 DOI:10.3217/978-3-85125-378-8-37

Echo State Networks for Brain Computer Interface Classification

Ryan McCardle[1] and Claire Davies[2]

[1] The University of Auckland, Auckland, New Zealand
ryan.mccardle@auckland.ac.nz
[2] The University of Auckland, Auckland, New Zealand
c.davies@auckland.ac.nz

Abstract

Echo state networks are good candidates for use as classifiers in brain computer interfaces. They can also be used to explore the data to some extent by manipulating their training. In this paper similar results were achieved to those in the Brain Computer Interface Competition IV in classifying four-class movements from magnetoencephalography.

1 Introduction

Artificial neural networks consist of a series of many *neurons* which are connected to each other in a network. Data is fed into the network and propagates from neuron to neuron. Each neuron generates an output from all the combined inputs according to an activation function built into it. How much influence each input to the neuron has is determined by a *weight* placed on each connection. The internal values of a number of special *output neurons* are the final result of the network. Most neural networks are *feed-forward*, meaning the propagation is in one direction (no feedback). A recurrent neural network (RNN) has a memory because the inputs propagate in all directions, blending with earlier inputs and being slowly drowned out by new inputs.

In a RNN the weights of all the connections between the neurons are modified. This creates a very complex non-linear optimization problem which can be very difficult to solve. RNNs and echo state networks (ESNs) are very similar architecturally, but differ in training. In an ESN the weights of all the connections between the neurons are randomly generated at the start of training and set permanently. This creates a large *reservoir* (i.e. reservoir computing) of randomly connected neurons with random weights. Only the weights leading to the output neurons are modified in training, turning a complex non-linear optimization problem into a simple one solvable with linear regression [4].

ESNs are good candidates for brain computer interfaces (BCIs). They perform well in tasks such as speech and handwriting recognition where there is high variability in the input, and a noisy signal. Their *memory* allows recognition of events over time (e.g. slow rise and fall in electrical potential), and events in sequence (e.g. propagation of brain activity). ESNs naturally produce a time series output, allowing continuous BCI output such as arm velocity over time. Crucially, they are easy to train; a dataset can be plugged into an ESN classifier with little pre-processing, and produce a classification level above chance. ESNs have performed well when compared to more commonly used classification techniques in several BCI studies [2, 1, 5].

ESNs are *black box* classifiers i.e. it is difficult or impossible to determine how the input signal is manipulated to produce the output. This is not much of a problem in practical applications, but can make it difficult to optimise the BCI in research. The classifier can also *cheat*. For example, the classifier could exploit irregularities in the data, instead of recorded signals. Non-brain signals could also be exploited e.g. blinking, or bumping EEG cables.

 DOI:10.3217/978-3-85125-378-8-37

We propose that with careful manipulation of the training these problems can mostly be avoided. For example, if the classification accuracy is low when trained only on non-brain signals and high when trained on signals from anatomically relevant brain areas, we can infer that cheating is not occurring. The characteristics of the signal the classifier is using can also be approximated by removing parts of the dataset and evaluating the classification accuracy. The purpose of this paper is to determine the effectiveness of ESNs in a 4-class BCI classification problem, also to investigate this manipulation of the training.

2 Methods

2.1 Training Dataset

Recordings from participants were extracted from one of the datasets distributed as part of the BCI Competition IV [6]. To generate this dataset two right handed subjects moved a joystick 4.5 cm from the centre position in self-chosen cardinal directions. Trial onset was cued with a grey circle. After a variable delay the grey circle would disappear signalling the subject to move. During the trial the subjects were instructed to fixate on a red cross and not to blink.

Magnetoencephalography (MEG) was recorded at 625 Hz, using ten channels (LC21, LC22, LC23, LC31, LC32, LC41, LC42, RC41, ZC01, ZC02) located above the motor areas. The MEG signals were band-pass filtered 0.5 Hz to 100 Hz and re-sampled at 400 Hz. The recordings were then cut into 1 s second trials, with movement onset at 200 ms. The training data set contained 40 trials for each of the four classes per subject, the testing dataset contained 73 trials (true classification unknown during the BCI Competition). As the BCI Competition is over, true labels are available for the testing dataset, and so the training and testing datasets were combined to enlarge the training dataset. The trial order was randomised before training.

2.2 Classifier Training

MEG data was processed with the open-source OrGanic Environment for Reservoir computing (Oger) toolbox that provides a robust implementation of ESNs. To prevent saturation of the network the mean was shifted to zero for each channel. It was then normalised to ± 1, to keep it in the range of the neuron activation function. Data before movement onset was cut because Waldert et al. [6] found that classification was at chance level until then.

ESNs cannot be used directly for classification as the output neurons produce fluctuating values. By representing the movement direction in the training data as four outputs from the ESN (1.0 for the desired class), the channel with the highest output became the classification. Standard deviation was reported to reflect the varying performance of the randomly generated reservoirs. The accuracy of the trained ESN was determined by ten-fold cross validation.

The time taken to train a network increases exponentially with the reservoir size (Figure 1). To achieve sufficient accuracy with a reasonable training time, a reservoir size of 500 neurons was chosen. The remaining parameters of the ESN were optimised by sweeping through a range of values and choosing those with the highest classification accuracy. This gave final values of 0.2 for the leaking rate, 0.1 for the input scaling, and 0.9 for the spectral radius.

3 Results

Figures 1 to 3 show results generated during the optimisation of the classifier, and the final results in 4. Figure 2 shows the accuracy when the network is trained on each electrode indi-

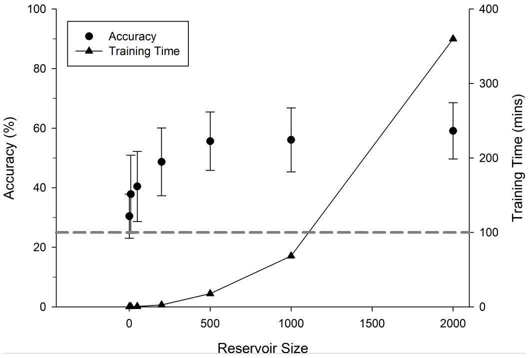

Figure 1: Size of echo state reservoir and its influence on accuracy and training time.

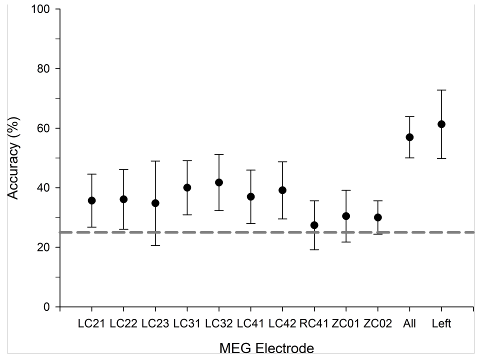

Figure 2: Accuracy for each electrode, all electrodes, and left hemisphere only.

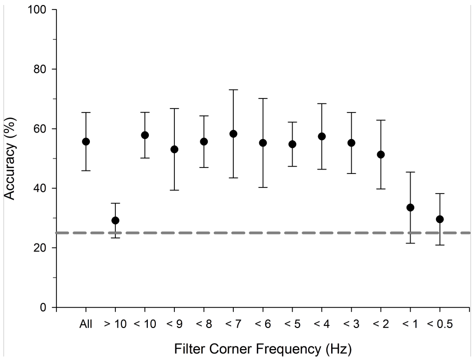

Figure 3: Accuracy of unfiltered data compared to filtering with a 3rd order Butterworth filter at different thresholds.

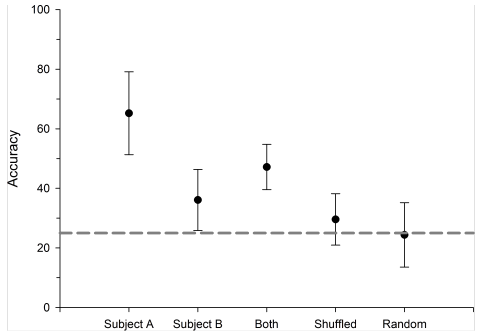

Figure 4: Final accuracy for subjects individually and combined, compared to shuffled and randomised classifiers.

vidually, all the electrodes, and only the electrodes on the left hemisphere. The left electrodes all perform similarly at around 35 % accuracy. The electrodes located along the mid-line, and the right hemisphere (RC41, ZC01, ZC02) perform at close to chance level, suggesting that the classification is not due to EMG from the right wrist. This difference in accuracy also indicates that the classifier is not using EOG artefacts, as these artefacts would effect both hemispheres equally. An accuracy of 57 % is achieved when all electrodes are combined, and with only the electrodes of the left hemisphere accuracy increases to 61 %. A different subset of the electrodes may perform better, but there are too many combinations to test exhaustively.

When the MEG signal is high-pass filtered at 10 Hz the classification accuracy drops off to chance level (Figure 3). When low-pass filtered the accuracy remains the same until it drops off at 2 Hz to 3 Hz. From these drops in accuracy we can infer that the signal being used by the ESN for classification is in the range 1 Hz to 3 Hz. However, this may be because the ESN used is most sensitive to frequencies in this range, and using different parameters when creating the ESN may yield different results.

Using the results described in Figures 1 to 3, only the subjects' left hemisphere electrodes were used, and then bandpass filtered (1 Hz to 3 Hz). The final classification accuracy was compared to shuffled and random classifiers (Figure 4). Shuffling the inputs so they do not match the outputs should make it impossible for the ESN to produce the correct output, and

 DOI:10.3217/978-3-85125-378-8-37

it should function like a random classifier with a mean accuracy of 25 %. In this case the average accuracy for shuffled data is 27 % - indicating that there are some irregularities in the dataset which can be exploited by the ESN. Therefore this is a better benchmark for chance level accuracy, which the reported results still exceed.

4 Discussion

Using the methods described, a classification accuracy of 65 % was achieved for subject 1, and for subject 2 a lower classification accuracy of 36 % (Figure 4). This is higher than, but consistent with, the winning BCI Competition IV entrant who achieved 59.5 % accuracy for subject 1, and 34.3 % for subject 2, with a smaller training dataset and without the benefit of checking their results [3]. The other entrants did not achieve results above chance level.

Manipulating the training dataset showed that brain activity related to the task was localised to the left hemisphere. More training rounds, and a higher electrode density could potentially be used to localise it more precisely. The approximate frequency of the activity was also determined, which shows the potential for using an ESN as a crude way to investigate brain activity with unknown characteristics, or find activity in new frequency bands.

When the ESN was trained with both subjects data combined it achieved an accuracy of 47 %. This is unusual as typically BCIs must be trained for each individual due to variation in brain activity. This may suggest that a single ESN BCI can be trained to work with multiple subjects, or be generalised to work with any subject. However it is impossible to tell without data from more subjects. A result of 47 % is comparable to the accuracy from simply combining the results from both subjects. This means the ESN may simply have been trained to differentiate between subjects and classify them accordingly (this in itself would be interesting).

References

[1] GP Cédric, M Sebag, A Souloumiac, and A Larue. Ensemble learning for brain computer-interface using uncooperative democratic echo state communities. In *Cinquième conférence plénière française de Neurosciences Computationnelles*, 2010.

[2] Aysegul Gunduz, Mustafa C. Ozturk, Justin C. Sanchez, and Jose C. Principe. Echo State Networks for Motor Control of Human ECoG Neuroprosthetics. *2007 3rd International IEEE/EMBS Conference on Neural Engineering*, pages 514–517, May 2007.

[3] Sepideh Hajipour Sardouie and Mohammad Bagher Shamsollahi. Selection of Efficient Features for Discrimination of Hand Movements from MEG Using a BCI Competition IV Data Set. *Frontiers in neuroscience*, 6(April):42, January 2012.

[4] Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks - with an erratum note. Technical report, German National Research Center for Information Technology, St. Augustin, 2001.

[5] Pieter-jan Kindermans, Pieter Buteneers, David Verstraeten, and Benjamin Schrauwen. An Uncued Brain-Computer Interface Using Reservoir Computing. In *Workshop: machine learning for assistive technologies, Proceedings*, page 8. Ghent University, Department of Electronics and information systems, 2010.

[6] Stephan Waldert, Hubert Preissl, Evariste Demandt, Christoph Braun, Niels Birbaumer, Ad Aertsen, and Carsten Mehring. Hand movement direction decoded from MEG and EEG. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(4):1000–8, January 2008.