational complexity, i.e. the number of differential equations describing the network.

There are a few lines of potential improvement that we plan to explore in the further work. Most notably, we aim to expand the SNN encoder towards Deep SNN architecture by adding more layers of spiking neurons which should possibly allow to learn more detailed features of the input data.

References

1. Abbott, L.F., Nelson, S.B.: Synaptic plasticity: taming the beast. *Nat. Neurosci.* **3**(Suppl 1), 1178–1183 (2000)
2. Aggarwal, C.C.: *Data Mining*. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-14142-8>
3. Aggarwal, C.C.: *Machine Learning for Text*, 1st edn. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73531-3_9
4. Asif, M., Ishtiaq, A., Ahmad, H., Aljuaid, H., Shah, J.: Sentiment analysis of extremism in social media from textual information. *Telematics Inform.* **48**, 101345 (2020). <https://doi.org/10.1016/j.tele.2020.101345>

5. Ayesha, S., Hanif, M.K., Talib, R.: Overview and comparative study of dimensionality reduction techniques for high dimensional data. *Inf. Fusion* **59**, 44–58 (2020)
6. Bengio, Y., Courville, A., Vincent, P.: Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1798–1828 (2013). <https://doi.org/10.1109/tpami.2013.50>
7. Chen, Y., Zaki, M.J.: KATE: k-competitive autoencoder for text. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, 13–17 August 2017, pp. 85–94. ACM (2017). <https://doi.org/10.1145/3097983.3098017>
8. Diehl, P., Cook, M.: Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* **9**, 99 (2015). <https://doi.org/10.3389/fncom.2015.00099>
9. Dominguez-Morales, J.P., et al.: Deep spiking neural network model for time-variant signals classification: a real-time speech recognition approach. In: 2018 International Joint Conference on Neural Networks (IJCNN). IEEE, July 2018. <https://doi.org/10.1109/ijcnn.2018.8489381>
10. Gao, Y., Wang, W., Qian, L., Huang, H., Li, Y.: Extending embedding representation by incorporating latent relations. *IEEE Access* **6**, 52682–52690 (2018). <https://doi.org/10.1109/ACCESS.2018.2866531>
11. Gerstner, W., Kistler, W.M.: Spiking Neuron Models. Cambridge University Press, Cambridge (2002). <https://doi.org/10.1017/cbo9780511815706>
12. Haddoud, M., Mokhtari, A., Lecroq, T., Abdeddaïm, S.: Combining supervised term-weighting metrics for SVM text classification with extended term representation. *Knowl. Inf. Syst.* **49**(3), 909–931 (2016). <https://doi.org/10.1007/s10115-016-0924-1>
13. Hartmann, J., Huppertz, J., Schamp, C., Heitmann, M.: Comparing automated text classification methods. *Int. J. Res. Market.* **36**(1), 20–38 (2019). <https://doi.org/10.1016/j.ijresmar.2018.09.009>
14. Hebb, D.O.: The organization of behavior: A neuropsychological theory. New York (1949)
15. Hosmer, D.W., Lemeshow, S.: Applied Logistic Regression. Second edn. Wiley, Hoboken (2000). <https://doi.org/10.1002/0471722146>
16. Hu, J., Zhang, J., Ji, N., Zhang, C.: A new regularized restricted boltzmann machine based on class preserving. *Knowl.-Based Syst.* **123**, 1–12 (2017). <https://doi.org/10.1016/j.knosys.2017.02.012>
17. James, G., Witten, D., Hastie, T., Tibshirani, R.: An Introduction to Statistical Learning: with Applications in R. Springer, Heidelberg (2013). <https://doi.org/10.1007/978-1-4614-7138-7>
18. Kasabov, N., Capecchi, E.: Spiking neural network methodology for modelling, classification and understanding of EEG spatio-temporal data measuring cognitive processes. *Inf. Sci.* **294**, 565–575 (2015). <https://doi.org/10.1016/j.ins.2014.06.028>
19. Kaski, S., Kohonen, T.: Winner-take-all networks for physiological models of competitive learning. *Neural Netw.* **7**, 973–984 (1994). [https://doi.org/10.1016/S0893-6080\(05\)80154-6](https://doi.org/10.1016/S0893-6080(05)80154-6)
20. Keogh, E., Mueen, A.: Curse of Dimensionality, pp. 314–315. Springer, Boston (2017). https://doi.org/10.1007/978-1-4899-7687-1_192
21. Lang, K.: NewsWeeder: learning to filter netnews. In: Proceedings of the Twelfth International Conference on Machine Learning, pp. 331–339 (1995)
22. Manning, C.D., Raghavan, P., Schütze, H.: Introduction to Information Retrieval. Cambridge University Press, New York (2008)

23. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Burges, C.J.C., Bottou, L., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5–8, 2013, Lake Tahoe, Nevada, United States*, pp. 3111–3119 (2013)
24. Miller, K.D., MacKay, D.J.C.: The role of constraints in Hebbian learning. *Neural Comput.* **6**(1), 100–126 (1994). <https://doi.org/10.1162/neco.1994.6.1.100>
25. Mladenović, D., Brank, J., Grobelnik, M.: *Document Classification*, pp. 372–377. Springer, Boston (2017). https://doi.org/10.1007/978-1-4899-7687-1_75
26. Murphy, K.P.: *Machine Learning - A Probabilistic Perspective*. Adaptive Computation and mACHine Learning Series. MIT Press, Cambridge (2012)
27. Nawrocki, R.A., Voyles, R.M., Shaheen, S.E.: A mini review of neuromorphic architectures and implementations. *IEEE Trans. Electron Devices* **63**(10), 3819–3829 (2016). <https://doi.org/10.1109/TED.2016.2598413>
28. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. In: Moschitti, A., Pang, B., Daelemans, W. (eds.) *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25–29, 2014, Doha, Qatar, A Meeting of SIGDAT, a Special Interest Group of the ACL*, pp. 1532–1543. ACL (2014). <https://doi.org/10.3115/v1/d14-1162>
29. Querlioz, D., Bichler, O., Dollfus, P., Gamrat, C.: Immunity to device variations in a spiking neural network with memristive nanodevices. *IEEE Trans. Nanotechnol.* **12**(3), 288–295 (2013). <https://doi.org/10.1109/TNANO.2013.2250995>
30. Raza, M., Hussain, F.K., Hussain, O.K., Zhao, M., ur Rehman, Z.: A comparative analysis of machine learning models for quality pillar assessment of SaaS services by multi-class text classification of users’ reviews. *Future Gen. Comput. Syst.* **101**, 341–371 (2019)
31. Reid, D., Hussain, A.J., Tawfik, H.: Financial time series prediction using spiking neural networks. *PLoS ONE* **9**(8), e103656 (2014). <https://doi.org/10.1371/journal.pone.0103656>
32. Silva, R.M., Almeida, T.A., Yamakami, A.: MDLText: an efficient and lightweight text classifier. *Knowl.-Based Syst.* **118**, 152–164 (2017). <https://doi.org/10.1016/j.knosys.2016.11.018>
33. Song, S., Miller, K., Abbott, L.: Competitive Hebbian learning through spike timing-dependent plasticity. *Nat. Neurosci.* **3**, 919–926 (2000). <https://doi.org/10.1038/78829>
34. Stimberg, M., Goodman, D., Benichoux, V., Brette, R.: Equation-oriented specification of neural models for simulations. *Front. Neuroinform.* **8**, 6 (2014). <https://doi.org/10.3389/fninf.2014.00006>
35. Tavanaei, A., Ghodrati, M., Kheradpisheh, S.R., Masquelier, T., Maida, A.: Deep learning in spiking neural networks. *Neural Netw.* **111**, 47–63 (2019). <https://doi.org/10.1016/j.neunet.2018.12.002>
36. Vlachos, M.: *Dimensionality Reduction*, pp. 354–361. Springer, Boston (2017). https://doi.org/10.1007/978-1-4899-7687-1_71
37. Wang, Y., Zeng, Y., Tang, J., Xu, B.: Biological neuron coding inspired binary word embeddings. *Cogn. Comput.* **11**(5), 676–684 (2019). <https://doi.org/10.1007/s12559-019-09643-1>
38. Webb, G.I.: *Overfitting*, pp. 947–948. Springer, Boston (2017). https://doi.org/10.1007/978-1-4899-7687-1_960

39. Wu, J., Chua, Y., Zhang, M., Li, H., Tan, K.C.: A spiking neural network framework for robust sound classification. *Front. Neurosci.* **12** (2018). <https://doi.org/10.3389/fnins.2018.00836>
40. Wysoski, S.G., Benuskova, L., Kasabov, N.: Evolving spiking neural networks for audiovisual information processing. *Neural Netw.* **23**(7), 819–835 (2010). <https://doi.org/10.1016/j.neunet.2010.04.009>
41. Zheng, S., Bao, H., Xu, J., Hao, Y., Qi, Z., Hao, H.: A bidirectional hierarchical skip-gram model for text topic embedding. In: 2016 International Joint Conference on Neural Networks, IJCNN 2016, Vancouver, BC, Canada, 24–29 July 2016, pp. 855–862. IEEE (2016). <https://doi.org/10.1109/IJCNN.2016.7727289>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

