A Generative Adversarial Approach for the Detection of Typical and Drowned Action Potentials

Alexandru Dodon
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
alex@alexdodon.net

Marius Adrian Calugar
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
calugar.marius25@gmail.com

Rodica Potolea
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
rodica.potolea@cs.utcluj.ro

Camelia Lemnaru
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
camelia.lemnaru@cs.utcluj.ro

Mihaela Dînşoreanu
Technical University of Cluj-Napoca
Cluj-Napoca, Romania
mihaela.dinsoreanu@cs.utcluj.ro

Vasile V. Moca
Transylvanian Institute of Neuroscience
Cluj-Napoca, Romania
moca@tins.ro

Raul C. Mureşan
Transylvanian Institute of Neuroscience
Cluj-Napoca, Romania
muresan@tins.ro

***Abstract*—Spike sorting methods are central for the interpretation of extracellular recordings and provide insight into more complex brain processes. To be able to apply the sorting methods spikes are detected, conventionally, by using an amplitude threshold that takes into account the features of the signal. This offers a straightforward detection solution that is able to discriminate the background noise, but comes at the cost of ignoring spikes with an amplitude below the threshold and limiting the picture provided by the recording. We propose a machine learning approach that captures spikes overlooked when considering only the amplitude. The proposed method leverages Generative Adversarial Networks (GANs) to learn particularities of the spike and noise components of the recordings and then extract the underlying components of novel samples with the goal of detecting the presence of spikes. We quantify the detection metrics with respect to the signal to noise ratio of the recording and show that the proposed method exhibits significantly higher sensitivity while providing comparable specificity to conventional spike detection methods across all signal environments, especially when spikes are almost drowned in the background noise.**

***Index Terms*—Neuroscience, Signal processing, Machine learning**

I. Introduction

The recent advancements in the processes of extracting and storing datasets from various biomedical fields, led to the application of some Machine Learning techniques, shown to be effective in situations where the data is plentiful, on tasks in these domains. *Spike sorting* is a class of techniques that operate mainly on electrophysiological data, with the goal of making sense of the neuronal chatter by identifying the moments when certain neurons fire. The task of *Spike Detection* is a step in the process of Spike Sorting, in which one tries to detect the presence of action potentials in data signals. This detection is usually done with a threshold based method. This approach is not ideal because it will either be too strict in including samples, thus losing a significant part of the recording or it will allow more samples with the trade off that a part of these will be polluted with noise. For addressing these issues, this paper proposes a new method of performing Spike Detection by applying Source Signal Separation.

Source Signal Separation is the task of obtaining a decomposition formed from several source signals from a set of mixed signals. Generally, the difficulty of this task stands in its underdetermination and is often addressed by imposing constraints on the modeled mixing process or source signals that may be derived either from a generative model or justified by good empirical performance.

In a recent advancement [1], these source signals are estimated with the help of a machine learning framework called Generative Adversarial Networks [2]. GANs are able to learn output data distributions without the need of explicitly defining a parametric formula for them. In our scenario the additive sources are represented by the background noise and the actual spike detected in each data sample. We noticed that the initial GAN formulation was hard to train so we chose to perform our experiments with a Wassertein-GAN formulation [3] that implemented gradient penalty [4].

With these configurations we showed a performance improvement on the task of Spike Detection in comparison with the state of the art threshold method, especially in scenarios where the spike component of the signal is deeply hidden in the background noise.

II. Related Work in Spike Detection and Source Separation

As was stated above, the most common method of detecting action potentials in extracellular signals is done by setting a

978-1-6654-0976-6/21/$31.00 ©2021 IEEE

hard threshold that is calculated either manually or automatically as the signal's mean plus a factor of standard deviation. Taking standard deviation of the signal for the threshold, in scenarios where the spike amplitude is large leads to values that are too high. To address this issue, an approximation for calculating the threshold while taking into account the standard deviation of the background noise was proposed in [5]:

$$Thr = 5\sigma_n \quad \sigma_n = \text{median}\left\{\frac{|x|}{0.6745}\right\} \tag{1}$$

where x is the band-pass filtered signal. Another approach of tackling the Spike Detection problem was made in [6], with a template matching method.

Regarding the initial step of training the GAN models, the inclusion of a **Wassertein** distance in the loss of the network was shown [3] to diminish some of the most common issues that could appear while training. In this formulation the loss function minimizes the Wasserstein-1 distance between the generated and original data distributions. Furthermore the authors in [4] have found a new way of imposing the Lipschitz constraint on the loss function of the critic, by including a new term called gradient-penalty that also replaces the need for parameter clipping.

In the paper "Generative adversarial source separation" [1], the authors have used a Wassertein GAN approach in order to compute the estimation of the additive sources of a given mixed signal. The main advantage of using WGANs for this task is that there is no need for the specification of output distribution. A mixed signal $x \in \mathbb{R}^L$ is modeled by the following generative process:

$$\begin{aligned} h_1 \sim p_{\text{latent}}\left(h_1\right), s_1 \mid h_1 \sim p_{\text{forward }_1}\left(s_1 \mid h_1\right) \\ h_2 \sim p_{\text{latent}}\left(h_2\right), s_2 \mid h_2 \sim p_{\text{forward }_2}\left(s_2 \mid h_2\right) \\ x \mid s_1, s_2 \sim p_{\text{mixture}}\left(x \mid s_1+s_2\right) \end{aligned} \tag{2}$$

Where h1 and h2 are latent variables and $p_{\text{forward }_1}\left(s_1 \mid h_1\right)$ is the forward model for the sources.

The sources are then approximated by the the respective generators and the synthetic mixture is compared to the original one and a loss is computed and used for gradient descent on the latent variables of the generators. This results in exploring the learned distributions and producing a plausible approximation of the given mixture, the respective sources being extracted in the generated components.

III. Materials and Methods Used

A. Dataset Description

To be able to train the generative models and to evaluate the detection performance, we employed a synthetic dataset [7] that was generated with the purpose of benchmarking spike sorting algorithms. By the generation process, the dataset is labeled and this offered a good environment for both training the generative models and to validate the performance of the method.

This dataset consists of 95 simulations, generated from a database of monkey recordings, of a length of 10 minutes each, out of which we used 24 for a cumulative recording time of 4 hours. Each simulation was created at 96 KHz and was then resampled at 24KHz to mimic real recordings that sample continuous signals at discrete time intervals. The dataset includes both single unit and multi unit activity beside the noise. These are different categories of activity with regards to spike sorting, single unit activity falling more closely to the electrode tip - up to $50\mu m$- while multi unit activity falls from $50\mu m$ to $140\mu m$. The original waveforms of the spikes had 316 points at 96 KHz. They have been mixed and scaled in the generated signal and down sampled at 24 KHz resulting in a length of 79 points. For each original waveform we know the neuron to which it is related and the position of its first sample in the generated signal.

Layer	Output Shape
Latent Variable	1 X 200
Linear	1 X 80
SoftPlus	1 X 80
Linear	1 X 80
SoftPlus	1 X 80
Linear	1 X 80

TABLE I
Architecture of the Generator

Layer	Output Shape
Input Sample	1 X 80
Linear	1 X 40
Tanh	1 X 40
Linear	1 X 1

TABLE II
Architecture of the critic

B. Generative Adversarial Networks

The utilized architectures and activation functions for the generators and critics are shown in Tables I and II respectively. As the neural networks that are very deep tend to need huge amounts of data in order to perform their tasks, we decided to keep the models fairly shallow with only 3 layers for the critic and 5 for the generator.

C. Training Setup

For the training of the models we adopted a WGAN-GP configuration, where the loss function of the critic over a batch of size m consisting in real samples x^i, samples of the random distribution h^i, and a random number ϵ sampled uniformly between 0 and 1 is:

$$\begin{aligned} &\tilde{\boldsymbol{x}} \leftarrow G_\theta(\boldsymbol{h}) \\ &\hat{\boldsymbol{x}} \leftarrow \epsilon \boldsymbol{x}+(1-\epsilon) \tilde{\boldsymbol{x}} \\ &\min_w \frac{1}{m} \sum_{i=1}^m C_w(G_\theta(h^i))-\frac{1}{m} \sum_{i=1}^m C_w(x^i)+ \\ &\lambda \frac{1}{m} \sum_{i=1}^m\left(\left\|\nabla_{\hat{\boldsymbol{x}}} C_w(\hat{\boldsymbol{x}})\right\|_2-1\right)^2 \end{aligned} \tag{3}$$

Here λ is the penalty coefficient and we have set it to 10. The Generator's loss is then:

$$\min_\theta -\frac{1}{m} \sum_{i=1}^m C_w(G_\theta(h^i)) \tag{4}$$

We used the RMSProp optimizer with a learning rate of 0.0001 and performed 5 training epochs with a batch size of 64. For each training iteration of the generator the critic was trained 10 times.

D. Source Separation

The generated mixture consists in a linear combination of the outputs from the two generators and the separation process is driven by the chosen loss function:

$$\min_{h_1,h_2} \frac{1}{n} \sum_{t=0}^{n-1} \left(f_{\widehat{\theta}_1}\left(h_1^t\right) + f_{\widehat{\theta}_2}\left(h_2^t\right) - S_{signal}^t \right)^2 \qquad (5)$$

Where S_{signal} is the signal that is undergoing separation, $f_{\widehat{\theta}_1}$ and $f_{\widehat{\theta}_2}$ are the feed forwards of the generators for spikes and noise and h_1 and h_2 are the respective latent variables.

A hyperparameter of the separation process is the number of iterations of gradient descent, as well as the parameters for the optimizer. We employed the Adam optimizer with a learning rate of 0.001 and betas of 0.5 and 0.999. The number of epochs was set to 3000.

E. Interpretation and Classification of Separated Signals

As a result of the separation, the classification process is straightforward: we take the energy of the extracted spike component and threshold it based on a threshold selected for the best F1 score on the validation set.

F. Empirical Evaluation

In all our experiments, we applied a training / validation / testing split of 80%, 10% and 10% on a balanced dataset. For an apt evaluation of the method, we chose to compare it with amplitude thresholding methods that use a multiple for the standard deviation from 1 to 5.

To be able to quantify the detection behavior for spikes that are disregarded when considering just the amplitude, we test the methods at different signal to noise ratios (SNR) from 1 to 6.

A uniform signal to noise ratio is achieved in the validation and test dataset by artificially "drowning" spikes. We model a drowned spike window as a linear combination between a scaled spike and a noise window:

$$W_{\text{drowned}} = SF \cdot W_{\text{spike}} + W_{\text{noise}} \qquad (6)$$

Then, for a given SNR and for each pair of spike and noise window the scaling factor SF can be derived from the definition of SNR:

$$SF = \sqrt{SNR \cdot \frac{\sum w_{\text{noise}}^2}{\sum w_{\text{spike}}^2}} \qquad (7)$$

The results of the process for a target SNR of 1 are exemplified in Figure 1.

To create a robust evaluation of the proposed method, comparing it to existing methods, 10 iterations of the validation pipeline have been performed. For one iteration the following have occurred:

- The balanced dataset consisting of 400 thousand windows is shuffled and then split in 80% training, 10% validation and 10% test dataset.

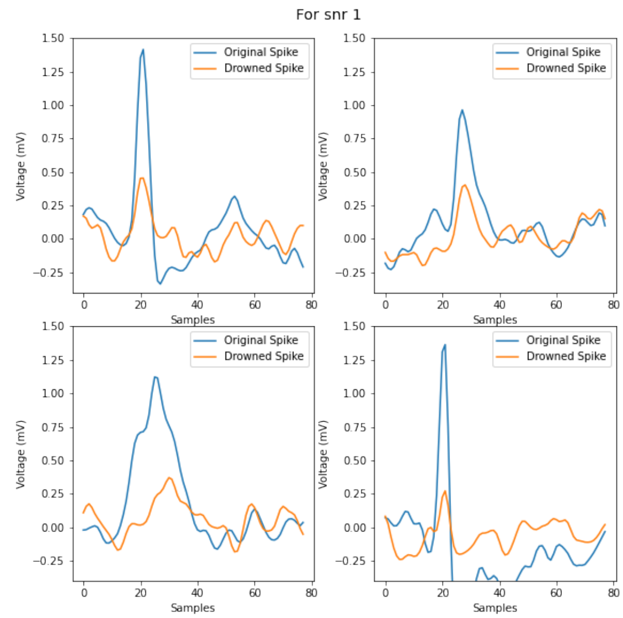

Fig. 1. Drowned Spike windows with a SNR of 1

- For each validation and test dataset variations with specific SNRs have been created.
- For each training dataset the GANs have been trained.
- For each combination of training dataset, validation and test dataset with a specific SNR the following have been run:
  - The separation and classification of the signal windows with our method.
  - The thresholding of signal windows with current thresholding methods.

At the end of each iteration, the predicted labels of each method are collected and then used to compute the respective sensitivities(true positive rates) and specificities(true negative rates).

IV. Results

We display the means and standard deviations of sensitivity and specificity for all methods after 10 iterations in Figure 2. The lightly shaded area represents the standard deviation.

Across all tested methods, the sensitivity decreases when applied on data sets with a lower Signal to Noise Ratio. This is to be expected for the thresholding methods because as SNR decreases, the amplitude of the spikes present in the signal also decreases and they become statistically similar to background noise. After all, the background noise is created by the action potentials of neurons whose firing is attenuated when it reaches the electrode. That being said, our proposed method does not singularly look at the maximum amplitude to detect spikes. Therefore, it holds better recognition as the spikes become drowned in the background noise.

The specificity of the methods also offer an interesting insight inside the workings of spike detection. Threshold based

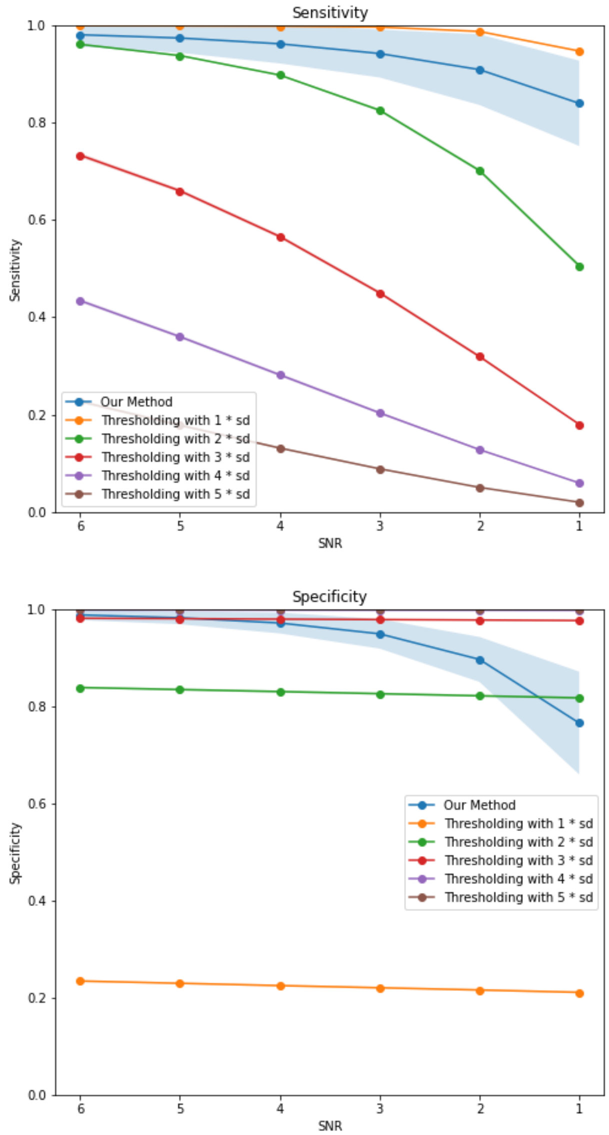

Fig. 2. Performance metrics.

algorithms for spike detection maintain almost constant specificity. This is to be expected as the underlying characteristics of the noise does not change. On the other hand, the proposed method's specificity drops as it targets spikes more and more embedded in the background noise.

In typical usage at a SNR of 2.5 to 3.5 it exhibits a better sensitivity than all other methods except the 1 SD threshold which due to its low specificity is not viable in a laboratory setting. Meanwhile, our method's specificity is comparable to the most conservative thresholds. Especially at a SNR of 1, the other methods become less attractive as while they maintain their specificity, they lose most of the positive samples. On the other hand, in the proposed configuration, our method, while its specificity becomes comparable to thresholding with 2 SD, it's sensitivity remains almost the same. This results in capturing a significantly bigger portion of spikes that historically could not be captured.

V. Discussion

The method's ability to distinguish spikes that are still well formed as they become drowned in the background noise opens up new avenues of research by expanding the available data from recordings that already exist.

The key advantage of the proposed method is that the generators were trained on typical spikes, so it can be bootstrapped on unlabeled recordings by existing spike detection methods. This can allow the method to be laboratory-ready by taking the samples relevant to the object-level problem and training the model to distinguish spikes recorded by the specific hardware available. The proposed method can also be applied on any number of channels from a given recording.

The information used to enhance classification can be increased by passing the extracted spike and noise sources through the feed forward networks of the respective critics to rate their plausibility.

These interpretations of the extracted sources can then be used in conjunction with the extracted energy to build the feature vector for a standard linear classifier or a multilayer perceptron.

After the generative models are trained, the classifier can be fed windows with a fixed length that are offsetted by a single sample in order to filter a continuous signal of any length and the spikes can be detected by setting a confidence threshold.

VI. Conclusion

We developed a method that is able to leverage generative models for the separation of electrophysiological signals into spike and noise components with the objective of detecting spikes. We showed promising performance increases compared to current amplitude threshold-based methods, especially when considering spikes drowned in the background noise.

VII. Acknowledgement

The research leading to these results was partially funded by: NO Grants 2014-2021, under Project contract number 20/2020 (RO-NO-2019-0504), three grants from the Romanian National Authority for Scientific Research and Innovation, CNCS-UEFISCDI (codes PN-III-P2-2.1-PED-2019-0277, PN-III-P3-3.6-H2020-2020-0118, PN-III-P3-3.6-H2020-2020-0109), and a H2020 grant funded by the European Commission (grant agreement 952096, NEUROTWIN).

References

[1] Y. C. Subakan and P. Smaragdis, "Generative adversarial source separation,"CoRR,vol. abs/1710.10779, 2017.

[2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio, "Generative adversarial nets," in NIPS, 2014.

[3] Mart´ın Arjovsky, Soumith Chintala, and L´eon Bottou, "Wasserstein GAN," CoRR, vol. abs/1701.07875, 2017.

[4] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein gans," 2017. arXiv:1710.10779 [cs, stat], Oct. 2017, arXiv: 1710.10779.

[5] Quian Quiroga R, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised Spike Detection and Sorting with Wavelets and Superparamagnetic Clustering. Neural Comp 16:1661-1687.

[6] K. J. Laboy-Juarez, S. Ahn, and D. E. Feldman, "A normalized template matching method for improving spike detection in extracellular voltage recordings," Scientific reports, vol. 9, no. 1, pp. 1–12, 2019.
[7] C. Pedreira, J. Martinez, M. J. Ison, and R. Quian Quiroga, "How many neurons can we see with current spike sorting algorithms?" vol. 211, no. 1, pp. 58–65