

Computationally efficient neural network classifiers for next generation closed loop neuromodulation therapy – a case study in epilepsy

Ali Kavvoosi^{1†}✉, Robert Toth^{2†}, Moaad Benjaber¹, Mayela Zamora¹,
Antonio Valentin³, Andrew Sharott² and Timothy Denison^{1,2}✉

Abstract—This work explores the potential utility of neural network classifiers for real-time classification of field-potential based biomarkers in next-generation responsive neuromodulation systems. Compared to classical filter-based classifiers, neural networks offer an ease of patient-specific parameter tuning, promising to reduce the burden of programming on clinicians. The paper explores a compact, feed-forward neural network architecture of only dozens of units for seizure-state classification in refractory epilepsy. The proposed classifier offers comparable accuracy to filter-classifiers on clinician-labeled data, while reducing detection latency. As a trade-off to classical methods, the paper focuses on keeping the complexity of the architecture minimal, to accommodate the on-board computational constraints of implantable pulse generator systems.

Clinical relevance—A neural network-based classifier is presented for responsive neurostimulation, with comparable accuracy to classical methods at reduced latency.

I. INTRODUCTION

Deep brain stimulation (DBS) first received approval for the symptomatic treatment of Parkinson's disease in 1997. While similar in design to cardiac pacemakers, the implantable pulse generators (IPG) of the time offered only an open-loop form of therapy, with typically a single stimulation pattern, set by a clinician for each patient. Real-time seizure detection and responsive neurostimulation (RNS) was first attempted using a computer-in-the-loop system by Gotman et al. in 1976 [1], it was Osorio et al. in 1998 [2], [3], [4] who introduced the more widely studied filter-based spectral biomarker detectors to the field of epilepsy research. With the continued development of IPGs and the maturation of low-power microprocessor technology, the first RNS system for epilepsy received approval for pre-clinical use in 2014. This system from Neuropace had the capabilities to sense bioelectric signals, and choose stimulation programs based on clinician-configured classification state [5].

This work was supported by the John Fell Fund of the University of Oxford, the UK Medical Research Council (MC_UU_00003/3, MC_UU_00003/6) and the Royal Academy of Engineering.

† Ali Kavvoosi and Robert Toth contributed equally to this work and share first authorship.

¹ Brain Network Dynamics Unit, Department of Pharmacology, University of Oxford, Oxford OX1 3TH, United Kingdom

² Institute of Biomedical Engineering, Old Road Campus Research Building, Department of Engineering Sciences, University of Oxford, Oxford OX3 7DQ, United Kingdom

³ Department of Basic and Clinical Neuroscience, King's College London, London SE5 9RT, United Kingdom

✉ ali.kavvoosi@linacre.ox.ac.uk, timothy.denison@eng.ox.ac.uk

Filter-based spectral-feature detectors have since been used successfully in other conditions, most notably for tremor suppression in Parkinson's disease, through the discovery of beta oscillations as a correlate of disease state [6]. However, the smaller ($1\mu\text{Vrms}$) signal size of beta oscillations, compared to epileptiform activity ($10\mu\text{Vrms}$), made deploying the detector algorithm in IPGs challenging due to the presence of stimulation and other artifacts. Contemporary work focuses on improving the robustness of the signal chains to enable simultaneous sensing and stimulation, thus true closed-loop operation across targeted diseases [7], [8]. Examples include the Medtronic Percept [9] and the Picostim-DyNeuMo research systems [10], [11]. A complementary avenue of refinement is the use of feed-forward predictors for adapting stimulation based on periodicities of disease state and patient needs, such as the circadian scheduling of the SenTiva system from LivaNova or the Picostim-DyNeuMo [12]. Taking advantage of more long term, weekly or even monthly rhythms are being investigated for epilepsy management [13].

Patient-specific filter design, while possible to aid with software, can be a complex problem, likely to limit both clinician-capacity and patient-throughput. Establishing and validating a neural-network (NN) training pipeline based on clinician-labeled data could offer a systematic classifier tuning process. Networks could be pre-trained on aggregate data from multiple patients, and refined based on individual labeled data at the point of deployment [14]. Of course, due to the black-box nature of neural network classifiers, extensive validation work will be required to establish safety before first-in-human studies. Advances in interpretable deep learning could facilitate building trust in NN-classifiers for medical use [15].

Liu et al. [16] demonstrated the feasibility of deploying high accuracy classifiers for seizure detection on modern microprocessors (ARM[®] Cortex-M4), through model compression and quantization techniques, showcasing several advanced NN topologies.

This paper is meant as an initial study to bring focus to the fundamental challenge of NN classifiers: computational cost. As state-of-the-art deep neural networks reach ever increasing model sizes [14], [17], we aim to explore whether lean NNs of only dozens of units could in fact compete in accuracy with classical, filter-based systems for bioelectric signal classification.

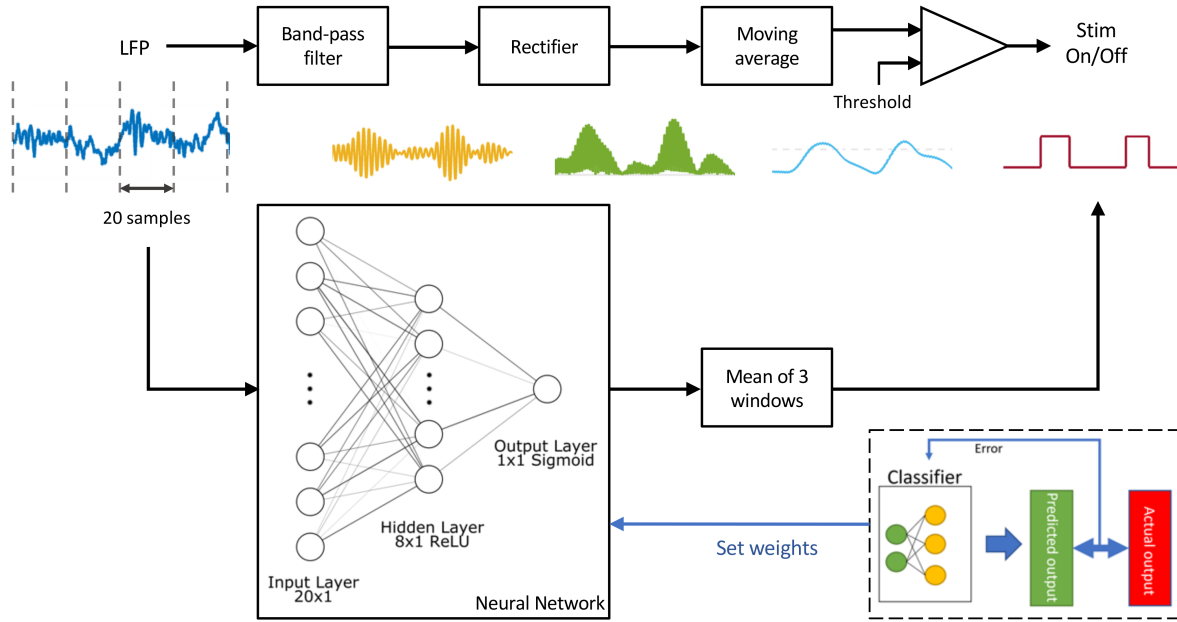


Fig. 1. Classifier architectures. *Top*: classical filter-based spectral power detector [6]. *Bottom*: the multi-layer perceptron architecture evaluated in this paper. Through training on labeled data, the neural network is expected to assume an overall transfer function similar to the hand-crafted filter topology.

II. DESIGN

A. Baseline Method

To establish a baseline for performance as well as for computational cost, we used a classical band-power estimation filter chain to detect epileptiform discharges [3], [6], that we have previously deployed with success in the Picostim-DyNeuMo experimental IPG system [11]. The processing steps of this method are shown in the top panel of Fig. 1. While this algorithm is computationally efficient and has a very favorable memory footprint (refer to Table I), the demodulated envelope signal, thus the detector output, will always lag the input signal to reduce output ripple – irrespective of processing speed. This trade-off arises from the very nature of causal filtering, and is necessary to prevent rapid switching of the detector output for input signals near the classification threshold. The reference classifier was configured as follows. Our band-pass stage was an 8–22 Hz, 4th order Butterworth filter, with a Direct Form I IIR implementation (16-bit coefficients, 32-bit accumulators). Envelope demodulation was achieved using an exponential moving average filter with a decay-coefficient of 32 samples. The filter chain, and all other classifiers were designed to operate at a sampling rate of 256 Hz.

B. New neural-based methods

In our search for a low-complexity classifier for time series input, we explored two main NN families. (1) Multi-Layer Perceptrons (MLP) are the simplest, and oldest family or of artificial neural networks [18], where the input vector is connected to ‘hidden’ layers of feedforward units, condensing information into an output unit. This architecture is shown in the bottom panel of Fig. 1. (2) As a step up in complexity, Convolutional Neural Network (CNN) introduce a convolution layer, also known as a filter bank, between

the input vector and the neural layers as an initial extra step [14]. The input to our networks is formed by a windowed set of past time samples of the local field potential (LFP) signal. The output signal, calculated once for each complete window of samples, is thresholded into a binary label. We denote this classifier the ‘standalone MLP’ model.

Recurrent neural networks, an otherwise natural choice for processing time series data, were dismissed from consideration as recurrence necessitates the introduction of dynamic state variables, which significantly increases memory footprint [14]. Without recurrence, we introduced coherence into our classifier in a different way. We settled on requiring a consensus of three subsequent outputs from the NN to define our final output label, providing the ‘adjusted MLP’ model.

C. Training and data

Our raw dataset consisted of LFP signals recorded from two patients, for a combined 24 hours, with 30 professionally labeled events of clinical significance. The recordings were resampled to a 256 Hz sampling frequency for uniformity.

As seizures are comparatively rare events scattered among very long periods of normal activity, we decided to introduce class imbalance into our training sets to best prepare the NNs for real-life use. The training set was biased towards negative samples in a 3:1 ratio, based on clinician annotations. The dataset was split in the common 70:30 ratio between a training and a validation set. Network weights and biases were quantized to 8-bit integers.

D. Technical equipment used

Neural networks were modeled and trained in Tensorflow Lite version 2.7.0, using an Intel Core i7 CPU with 16 GB of RAM. Embedded performance was tested on an Arduino Nano 33 BLE Sense evaluation board for the nRF52840 ARM Cortex-M4F microprocessor.




Fig. 2. ROCs for different classifiers. Note that performance converges towards high TPR and FPR, which is the desirable operating point of seizure detectors as FNs pose significantly greater risk of harm to patients than FPs.

E. Comparison of CNN and MLP models

Fig. 2 shows the performance of our CNN and MLP classifiers. While the CNN outperforms the standalone MLP, it performs with similar accuracy to the adjusted MLP model above 80% true positive rate (TPR). For safe use, the operating point of a seizure detection system should be biased towards high TPR – missed seizures (false negatives) pose significantly more risk to the patient than false positives, which merely result in unnecessary stimulation. Overall, in targeting resource constrained IPGs, we judged the minor edge of the CNN insufficient to justify the added computational burden of the convolutional layer.

F. Tuning the MLP classifier

The performance of the MLP model, when trained on a given dataset, is primarily determined by two hyperparameters: the number of timepoints in the input window, and the hidden layer’s size. We found that varying the number of hidden layers had very modest effects on accuracy (not shown in this paper). Fig. 3 systematically explores the effect of the two key hyperparameters on the classification error of a single output, single hidden layer MLP model. As expected, the network requires a certain size and complexity to encode a feature space sufficient for reliable classification, though increasing the number of units in either layer beyond a certain point leads to diminishing returns. To select one of the possible models from the error surface, one could define a scoring scheme including network size, computational time and the loss itself, to make an educated choice, however, this is beyond the scope of this paper. Favoring low complexity, we settled on using a 20-point input window and 8 hidden neurons, in the ‘transition zone’ of the error surface.

III. MODEL PERFORMANCE AND INTERPRETATION

The next step is to compare our best MLP result to the baseline classifier. The ROC of Fig. 2 reveals that a well-tuned filter chain outperforms the small MLP model below 60% false positive rate, beyond which they converge in accuracy. Identifying a seizure does not present a holistic view of performance though. In Fig. 4 we highlight two




Fig. 3. Tuning the MLP classifier. Grid search on the two main model hyperparameters: the input vector length and the number of hidden layer units. Loss is represented as binary cross entropy across all samples.

additional characteristics to consider in classifier evaluation: (1) latency at event onset, and (2) the overlap between classifier and clinician labels. As shown, the MLP responds on average more rapidly to a commencing seizure (mean latency of 0.6 sec vs 1.7 sec), and tracks the clinician label more closely overall, compared to the baseline method.

To explore the MLP classifier’s internal representation of a seizure, we present a small interpretation experiment in Fig. 5. We presented the classifiers with second-long sinusoidal bursts of activity, performing parameter sweeps along both test frequency and test amplitudes. As seen, the MLP model (right) was successful in internalizing a notion of the spectral characteristics of epileptiform activity (low frequency lobe), that encompasses the pass-band of the filter classifier (left). The greater effective bandwidth of the MLP could explain the lower false positive rates seen on the ROC of Fig. 2. The activation lobes at higher frequencies likely represent a process analogous to aliasing, and we expect this periodicity to be a correlate of the input window size, which should be investigated further.

IV. DISCUSSION

The example design of the MLP classifier demonstrates that even tiny neural networks can be effective at simple signal processing tasks. As the final step, we should reflect




Fig. 4. Detailed performance of classifiers. *Top*: histogram of classification latency. *Bottom*: histogram showing the percentage of overlap between positive classifier output and clinician-labeled event.




Fig. 5. Frequency response of classifiers to input signals of different magnitudes. *Left*: filter chain classifier. *Right*: MLP classifier with average of 3 windows. Accuracy is presented as the mean classifier output for a 1 second sinusoidal test tone over 10 repeats.

TABLE I
RESOURCE USAGE

Classifier	Execution time (Cycles/Sample)	Code (bytes)	Memory (bytes)
Filter (custom)	750	675	100
MLP (TF Lite)	2250	31k	8k
CNN (TF Lite)	3200	31K	13k
MLP (custom)	2250*	2.0k*	300*

*predicted

on the embedded resource usage achieved, so we refer the reader to Table I. Importantly, the network achieved sufficiently low complexity for real time use. Note that NN execution times are reported per sample, though the output only changes at the end of a window of 20 samples. Notably, the true memory footprint of the classifier could not be determined with this evaluation system – Tensorflow Lite does not generate network code, rather it provides a network description file, to be run by a relatively large, general purpose interpreter library in the embedded system. For a more realistic, yet conservative outlook, we present estimates for the memory usage of the same network deployed using customized library, trimmed down to the features used in our design. In summary, the MLP could provide an alternative to existing tuned filter methods used in commercial devices.

DISCLOSURES

The University of Oxford has research agreements with Bioinduction Ltd. Tim Denison also has business relationships with Bioinduction for research tool design and deployment, and stock ownership (< 1 %).

ACKNOWLEDGMENT

The authors would like to thank Bence Mark Halpern for reviewing the manuscript, and Tom Gillbe at Bioinduction Ltd. for feedback on the filter-classifier design.

REFERENCES

- [1] J. Gotman and P. Gloor, "Automatic recognition and quantification of interictal epileptic activity in the human scalp EEG," *Clin. Neurophysiol.*, vol. 41, no. 5, pp. 513–529, Nov. 1976.
- [2] I. Osorio, M. G. Frei, and S. B. Wilkinson, "Real-time automated detection and quantitative analysis of seizures and short-term prediction of clinical onset," *Epilepsia*, vol. 39, no. 6, pp. 615–627, Jun. 1998.
- [3] I. Osorio *et al.*, "Performance reassessment of a real-time seizure-detection algorithm on long ECoG series," *Epilepsia*, vol. 43, no. 12, pp. 1522–1535, Dec. 2002.
- [4] N. Bhavaraju, M. Frei, and I. Osorio, "Analog seizure detection and performance evaluation," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 2, pp. 238–245, Feb. 2006.
- [5] F. T. Sun and M. J. Morrell, "The RNS system: responsive cortical stimulation for the treatment of refractory partial epilepsy," *Expert Rev. Med. Devices*, vol. 11, no. 6, pp. 563–572, Nov. 2014.
- [6] S. Little *et al.*, "Adaptive deep brain stimulation in advanced parkinson disease: Adaptive DBS in PD," *Ann. Neurol.*, vol. 74, no. 3, pp. 449–457, Sep. 2013.
- [7] M. M. Sorkhabi, M. Benjaber, P. Brown, and T. Denison, "Physiological artifacts and the implications for brain-machine-interface design," *In Proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1498–1498, 2020.
- [8] J. Ansó *et al.*, "Concurrent stimulation and sensing in bi-directional brain interfaces: a multi-site translational experience," *J. Neural Eng.*, vol. 19, no. 2, p. 026025, Mar. 2022.
- [9] W.-J. Neumann *et al.*, "The sensitivity of ECG contamination to surgical implantation site in brain computer interfaces," *Brain Stimul.*, vol. 14, no. 5, pp. 1301–1306, Sep. 2021.
- [10] M. Zamora *et al.*, "DyNeuMo Mk-1: Design and pilot validation of an investigational motion-adaptive neurostimulator with integrated chronotherapy," *Exp. Neurol.*, vol. 351, p. 113977, May 2022.
- [11] R. Toth *et al.*, "DyNeuMo Mk-2: an investigational circadian-locked neuromodulator with responsive stimulation for applied chronobiology," *In Proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 3433–3440, 2020.
- [12] M. Zamora *et al.*, "Case report: Embedding "digital chronotherapy" into medical devices — a canine validation for controlling status epilepticus through multi-scale rhythmic brain stimulation," *Front. Neurosci.*, vol. 15, p. 734265, 2021.
- [13] M. O. Baud *et al.*, "Multi-day rhythms modulate seizure risk in epilepsy," *Nature Commun.*, vol. 9, no. 1, p. 88, Jan. 2018.
- [14] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [15] S. Lapuschkin *et al.*, "Unmasking Clever Hans predictors and assessing what machines really learn," *Nat. Commun.*, vol. 10, no. 1, p. 1096, Dec. 2019.
- [16] X. Liu, and A. G. Richardson, "Edge deep learning for neural implants: a case study of seizure detection and prediction," *J. Neural Eng.*, vol. 18, no. 4, p. 046034, Apr. 2021.
- [17] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" *In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, pp. 610–623, Mar. 2021.
- [18] F. Rosenblatt, "The Perceptron: A probabilistic model for information storage and organization in the brain," *Psychol. Rev.*, vol. 65, no. 6, pp. 386–408, 1958.