



Published in final edited form as:

Proc IEEE Int Conf Acoust Speech Signal Process. 2016 March ; 2016: 754–758. doi:10.1109/ICASSP.

2016.7471776

EPILEPTIFORM SPIKE DETECTION VIA CONVOLUTIONAL NEURAL NETWORKS

Alexander Rosenberg Johansen^{a,b}, Jing Jin^b, Tomasz Maszczyk^b, Justin Dauwels^b, Sydney S. Cash^c, and M. Brandon Westover^c

^aTechnical University of Denmark, DTU Compute, Lyngby, Denmark

^bNanyang Technological University, School of Electrical and Electronic Engineering, Singapore

^cMassachusetts General Hospital Neurology Department, and Harvard Medical School, USA

Abstract

The EEG of epileptic patients often contains sharp waveforms called “spikes”, occurring between seizures. Detecting such spikes is crucial for diagnosing epilepsy. In this paper, we develop a convolutional neural network (CNN) for detecting spikes in EEG of epileptic patients in an automated fashion. The CNN has a convolutional architecture with filters of various sizes applied to the input layer, leaky ReLUs as activation functions, and a sigmoid output layer. Balanced mini-batches were applied to handle the imbalance in the data set. Leave-one-patient-out cross-validation was carried out to test the CNN and benchmark models on EEG data of five epilepsy patients. We achieved 0.947 AUC for the CNN, while the best performing benchmark model, Support Vector Machines with Gaussian kernel, achieved an AUC of 0.912.

Index Terms

Epilepsy; Spike detection; EEG; Deep learning; Convolutional neural network

1. INTRODUCTION

Epilepsy refers to a group of chronic brain disorders characterized by recurrent seizures, affecting approximately 65 million people worldwide [1]. Electroencephalography (EEG) is the primary diagnostic test for epilepsy, which provides a continuous measure of cortical function with excellent temporal resolution. Significant efforts are spent on interpreting EEG data for clinical purposes. In current clinical practice, visual inspection and manual annotation are still the gold standard for interpreting EEG, which is tedious and ultimately subjective. In addition, experienced electroencephalographers are in short supply [2]. As a result, a great need exists for automated systems for EEG interpretation.

The finding of primary importance for the diagnosis for epilepsy is the presence of interictal discharges, also known as “spikes” and “sharp waves”, hereafter referred to collectively simply as “spike(s)”. Automated spike detection would enable wider availability of EEG diagnostics and more rapid referral to qualified physicians who can provide further medical investigation and interventions. However, spikes are difficult to detect in a consistent manner due to the large variability of spike waveforms between patients among other factors [3].

Great attempts have been made to detect spikes by general classifications such as mimetic, linear predictive, and template based methods [4]. Many recent spike detection algorithms combine multiple methodologies, such as local context [5, 6, 7], morphology [8, 6, 7, 9, 10, 11], field of spike [8, 7, 12, 11, 13], artifact rejection [8, 6, 7], and temporal and spatial contexts [8, 13, 14]. Unfortunately, none of them are universally accepted or tested on a significantly large dataset of patients and spikes. To date, no algorithmic approach has overcome these challenges to yield expert-level detection of spikes [3].

In this study, we analyze the scalp EEG recordings of five patients diagnosed with epilepsy. Suspected interictal epileptiform spikes were cross-annotated by two neurologists. We applied convolutional neural networks (CNNs) to learn the discriminative features of spikes. CNNs are statistical models incorporating prior knowledge about the discriminative features of spikes. CNNs are commonly applied for finding local pixel dependencies [15]. Furthermore, CNNs have been proven successful in surpassing human accuracy in image recognition tasks [16], time-series tasks for text analysis [17] and biological sequences [18]. In addition, CNNs possess much fewer connections compared to a fully-connected neural network, due to the sparsity and the parameter sharing across the filters. Nevertheless, the computational complexity remains a key challenge to implementing CNNs. To address this issue, we utilize Graphical Processing Units (GPUs) and high performance libraries for modeling CNNs.

We benchmarked the CNN approach with several standard classification methods. Leave-one-patient-out cross-validation was conducted to generate the receiver operating characteristic (ROC) curve for each model, with the average area-under-the-curve (AUC) as the benchmark criterion. We achieved 0.947 AUC for the CNN, while the best performing benchmark model, Support Vector Machines with Gaussian kernel, achieved an AUC of 0.912.

This paper is organized as follows. In Section 2, we provide information on the EEG data considered in this study. We also elaborate on the design of the CNN model. In Section 3, we provide results for the CNN approach and four standard classifiers. In Section 4, we offer concluding remarks and ideas for future research.

2. METHODS

2.1. Epileptiform EEG

We analyzed the scalp EEG recordings of five patients with known epilepsy. Each EEG data lasts 30 minutes recorded from 19 standard 10-20 scalp electrodes. The sampling rate is 128Hz. A high-pass filter at 1Hz was applied to remove the baseline drifts. A notch filter centered at 60Hz was applied to remove the power line interference. The common average referencing montage was applied to remove the common EKG artifacts.

As shown in Fig. 1a, interictal epileptiform spikes are morphologically defined events with an outstanding sharp peak distinguishable from the background fluctuations (see Fig.1b). Spikes were cross-annotated independently by two neurologists, and extracted with a fixed duration of 0.5 s. In order to reduce the computational load, we randomly sampled 1,500

spikes and 150,000 background waveforms from each subject. Leave-one-patient-out cross-validation was conducted to evaluate the CNN model. Moreover, 10% of the training data was further extracted under stratified sampling for validation, in order to tune the CNN model.

2.2. Convolutional Neural Networks

The problem of spike detection is typically ill-posed, i.e., there are many models which fit the training patients or spikes well but do not generalize well. In other words, there is often not sufficient training data to allow accurate estimation of class probabilities throughout the input space. Convolutional neural networks (CNNs) [19] are suitable for such scenario, since they incorporate constraints and can achieve some degree of shift and deformation invariance. The architecture of a CNN (see Fig. 2) typically contains multiple layers: convolutional layers, dense layers, and an output layer. We applied the backpropagation algorithm to train the CNN. In this way, we can extract both low-level and high-level features from the input layer.

2.2.1. Activation function—In computational networks, the activation function defines the output of a layer given an input or a set of inputs. Nonlinear activation functions are commonly used in neural networks to improve the learning capacity and system robustness for nontrivial problems. Let us denote the output of layer ℓ as h_ℓ (with h_0 the input layer), $\Theta_{\ell+1}$ as the weight matrix of layer $\ell+1$, and $z_{\ell+1}$ as the linear combination of the weighted input to each neuron. The output of layer $\ell+1$ is computed as follows:

$$\begin{aligned} z_{\ell+1} &= h_\ell \Theta_{\ell+1}, \\ h_{\ell+1} &= a(z_{\ell+1}). \end{aligned} \quad (1)$$

The activation function is denoted by $a(z)$. The most commonly used nonlinear activation functions in neural networks are the Logistic Sigmoid $a(z) = 1/(1 + e^{-z})$ and the Hyperbolic Tangent $a(z) = (e^z - e^{-z})/(e^z + e^{-z})$. Non-saturating activation functions such as the Rectifier Linear Unit (ReLU) with $a(z) = \max(0, z)$ have become more popular for their non-vanishing gradients and computational efficiency [15]. However, the ReLU also suffers from “dead” gradients, i.e. from large gradients that may deactivate the neurons, and further disable the network. In order to avoid such problems, we applied the Leaky ReLU [16] as our activation function in CNN, which is defined as:

$$a(z) = \max(\alpha z, z), \quad (2)$$

with $\alpha = 0.1$. “Leaky” refers to the additional slope when $a(z) < 0$.

2.2.2. The convolutional architecture—In 1D Convolutional architectures, the processing steps of (1) are modified as follows:

$$\begin{aligned} z_{\ell+1}^{id} &= x_{\ell}^i \Theta_{\ell+1}^d, \\ h_{\ell+1}^{id} &= a(z_{\ell+1}^{id}), \end{aligned} \quad (3)$$

where the input $x_{\ell}^i \in \mathbb{R}^{kc \times 1}$ represents a co-located vector of length k in c channels and i refers to the specific co-located vector. The weight matrix $\Theta_{\ell}^d \in \mathbb{R}^{1 \times kc}$ corresponds to a spatial weight filter with kc connections, in the d^{th} output channel, between the input layer and each neuron in the output layer. As a result, convolutional layers are configurable to their filter size k and number of filters d [15], leading to many possible configurations of the CNN architecture.

Multiple convolutional layers can help to reduce the parameter space, and to model non-linear mappings [20]. Moreover, combining different filter sizes can improve the overall performance [18]. Therefore, in this paper, we did not only test a single convolutional layer, but also stacked different convolutional layers on top of each other, and tested convolutional layers with different filter sizes. Based on our results, the implementation of multiple convolutional layers with different filter sizes provided the best validation performance, and thus was used to build the final CNN model.

2.2.3. The final CNN model—As depicted in Fig. 2, the CNN model contains five layers. The first three layers are convolutional layers with different filter sizes. The fourth layer is a fully-connected layer, which feeds into a fifth binary logistic layer (for normalization purpose) to make final decisions on whether an input waveform is a spike or not. During the training process, the CNN network is optimized by minimizing the binomial cross-entropy on the basis of the probability output from the neural network as:

$$L(x, y) = -f(x) \log(y) - (1 - y) \log(1 - f(x)), \quad (4)$$

where $f(x)$ is the prediction of the neural network given x .

After being activated by inputs, the convolutional layers are merged and passes the inputs to the neurons of the next fully-connected layer. The output of the fully-connected layer is then passed to the logistic output function. The leaky ReLU nonlinear activation function is applied after each convolutional layer as well as the fully-connected layer. The three convolutional layers in the CNN model contains one-dimensional filter of size 3, 5 and 7 respectively, and a stride of 1 with 8 filters each. The fully-connected layer has 150 hidden neurons. In order to avoid overfitting, we are using dropouts [21] with $p = 0.5$ on the fully-connected layer.

The data is extremely skewed (or “imbalanced”), since there are vastly more background waveforms than spikes. Therefore, the CNN may mostly model the background waveforms, instead of the spikes. This problem is addressed by training the CNN by means of balanced mini-batches.

Our CNN model is relatively small in size (220k parameters) compared to conventional CNN models for image processing [15]. Due to the small dataset used in this study, we limited the number of parameters in the CNN model to avoid overfitting while keeping the model simple. For future work, data augmentation [15], such as shifting the spikes by a few data points, may alleviate the problem of overfitting by generating synthetic data.

2.2.4. Details of learning—Our CNN model was trained with stochastic gradient descent using a batch-size of 4096 (2048 spikes and 2048 background waveforms). The update rule for the weight matrix Θ involves the Nesterov momentum [22] as follows:

$$\begin{aligned} v_{t+1} &= \mu v_t - \varepsilon \Delta f(\Theta_t + \mu v_t), \\ \Theta_{t+1} &= \Theta_t + v_{t+1}, \end{aligned} \quad (5)$$

with $\varepsilon > 0$ the learning rate, $\mu \in [0, 1]$ the momentum co-efficient, and $\Delta f(\Theta_t + \mu v_t)$ the gradient. In this study, the learning rate ε was set by grid searching from 10^{-3} to 10^{-5} , where $2 \cdot 10^{-4}$ gave the best convergence in the training data. The momentum applied was set to $\mu = 0.9$ and not optimized further on.

The CNN network was trained for a maximum of 50 epochs. The training was stopped when the validation error and training error diverged from one another. Leave-one-patient-out crosstesting was carried out to test the CNN model. The training was performed on a Tesla K40 GPU. We applied the Python built Theano library [23, 24] to compile to CUDA (a GPU interpretable language). We made use of the Lasagne library [25] to build and configure our CNN model.

3. RESULTS

We carried out a benchmark experiment to compare CNN against a variety of classifiers such as Support Vector Machines (SVM) with Gaussian kernel [26], Random Forest (RF) [27], k-Nearest Neighbor (KNN) [28], and C4.5 Decision Tree (DT) [29]. The receiver operating characteristic (ROC) curve was computed for each model, and the area under the ROC curve (AUC) was calculated as benchmark criterion. The larger the AUC, the better the performance. The AUC values are listed in Table 1.

Our numerical results show that the CNN achieves the largest AUC of 0.947, which outperforms the other four classifiers. As illustrated in Fig. 3, the ROC curves show that the CNN has favorable specificity and sensitivity compared to the other classifiers. However, due to the small dataset of only 5 patients, AUC scores varied a lot among different patients. Investigation with a much larger dataset is necessary to further explore the capabilities of CNNs in detecting epileptiform spikes from EEG signals.

4. CONCLUSION

In this paper, we develop a CNN model for detecting spikes in the EEG of epileptic patients. Our numerical results for a small pool of 5 patients show that the CNN performs better than four standard classifiers. The CNN model was relatively small compared to previous

contributions in the field of convolutional neural networks, since we limited the number of parameters in order to avoid overfitting. In future work, we will consider datasets of hundreds of epilepsy patients, and will train larger CNN models. Large-scale CNN models may yield even better detection results.

References

1. Epilepsy Foundation of America. Epilepsy Foundation; 2014. About Epilepsy: The Basics. <http://www.epilepsy.com/learn/about-epilepsy-basics>
2. Racette, Brad A., Holtzman, David M., Dall, Timothy M., Drogan, Oksana. Supply and demand analysis of the current and future us neurology workforce. *Neurology*. 2014; 82(24):2254–2255. [PubMed: 24935638]
3. Jin, Jing, Dauwels, Justin, Cash, Sydney, Westover, M Brandon. Spikegui: Software for rapid interictal discharge annotation via template matching and online machine learning. *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE; IEEE; 2014. p. 4435–4438.*
4. Wilson, Scott B., Emerson, Ronald. Spike detection: a review and comparison of algorithms. *Clinical Neurophysiology*. 2002; 113(12):1873–1881. [PubMed: 12464324]
5. Gotman, Jean, Wang, Li-Yan. State dependent spike detection: validation. *Electroencephalography and clinical neurophysiology*. 1992; 83(1):12–18. [PubMed: 1376661]
6. De Oliveira, Pedro Guedes, Queiroz, Carlos, Da Silva, Fernando Lopes. Spike detection based on a pattern recognition approach using a microcomputer. *Electroencephalography and clinical neurophysiology*. 1983; 56(1):97–103. [PubMed: 6190638]
7. Wilson, Scott B., Turner, Christine A., Emerson, Ronald G., Scheuer, Mark L. Spike detection ii: automatic, perception-based detection and clustering. *Clinical neurophysiology*. 1999; 110(3):404–411. [PubMed: 10363762]
8. Gotman J, Wang LY. State-dependent spike detection: concepts and preliminary results. *Electroencephalography and clinical neurophysiology*. 1991; 79(1):11–19. [PubMed: 1713547]
9. Faure, Claudie. Attributed strings for recognition of epileptic transients in EEG. *International journal of bio-medical computing*. 1985; 16(3):217–229. [PubMed: 3891638]
10. Davey BLK, Fright WR, Carroll GJ, Jones RD. Expert system approach to detection of epileptiform activity in the EEG. *Medical and Biological Engineering and Computing*. 1989; 27(4):365–370. [PubMed: 2689807]
11. Webber WRS, Litt Brian, Wilson K, Lesser RP. Practical detection of epileptiform discharges (EDs) in the EEG using an artificial neural network: a comparison of raw and parameterized EEG data. *Electroencephalography and clinical Neurophysiology*. 1994; 91(3):194–204. [PubMed: 7522148]
12. Gabor, Andrew J., Seyal, Masud. Automated in-terictal EEG spike detection using artificial neural networks. *Electroencephalography and clinical Neurophysiology*. 1992; 83(5):271–280. [PubMed: 1385083]
13. Ramabhadran, Bhuvana, Frost, James D., Jr, Glover, John R., Ktonas, Periklis Y. An automated system for epileptogenic focus localization in the electroencephalogram. *Journal of clinical neurophysiology*. 1999; 16(1):59–68. [PubMed: 10082093]
14. Black, Michael A., Jones, Richard D., Carroll, Grant J., Dingle, Alison A., Donaldson, Ivan M., Parkin, Philip J. Real-time detection of epileptiform activity in the EEG: a blinded clinical trial. *Clinical EEG and Neuroscience*. 2000; 31(3):122–130.
15. Krizhevsky, Alex, Sutskever, Ilya, Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In: Pereira, F.Burges, CJC.Bottou, L., Weinberger, KQ., editors. *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc; 2012. p. 1097–1105.
16. He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, Sun, Jian. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *CoRR*. 2015 abs/1502.01852.
17. Zhang, Xiang, LeCun, Yann. Text understanding from scratch. *CoRR*. 2015 abs/1502.01710.

18. Sønderby, Søren Kaae, Sønderby, Casper Kaae, Nielsen, Henrik, Winther, Ole. Convolutional lstm networks for subcellular localization of proteins. arXiv preprint arXiv:1503.01919. 2015
19. Lawrence, Steve, Giles, C Lee, Tsoi, Ah Chung, Back, Andrew D. Face recognition: A convolutional neural-network approach. *Neural Networks, IEEE Transactions*. 1997; 8(1):98–113.
20. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. *CoRR*. 2014 abs/1409.1556.
21. Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, Salakhutdinov, Ruslan. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*. 2014; 15:1929–1958.
22. Sutskever, Ilya, Martens, James, Dahl, George, Hinton, Geoffrey. On the importance of initialization and momentum in deep learning. In: Dasgupta, Sanjoy, Mcallester, David, editors. *Proceedings of the 30th International Conference on Machine Learning (ICML-13); JMLR Workshop and Conference Proceedings*; May. 2013 p. 1139-1147.
23. Bergstra, James, Breuleux, Olivier, Bastien, Frédéric, Lamblin, Pascal, Pascanu, Razvan, Desjardins, Guillaume, Turian, Joseph, Warde-Farley, David, Bengio, Yoshua. Theano: a CPU and GPU math expression compiler. *Proceedings of the Python for Scientific Computing Conference (SciPy)*; June 2010; Oral Presentation
24. Bastien, Frédéric, Lamblin, Pascal, Pascanu, Razvan, Bergstra, James, Goodfellow, Ian J., Bergeron, Arnaud, Bouchard, Nicolas, Bengio, Yoshua. Theano: new features and speed improvements. *Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop*; 2012.
25. Dieleman, Sander, Schlter, Jan, Raffel, Colin, Olson, Eben, Snderby, Sren Kaae, Nouri, Daniel, Maturana, Daniel, Thoma, Martin, Battenberg, Eric, Kelly, Jack, De Fauw, Jeffrey, Heilman, Michael, diogo149, McFee, Brian, Weideman, Hendrik, takacsg84, peterderivaz, instagibbs, Jon, Rasul, Kashif, Dr, Britefury, CongLiu, Degrave, Jonas. Lasagne: First release. Aug.2015
26. Hearst, Marti A., Dumais, Susan T., Osman, Edgar, Platt, John, Scholkopf, Bernhard. Support vector machines. *Intelligent Systems and their Applications, IEEE*. 1998; 13(4):18–28.
27. Svetnik, Vladimir, Liaw, Andy, Tong, Christopher, Culberson, J Christopher, Sheridan, Robert P., Feuston, Bradley P. Random forest: a classification and regression tool for compound classification and qsar modeling. *Journal of chemical information and computer sciences*. 2003; 43(6):1947–1958. [PubMed: 14632445]
28. Peterson, Leif E. K-nearest neighbor. *Scholarpedia*. 2009; 4(2):1883.
29. Quinlan, J Ross. Induction of decision trees. *Machine learning*. 1986; 1(1):81–106.




Fig. 1.
Illustration of (a) an interictal epileptiform spike, and (b) a background waveform.




Fig. 2.
The architecture of a the CNN model applied.




Fig. 3.
ROC curves for various statistical models.

Table 1

AUC for the CNN model and standard classifiers.

Models	AUC
CNN	0.947
SVM	0.912
RF	0.883
KNN	0.835
DT	0.817

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript