 DOI:10.3217/978-3-85125-260-6-4

Transfer Learning for Accelerated P300 Speller Classifier Training

K. Colwell[1], C. Throckmorton[1], L. Collins[1], K. Morton[1]
[1]*Duke University, Durham, NC, USA*

Correspondence: K. Colwell, Duke University, 129 Hudson Hall, Box 90291, Durham, NC, USA. E-mail: mkenneth.colwell@duke.edu

Abstract. The P300 Speller provides a means of communication by recognizing evoked potentials on the scalp with a software classifier. However, collecting data for classifier training requires an extensive amount of time and effort and delays the user's ability to use the system. We propose accelerating the training process with transfer learning, which leverages previously-collected data and classifiers to reduce training time and improve classification. We perform clustering on the set of classifiers corresponding to 94 P300 Speller training sessions and evaluate the performance of the resulting set of classifiers on the associated datasets.

Keywords: P300 Speller, Transfer Learning, C lassification

1. Introduction

The P300 Speller is a brain-computer interface that allows users to type on a virtual keyboard using brain signals, requiring no neuromuscular control. Users focus on the target character that they intend to spell from a grid of characters on a computer monitor; characters then flash in a controlled random pattern, and the flashing of the target is assumed to elicit a P300 evoked potential that can be observed via electroencephalogram (EEG). A computer then uses a classifier to distinguish target responses from non-target responses, and thereby determine the target character. Each classifier is trained on a selection of training data that must be collected for each subject before he or she can use the system for free spelling. Although collecting more training data can improve the accuracy of the classifier (and therefore the spelling rate), it also prevents the user from using the system to communicate for a period of time and requires a significant amount of effort from the user. Hence, decreasing the amount of training time makes the system more practical for general use.

Because neural responses vary from subject to subject, classifier training is typically performed as an isolated task, using only the current subject's training data. However, since the basis for operation for the P300 Speller is a particular evoked potential, it is reasonable to hypothesize that out of a collection of subjects, some will exhibit similar target responses. Therefore, we propose to accelerate classifier training using transfer learning, which uses knowledge obtained in a previous task to inform the completion of a new task: in this case, we use previously collected training data or classifiers to inform the training of new classifiers for new data. To allow for variation in target response while still providing generalization, we perform clustering on the linear classifiers corresponding to each available training dataset, then determine whether the resulting clusters describe their cluster members' data well enough to be used as classifiers themselves. Classification and clustering are performed using a mixture of logistic discriminants, and applied to a corpus of 94 individual P300 Speller training sessions. This will demonstrate whether transfer learning is a promising method for further development.

2. Material and Methods

2.1. Data

The data collection is comprised of 94 individual P300 Speller training sessions. Each session was collected using the row/column paradigm described in [Farwell and Donchin, 1988] with a 9x8 grid from abled subjects in a laboratory environment. Sessions contain between 30 and 40 spelled characters, collected between the years 2009 and 2012. EEG data were collected at 256 Hz; each flash triggers an 800ms response window that is decimated to 20Hz; then, of the 32 available channels, the 8-channel set determined in [Krusienski et al., 2008] is used for classification: $\{F_z, C_z, P_z, O_z, P_3, P_4, PO_7, PO_8\}$.

2.3. Methods

The data from each training session are used to train a linear classifier via logistic regression, as this classification method does not enforce weight sparsity, which could hamper the classifiers' ability to cluster. Clustering of classifier weights is then performed using the *k*-means algorithm, for several values of *k* ranging from

 DOI:10.3217/978-3-85125-260-6-4

1 to 30, such that each subject's classifier is assigned a cluster. Each value of *k* is randomly initialized 10 times; each cluster mean is applied as a linear classifier to the datasets assigned to its cluster, and the initialization that produces the highest mean area-under-curve (AUC) is selected for further analysis. Each cluster mean is then applied to all datasets belonging to other clusters, and AUC is calculated for comparison. In addition, an AUC is calculated via cross-validation for each session using only its own data as a baseline.

3. Results

Figure 1 demonstrates that with a sufficient number of clusters, the linear classifiers described by the *k*-means cluster means are effective classifiers for the datasets that would be assigned to them, and are nearly as effective as subject-specific classifiers: using the non-subject-specific cluster means results in an average loss of AUC of less than 0.05 when *k*-means is allowed to determine at least 10 clusters. Cluster membership is increasingly important as the number of clusters increases: cluster-mean classifiers applied to non-member datasets perform markedly worse than subject-specific classifiers.

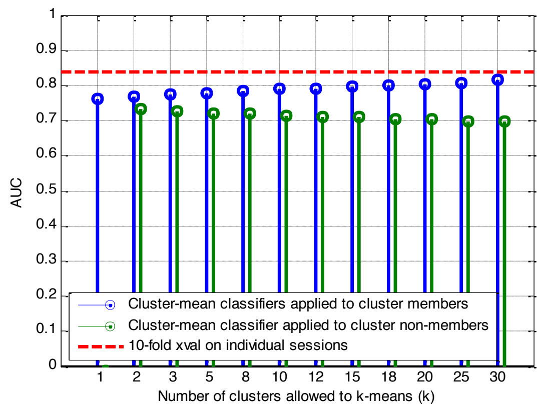

Figure 1. *Mean AUC obtained by applying each cluster mean as a linear classifier to datasets in its cluster (blue) and in all other clusters (green). Compare to mean 10-fold cross-validated logistic regression performed per-session (red).*

4. Discussion

The results in Fig. 1 suggest that separate subjects and sessions share enough similarities to be used for transfer learning, and that further research into using previously-collected training data is merited. Most directly, training time could be reduced for a new participant by collecting only enough training data to accurately select a cluster, then using the cluster-mean classifier. However, more sophisticated methods could use the subject's training data to tune the cluster-mean classifier, potentially improving training speed and accuracy, or incorporate information from all clusters to varying degrees during training. Further, a hierarchical Bayesian model could combine cluster and classifier training on the previously-collected corpus to provide more informative cluster information than *k*-means, while providing a natural method for training classifiers for new subjects.

Acknowledgements

We would like to thank Dr Eric Sellers and David Ryan for contributing data they collected at East Tennessee State University, and Boyla Mainsah for contributing data she collected in Dr Collins' laboratory at Duke University.

This work was funded by NIH/NIDCD (R33 DC010470-03).

References

Farwell L, Donchin E. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. *Electroenceph Clin Neurophysiol*, 70:510-523, 1988.

Krusienski D, Sellers E, McFarland DJ, Vaughan TM, Wolpaw JR. Toward enhanced P300 speller performance. *J Neurosci Meth*, 167:15-21, 2008.