Epileptic Seizure Prediction using Big Data and Deep Learning: Toward a Mobile System

Dr. Isabell Kiral-Kornek, IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Dr. Subhrajit Roy, IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Ewan Nurse, **MEng**, The University of Melbourne, 3010 Parkville, VIC, AU and IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Dr. Benjamin Mashford, IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Philippa Karoly, MEng, The University of Melbourne, 3010 Parkville, VIC, AU and IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Thomas Carroll, MEng, The University of Melbourne, 3010 Parkville, VIC, AU and IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Daniel Payne, MEng, The University of Melbourne, 3010 Parkville, VIC, AU and IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Dr. Susmita Saha, IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU
Steven Baldassano, **MD**, The University of Melbourne, 3010 Parkville, VIC, AU

Prof. Terence O'Brien, The University of Melbourne, 3010 Parkville, VIC, AU
Prof. David Grayden, Department of Biomedical Engineering, The University of Melbourne, 3010 Parkville, VIC, AU
Prof. Mark Cook, The University of Melbourne, 3010 Parkville, VIC, AU

Corresponding Authors:
Dr. Dean Freestone, The University of Melbourne, 3010 Parkville, VIC, AU, email: deanrf@unimelb.edu.au, phone: +61-3-0001-3949
Dr. Stefan Harrer, IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, AU, email: sharrer@au.ibm.com, phone: +61-3-9921-3871

All the figures presented in this paper can be explored and examined interactively in the following link:

http://ebiomedicine-12-2017-epileptic-seizure-prediction.surge.sh/

SUPPLEMENTARY MATERIAL

S1) Deep Learning – a brief introduction

Deep learning (LeCun et al., 2015) is a data-driven process that allows an algorithm to adjust itself to perform specific tasks. It does so by exposing the algorithm to a large set of data that exhibits a desired and distinct behaviour. Deep learning as used in our study requires the relevant parts of the data to be appropriately labelled. Typically, deep learning is used to train a class of algorithms known as neural networks to perform specific tasks, such as image recognition or anomaly detection. Deep neural networks have proven to be highly effective for a diverse range of applications, including aiding medical diagnoses (Litjens et al., 2017; Miotto et al., 2017). Similar to the brain, neural networks can contain millions of tunable connectivity strengths that must be able to adapt to perform the task at hand. Deep learning tunes these connections in a highly generalisable way so that a given network can be trained to perform almost any arbitrary task. Learning occurs via a process of reinforcement, whereby the network is optimised based on the errors it is making while analysing the training dataset.

In the context of seizure prediction, the training dataset consists of preictal and interictal iEEG segments and the corresponding labels. The neural network is trained on this training dataset to discriminate between these two classes. On the other hand, the test or inference dataset consists of continuous iEEG signal including all interictal, preictal and ictal segments.

S2) System Architecture

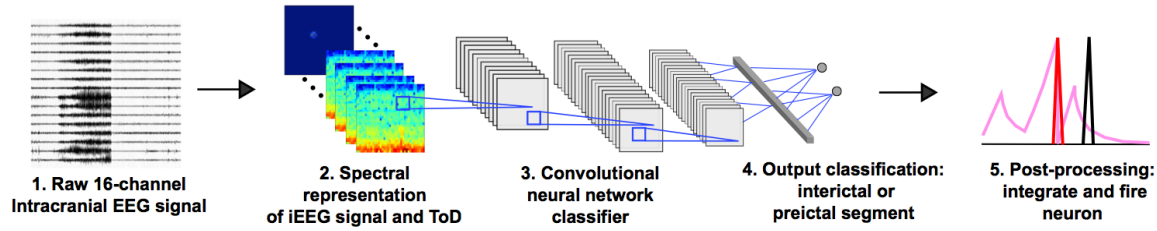

***Fig. S1*: System architecture.** The 16 channels of the raw iEEG signal (**1**) collected through an electrode array are transformed into 16 spectrograms (**2**). An image representing time of day (ToD) information is appended as a 17th channel (**2**). Data are then passed into a convolutional neural network (CNN) for training and classification (**3**). The output of this CNN, a prediction of whether a data segment is inter- or preictal (**4**) is then passed into the post-processing layer (**5**), where the final output is computed using an integrate-and-fire neuron.

a. Data Preparation

i. Conversion of iEEG Signals into Spectrograms

The dataset used was collected during a previous trial undertaken in (Cook et al., 2013). Upon completion of this initial study, the dataset was labelled by expert investigators, using iEEG data as well as secondary seizure evidence (such as audio recordings and patients' seizure diaries). For the purpose of this study, clinical (identified using iEEG and audio recordings/patient diaries) and clinically equivalent (identified using iEEG data only) seizures were treated equally.

For each patient, the recorded iEEG signal was processed before it was ready for being fed into a CNN. Processing involved the following steps:

1. *Dropout removal:* The recorded signal was first scanned for "dropouts" – occurrences of ongoing zeros or missing entries ('NaN' values) simultaneously in all electrode channels. Dropouts are typically the result of data acquisition system failure due to low battery power or the receiver being too far away. These missing values were removed from the signal.

2. *Filtering:* Next, the signal was bandpass filtered from 1 Hz to 140 Hz and notch filtered at 50 Hz (mains removal). Both operations were using zero phase second order non-causal Butterworth filters.
3. *Definition of preictals and interictals:* The previously annotated clinical or clinically equivalent seizures were assigned the ictal class. A 15-minute segment before each seizure was defined as being preictal, starting 16 minutes before an ictal segment and ending one minute before. Parts of the signal that were neither preictal nor ictal were defined as interictal.
4. *Conversion to spectrograms:* Previous studies have shown that analysing the iEEG signal in the frequency domain can be effective for seizure prediction (Howbert et al., 2014; Korshunova et al., 2017). Building on that, we converted the recorded iEEG time series signal into spectrograms. Each 30-sec segment of the time series was partitioned into overlapping 0.445 second frames and within each frame the amplitude spectrum was computed. The generated spectrograms were resized into 32x32 images for feeding into the neural network. Note that, each 30-sec segment of the signal led to the generation of 16 spectrograms corresponding to the 16 electrodes.
5. *Extraction of preictal and interictal training segments:* Once the preictals, interictals, and ictals were marked, the first step in creating the training set was to extract preictal segments. Next, an equal number of interictal samples was randomly selected. Note that although there is a much larger number of interictal than preictal segments in the original signal, a balanced set was extracted to train the network. Moreover, the training set didn't include any ictal samples. However, to mimic a real-world situation, the continuous signal, containing all interictal, preictal, and ictal segments, was used for testing.

ii. Inclusion of Time of Day

It has recently been shown that epilepsy is characterized by a periodic dynamics that increase seizure likelihood at certain times of day (Karoly et al., 2016; Loddenkemper et al., 2011). In the data, we used the specific time of day (24-hour cycle) of when a seizure occurred was noted for each seizure and each patient. This phenomenon suggests that, for these patients, seizures are likely to happen during a certain time of the day or night. For example, Fig. S2 shows the distribution for Patient 9 and Patient 10 which depicts its localized nature and motivates including this data on a patient-specific basis.

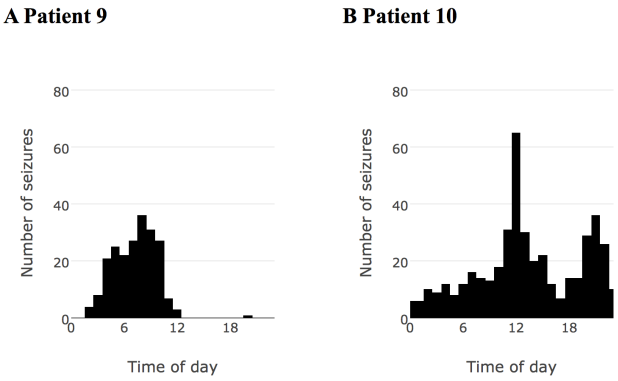

***Fig. S2*: Example of distribution of seizures over time of day for: A** Patient 9, **B** Patient 10

Hence, we decided to include this information with the input. As discussed in Section S2a(i), we first generated 16 spectrograms of size 32x32 corresponding to the 16 deployed electrodes for each sample. Next, the time of day information was appended as a 17^{th} channel to these samples. Note that the way we were preparing the data meant that a single preictal segment could lead to the generation of multiple samples. Each of these samples were annotated with a specific time of day information.

The steps for inclusion of a particular seizure at time x to a given preictal segment are outlined below:

1. A vector T of size 24x1 is created and the x-th position of T is set to k while the other 23 positions are set to 0, where k is a positive value.
2. A matrix TOD of size 24x24 is created by multiplying T and T' ($TOD = T * T'$). Hence, in TOD the entry at (x, x) is 1 while the others are 0.

3. TOD is resized to have a shape of 32x32 and appended to the 16 spectrograms we already had for the 16 electrodes.
4. Note that resizing TOD i.e. a 24x24 matrix to a 32x32 matrix through interpolation adds the time of day information as gradually diminishing probabilities to the neigbouring time segments.

iii. Accounting for Nonstationary iEEG Data in a Real-Time System

The proposed system has been developed for application in a real-world scenario where data is acquired continuously and in real-time. Typically, to develop machine learning algorithms, the split of the available data into training and testing set is done by using a cross-validation technique. During cross-validation, a dataset is randomly split into different sets of training and testing data to compute an average performance. However, in the context of a prediction system like ours, this method would result in the inclusion of future data in the training set. Hence, to mimic the real-world scenario, we created training and test sets in such a way that the test set was always chronologically ahead of the training set.

Since we are using iEEG signals, another important aspect of the data is its non-stationary nature (Sillay et al., 2013; Ung et al., 2017) i.e. the data gradually changes over time. This means that using randomly shuffled data (from cross-validation) would make the task unrealistically easy for a prediction system. Furthermore, to address the non-stationarity of the data, we retrain our model on a regular basis, using only the most recent data. The retraining interval is loosely based on the patient's seizure rate, which led us to retrain the models for patients 1, 2, and 14 every 90 days, while all other models were retrained every 30 days. If a given period chosen for extracting the training data doesn't have any seizures, we use the previous period.

b. Deep Neural Network Topology

We used Convolutional Neural Networks (CNN) to classify between preictals and interictals and thereby predict seizures. Different CNN structures were implemented in TensorFlow ("TensorFlow," 2015) on a desktop computer and in the custom developed IBM TrueNorth Neurosynaptic System (Esser et al., 2016). Both CNN structures are described in the tables below. Here Conv2D, Maxpooling2D, FC, Dropout, ReLU and SoftMax denote two-dimensional convolution, two-dimensional max-pooling, fully-connected layer, dropout, rectified linear unit non-linearity, and the softmax function respectively.

Layer Index	Name	Relevant parameters (number of filters, kernel size, number of output neurons, dropout)
1	Conv2D	16@3x3
2	ReLU	NA
3	Maxpooling2D	2
4	Dropout	0·7
5	Conv2D	32@3x3
6	ReLU	NA
7	Maxpooling2D	2
8	Dropout	0·7
9	Conv2D	32@3x3
10	ReLU	NA
11	Maxpooling2D	2
12	Dropout	0·7
13	FC	32
14	ReLU	NA
15	Dropout	0·5
16	FC	2
17	SoftMax	NA

ST1) ***TensorFlow Network.***

Layer Index	Name	Relevant parameters (number of filters, kernel size, number of output neurons, dropout, padding, stride)
1	Conv2D	12@3x3
2	Conv2D	126@4x4, P=1, S=2
3	Conv2D	128@1x1
4	Conv2D	128@2x2, S=2
5	Conv2D	256@3x3, P=1
5	Conv2D	256@1x1
6	Conv2D	256@1x1
7	Dropout	0·5
8	Conv2D	256@1x1
9	Conv2D	256@2x2, S=2
10	Conv2D	512@3x3, P=1
11	Dropout	0·5
12	Conv2D	512@1x1
13	Conv2D	512@2x2, S=2
14	Conv2D	512@1x1
15	Conv2D	512@1x1
16	Dropout	0·5
17	Conv2D	1020@1x1
18	FC	2

ST2) ***TrueNorth Network.*** *(a ReLU non-linearity is applied to the output of each Conv2D layer).*

The last fully connected layer in each of the networks (FC) outputs two distinct neurons, each indicating the likelihood of a given sample being interictal or preictal. In the Tensorflow implementation, a SoftMax layer forces the sum of the output of both neurons to be equal to one. The output of the preictal neuron was subsequently used during post-processing to predict ictal events. When using TrueNorth, the inverse of the interictal output neuron of the last fully connected layer is used.

While in TensorFlow ("TensorFlow," 2015) we used the Adaptive Moment Estimation (Adam) method to train the CNN, in TrueNorth a custom gradient descent backpropagation rule (Esser et al., 2016) was used. Since TrueNorth is composed of binary synapses, traditional optimisers cannot be directly used to train networks implemented in TrueNorth. Hence, we employed a custom learning rule capable of performing weight updates in a binary neural network. Moreover, to achieve extreme power efficiency, TrueNorth imposes hardware constraints on network structure and weight precision. Hence, the CNN implemented in TrueNorth requires more neurons and synapses that its TensorFlow counterpart.

c. Post-processing Layer:

A post-processing layer is used to translate network outputs to device alarms. It integrates the output of the CNN using an exponentially weighted moving average filter. A threshold crossing triggered a preictal alarm indicating an upcoming seizure. It was adjusted each month to reflect the changes over time in data and model. Performance of the post-processing layer relied on three adjustable parameters, each of which will be described in more detail below.

The 50%-thresholded output of the preictal neuron of the SoftMax layer (Tensorflow) or inverse of the predictions of the interictal neuron (TrueNorth) is fed into a leaky integrate and fire neuron with variable firing threshold V_{th} and decay time τ (see Fig. S3). The prediction output increments a running variable by one with each preictal prediction. With each interictal prediction, the value of the running variable decays exponentially with the time constant τ. Every time the threshold V_{th} is crossed, a spike is generated and the running variable is set to zero. A spike in the spiking output layer triggers an alarm of the system with a variable alarm length t_a.

To find the optimal values of the parameters V_{th}, τ, and t_a, a grid search is performed and sensitivity (S), time in warning (TiW), and improvement over chance ($IoC = S\text{-}TiW$) are computed (see Section S3 for details).

The optimal operating point, as determined through maximisation of an objective function for month i, $p_{opt,i} = f(V_{th,\ opt},\ \tau_{opt},\ t_{a,\ opt})$ is computed and stored. For the full validation, the objective was to maximise improvement over chance. For the case study, the objective function was modified slightly (see Section S3) to reflect the relative priority given to sensitivity or time in warning. In order to compute the operating point for the next month, p_{i+1}, each parameter was weighed by the number of seizures in the corresponding month, leading to a post-processing layer with memory that gradually adjusted to the updating models:
$p_{i+1} = 1/N_s * \Sigma_{m=(0\ to\ i)}\ (p_{opt,i} * n_{s,i})$,
where N_s is the total number of seizures until the next month, and $n_{s,i}$ is the number of seizures in month i.

In order to ensure we did not use data that is not yet available at the time of prediction, performance was computed using the optimal parameters derived from historical data only.

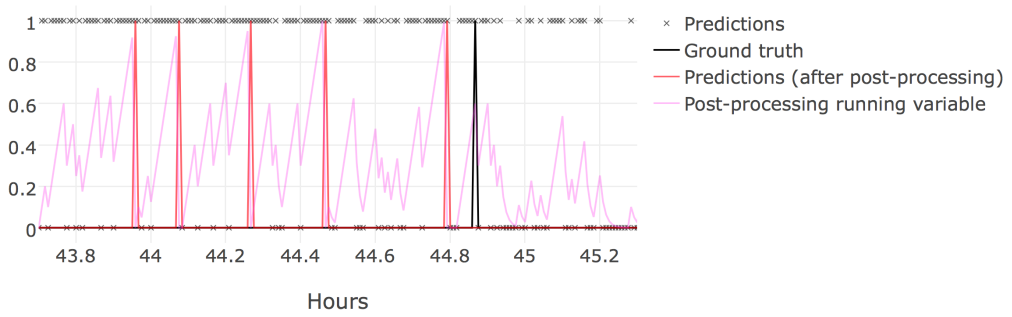

***Fig. S3*: Post-processing layer:** The preictal output neuron (TF) or the inverse interictal output neuron (TN) was passed to an integrate-and-fire neuron. A positive prediction (marked × at position 1) led to a linear increase of the running variable (pink trace). A negative prediction (marked × at position 0) led to an exponential decay of the running variable. Once the running variable crossed a threshold (1 in this example), the post-processing layer predicted an upcoming seizure (red spike). A positive prediction after post-processing then invoked an alarm of a certain alarm length (not displayed). Seizures (black spike) that fell within this alarm window were counted as true positives, seizures outside an alarm window as false negative. Threshold, decay time, and alarm window were optimized as described above.

S3) Performance Measures:

Sensitivity: While any seizure that occurred while the system was in the alarm state were regarded as a true positive (TP), seizures occurring outside alarms were regarded as false negatives (FN). Sensitivity S is defined by:

$S = TP/(TP+FN)$.

Time in warning: Time in warning (*TiW*) is the percentage of time (after removal of any dropout occurring in the data) the system spends in the warning state.

Improvement over chance: For the purpose of evaluating our system, we defined the improvement over chance *IoC* as:

$IoC = S\text{-}TiW$.

Intuitively, this can be explained by imagining a system that spends either 100% or 0% in the warning state, predicting 100%, and 0% of seizures, respectively (see Fig. S4). Sensitivity values for any given time in warning can be derived through interpolation between these two extremes. To find the optimal operating point during post-processing $p_{opt,I}$ for each month, S, *TiW*, and *IoC* were computed. The operating point was determined by selecting the values that satisfied $max(w_{Sens}*S - w_{TiW}*TiW)$, with w_{Sens} and w_{TiW} being the relative weights (importance) assigned to sensitivity and time in warning.

A Month 2 **B Month 3**

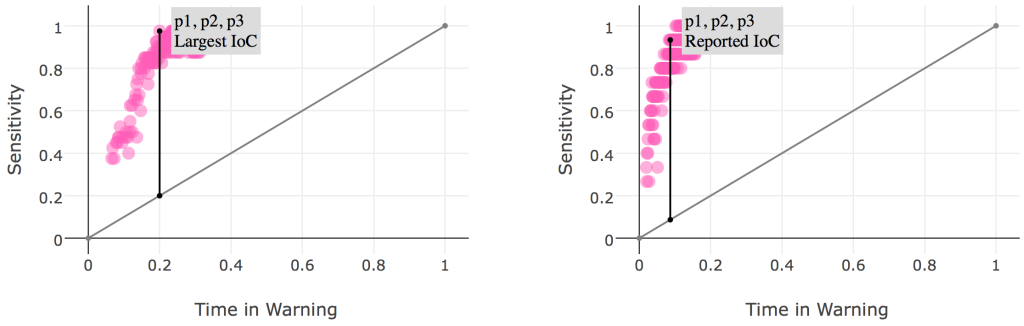

***Fig. S4*: Performance reporting:** The parameter set resulting in the largest improvement over chance in month 2 (and later all previous months) was used to inform which parameter set is used for testing in the following month in order to enable a pseudo-prospective study.

References

Cook, M.J., O'Brien, T.J., Berkovic, S.F., Murphy, M., Morokoff, A., Fabinyi, G., D'Souza, W., Yerra, R., Archer, J., Litewka, L., 2013. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. Lancet Neurol. 12, 563–571.

Esser, S.K., Merolla, P.A., Arthur, J.V., Cassidy, A.S., Appuswamy, R., Andreopoulos, A., Berg, D.J., McKinstry, J.L., Melano, T., Barch, D.R., Nolfo, C. di, Datta, P., Amir, A., Taba, B., Flickner, M.D., Modha, D.S., 2016. Convolutional networks for fast, energy-efficient neuromorphic computing. Proc. Natl. Acad. Sci. 113, 11441–11446. https://doi.org/10.1073/pnas.1604850113

Howbert, J.J., Patterson, E.E., Stead, S.M., Brinkmann, B., Vasoli, V., Crepeau, D., Vite, C.H., Sturges, B., Ruedebusch, V., Mavoori, J., others, 2014. Forecasting seizures in dogs with naturally occurring epilepsy. PloS One 9, e81920.

Karoly, P.J., Freestone, D.R., Boston, R., Grayden, D.B., Himes, D., Leyde, K., Seneviratne, U., Berkovic, S., O'Brien, T., Cook, M.J., 2016. Interictal spikes and epileptic seizures: their relationship and underlying rhythmicity. Brain 139, 1066–1078.

Korshunova, I., Kindermans, P.-J., Degrave, J., Verhoeven, T., Brinkmann, B., Dambre, J., 2017. Towards improved design and evaluation of epileptic seizure predictors. IEEE Trans. Biomed. Eng.

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444.

Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., Sánchez, C.I., 2017. A Survey on Deep Learning in Medical Image Analysis. ArXiv170205747 Cs.

Loddenkemper, T., Lockley, S.W., Kaleyias, J., Kothare, S.V., 2011. Chronobiology of epilepsy: diagnostic and therapeutic implications of chrono-epileptology. J. Clin. Neurophysiol. 28, 146–153.

Miotto, R., Wang, F., Wang, S., Jiang, X., Dudley, J.T., 2017. Deep learning for healthcare: review, opportunities and challenges. Brief. Bioinform. https://doi.org/10.1093/bib/bbx044

Sillay, K.A., Rutecki, P., Cicora, K., Worrell, G., Drazkowski, J., Shih, J.J., Sharan, A.D., Morrell, M.J., Williams, J., Wingeier, B., 2013. Long-Term Measurement of Impedance in Chronically Implanted Depth and Subdural Electrodes During Responsive Neurostimulation in Humans. Brain Stimulat. 6, 718–726. https://doi.org/10.1016/j.brs.2013.02.001

TensorFlow [WWW Document], 2015. TensorFlow. URL https://www.tensorflow.org/ (accessed 6.27.17).

Ung, H., Baldassano, S., Bink, H., 2017. Intracranial EEG Fluctuates Over Months After Implanting Electrodes in Human Brain. J. Neural Eng.