ARTICLE IN PRESS
EBIOM-01291; No of Pages 9
EBioMedicine xxx (2017) xxx–xxx
Contents lists available at ScienceDirect
EBioMedicine
journal homepage: www.ebiomedicine.com

Research Paper

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

Isabell Kiral-Kornek [a], Subhrajit Roy [a], Ewan Nurse [a,b], Benjamin Mashford [a], Philippa Karoly [a,b], Thomas Carroll [a,b], Daniel Payne [a,b], Susmita Saha [a], Steven Baldassano [b], Terence O'Brien [b], David Grayden [c], Mark Cook [b], Dean Freestone [b,*], Stefan Harrer [a,*]

[a] IBM Research – Australia, 204 Lygon Street, 3053 Carlton, VIC, Australia
[b] The University of Melbourne, 3010 Parkville, VIC, Australia
[c] Department of Biomedical Engineering, The University of Melbourne, 3010 Parkville, VIC, Australia

ARTICLE INFO

Article history:
Received 17 August 2017
Received in revised form 16 November 2017
Accepted 28 November 2017
Available online xxxx

Keywords:
Epilepsy
Seizure prediction
Artificial intelligence
Deep neural networks
Mobile medical devices
Precision medicine

ABSTRACT

Background: Seizure prediction can increase independence and allow preventative treatment for patients with epilepsy. We present a proof-of-concept for a seizure prediction system that is accurate, fully automated, patient-specific, and tunable to an individual's needs.
Methods: Intracranial electroencephalography (iEEG) data of ten patients obtained from a seizure advisory system were analyzed as part of a pseudoprospective seizure prediction study. First, a deep learning classifier was trained to distinguish between preictal and interictal signals. Second, classifier performance was tested on held-out iEEG data from all patients and benchmarked against the performance of a random predictor. Third, the prediction system was tuned so sensitivity or time in warning could be prioritized by the patient. Finally, a demonstration of the feasibility of deployment of the prediction system onto an ultra-low power neuromorphic chip for autonomous operation on a wearable device is provided.
Results: The prediction system achieved mean sensitivity of 69% and mean time in warning of 27%, significantly surpassing an equivalent random predictor for all patients by 42%.
Conclusion: This study demonstrates that deep learning in combination with neuromorphic hardware can provide the basis for a wearable, real-time, always-on, patient-specific seizure warning system with low power consumption and reliable long-term performance.

1. Introduction

Epilepsy is singularly unusual among other serious neurological conditions because seizures are brief and infrequent, so that for at least 99% of the time patients are unaffected by seizure activity. Although seizure activity is infrequent, the disability caused by epilepsy can be significant due to the uncertainty around the occurrence and the consequences of the events. The constant uncertainty impairs the quality of life for these individuals. A recent survey confirmed that the majority of patients find this unpredictability to be the most debilitating aspect of epilepsy ("2016 Community Survey," 2016). There is an unmet need for a device that provides a warning when there is an increased risk of a seizure.

A warning system could support new treatment approaches and improve a patient's quality of life. For example, such a system could inform patients' daily routines and help them to avoid dangerous situations when at higher risk of seizure. Tracking fluctuations in seizure likelihood could also be used to titrate therapeutic interventions, reducing the time spent using anti-epileptic drugs or electrical stimulation.

Given the nature of epilepsy, there are undeniably technological and theoretical hurdles to creating a viable warning system for seizures; however, such a system is no longer considered impossible to build (Freestone et al., 2015; Mormann and Andrzejak, 2016). A key development has been the use of long-term electroencephalography (EEG) data. After a long-term clinical trial, Cook et al. were able to demonstrate success of an implantable recording system, seizure prediction algorithm, and handheld patient advisory device (Cook et al., 2013). Using this device, the group recorded a dataset that comprises a total of over 16 years of continuous intracranial electroencephalography (iEEG) recording and thousands of seizures. Cook et al. established the feasibility of seizure prediction in a clinical setting, and provided inspiration for the development of further seizure prediction algorithms (Freestone et al., 2017).

Despite the trial's success, there were also limitations (Elger and Mormann, 2013). While pre-seizure patterns in the iEEG data were extracted in an automated fashion, it was based on a limited and predefined set of features, which may be one reason that prediction was not possible for all patients. After the initial design phase, the algorithm

* Corresponding authors.
E-mail addresses: deanrf@unimelb.edu.au (D. Freestone), sharrer@au.ibm.com (S. Harrer).

https://doi.org/10.1016/j.ebiom.2017.11.032
2352-3964/

was no longer tunable, making the system inflexible to patients' changing preferences regarding false alarm and missed seizure rates.

As preictal patterns are patient specific, no pre-determined set of features will be able to capture all possible preictal signatures. Therefore, standard feature engineering techniques are unsuitable for the creation of a generalizable predictor (Freestone et al., 2017). Instead of restricting the feature space a-priori, all data should be considered potentially relevant for recognizing preictal patterns – a task to which novel computational techniques are uniquely suited.

Deep learning, a machine learning technique, is a powerful computational tool that enables features to be automatically learnt from data (LeCun et al., 2015). Typically, deep learning is used to train a class of algorithms known as deep neural networks to perform specific tasks. The availability of big data has cemented the usefulness of deep learning for a diverse range of problems (LeCun et al., 2015). Applications range from self-driving cars via robotics to novel diagnostic and treatment options in medical imaging, healthcare, and genomics ("FACT SHEET," 2016; Gulshan et al., 2016; Litjens et al., 2017; Ratner, 2015; Stebbins, 2016). Recent open source seizure prediction competitions (Brinkmann et al., 2016; "Melbourne University AES/MathWorks/NIH Seizure Prediction | Kaggle," 2016) have shown that machine learning techniques are able to produce pre-eminent results, suggesting this method may provide a path to clinical translation of seizure prediction devices. However, the best performing algorithms in competitions often require an unrealistic amount of computing resources for a wearable device ("Melbourne University AES/MathWorks/NIH Seizure Prediction | Kaggle," 2016).

For seizure prediction to be implemented in a clinical device, it is necessary for algorithms to run on small, low-power technology. A number of recent advances in computing led to the development of sophisticated deep learning algorithms using ultra-low power chips (Furber, 2016). One example of such a chip is IBM's TrueNorth Neurosynaptic System (Esser et al., 2016; Merolla et al., 2014). TrueNorth is a specialized chip capable of implementing artificial neural networks in hardware and hence it is neuromorphic in nature. It is one of the most power-efficient chips to date, consuming <70 mW power at full chip utilization. The chip's neuromorphic technology allows for the deployment and testing of algorithms that were previously unrealizable in a clinically viable seizure warning system.

Seizure prediction has been established as clinically feasible and highly desirable for patients. In light of promising results (Brinkmann et al., 2016; Cook et al., 2013; Howbert et al., 2014), the development of a practical seizure warning device has been declared a grand challenge in epilepsy management ("Seizure Gauge Challenge," 2017). In this paper, we describe how deep learning and the TrueNorth processor can be leveraged to advance the task of patient-specific seizure prediction. Prediction results were benchmarked using data recorded during the trial undertaken by Cook and colleagues (Cook et al., 2013). The presented results address several limitations of this earlier study, and provide proof-of-concept for a deep learning system for seizure prediction.

2. Materials and Methods

The overall study design is shown in Fig. 1. The iEEG signal is recorded using intracranial electrodes (magenta circles). Annotated iEEG signals are processed by a deep neural network that is trained to distinguish between preictal and interictal signals. The resulting deep learning model is subsequently deployed onto the neuromorphic TrueNorth chip.

2.1. System Design Rationale

The objective of this study is the development, implementation, and evaluation of a clinically relevant seizure prediction system. In order for a system to be valuable to patients while being maintainable by clinicians, we defined the following goals:

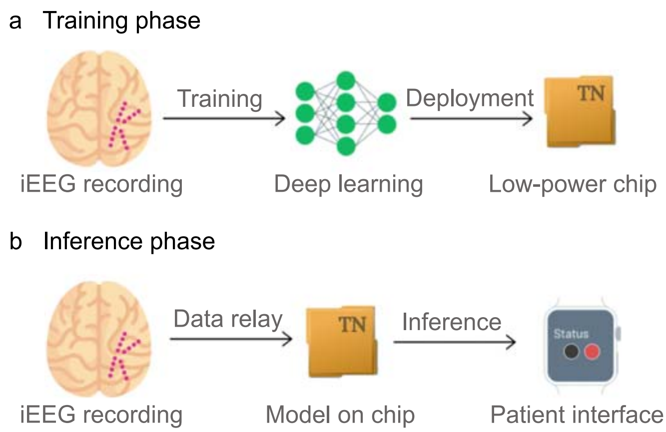

Fig. 1. Concept of seizure advisory system: a) Training phase: iEEG signal is recorded via intracranial electrodes (magenta circles indicate a possible configuration) and recordings are passed on to a deep learning network (green network graph). The model is subsequently deployed onto a TrueNorth chip. b) Inference phase: iEEG signal is recorded via intracranial electrodes (magenta circles) and recordings are passed on to the TrueNorth chip. Prediction of a seizure is indicated to the patient on a wearable device.

G1. The system needs to perform well and reliably across patients.
G2. The system needs to operate autonomously over long periods of time without a requirement for regular maintenance or reconfiguration by an expert.
G3. The system needs to allow for patients to set personal preferences with respect to sensitivity.
G4. The system must run in real-time on a low-power platform.

We addressed performance (G1) and long-term feasibility (G2) using deep learning, a technique that, in contrast to a more traditional feature engineering approach, does not rely on data analysis experts for the monitoring and adaptation of models. Unlike traditional computing systems that learn through instructions or explicit programming, deep learning algorithms learn from examples to automatically discriminate different classes of signals. In the context of a seizure prediction system, this is what allows the algorithm to distinguish between preictal and interictal data segments. By its nature, a system using an artificial neural network cannot only adjust to each individual patient's brain signals, but also to short- and long-term changes in the recording. It further allows for the integration of other patient-specific variables that have been shown to co-vary with seizure likelihood, such as time of day information. Moreover, a deep neural network can automatically learn to discriminate between different classes of data, for example, in this case, preictal and interictal.

Generally, a classification neural network such as the one used in our study will classify the signal on a sample-by-sample basis, leading to potentially very frequent but short alarms. In a real-time system, an additional processing layer is therefore required to balance the sensitivity of the system, number, and duration of alarms. In addition to forming the basis of system optimization, this processing layer also allows for instantaneous tuning of the system's sensitivity by the patient directly (G3).

Adaptation to changes of the signal over time (G2), as for example observed by Cook et al. (2013) were addressed in both processing layers.

Running a neural network classifier in a real-time environment requires specialized hardware (G4). TrueNorth is a highly power-efficient and specialized chip. The network needed to be adapted to run on the TrueNorth chip.

Please cite this article as: Kiral-Kornek, I., et al., Epileptic Seizure Prediction Using Big Data and Deep Learning: Toward a Mobile System, EBioMedicine (2017), https://doi.org/10.1016/j.ebiom.2017.11.032