Deep Learning Approaches for P300 Classification in Image Triage: Applications to the NAILS Task

Amelia J. Solon
DCS Corporation
Alexandria, VA USA
asolon@dcscorp.com

Stephen M. Gordon
DCS Corporation
Alexandria, VA USA
sgordon@dcscorp.com

Brent J. Lance
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD USA
brent.j.lance.civ@mail.mil

Vernon J. Lawhern
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD USA
vernon.j.lawhern.civ@mail.mil

ABSTRACT

This paper describes the rationale behind and results of five evaluation submissions to the NAILS (Neurally Augmented Image Labelling Strategies) challenge at the NTCIR-13 conference. Image triage is a time- and resource- intensive process for human labelers. Researchers have identified a potential P300-based BCI solution to alleviate the strain of manual labeling. The NAILS dataset was designed to capture the P300 signal over various image search activities and to act as a benchmark dataset for P300 detection methods. Here we describe approaches that utilize cross- and within-subject training using our in-house Convolutional Neural Network (CNN) EEGNet, and another state-of-the art event-related-potential approach based on xDAWN spatial filtering combined with Information Geometry using Riemannian manifolds. We show improved performance with within-subject training, more data, and modifications to the EEGNet model, and briefly discuss the implications of using certain training data over others.

Team Name

ARL17

Subtask

NAILS – Neurally Augmented Image Labeling Strategies

Keywords

Electroencephalography (EEG), Brain-Computer Interface (BCI), Deep Learning, P300, Event-Related Potential, Convolutional Network, Image Triage.

1. INTRODUCTION

The proliferation of machine learning solutions for image processing, especially those based on deep neural networks (DNNs), is due to in part to their success on a wide range of image recognition tasks, outpacing their human counterparts in both speed and accuracy in limited cases [1]. However, in many tasks, ranging from accurately recognizing the same object types across a variety of different imaging modalities, to segmenting an image into critical objects, to evaluating semantic context in a scene, DNNs still fail to achieve parity with human performance.

Although human image labelers are accurate, they are a slow and expensive solution to difficult large-scale image processing problems. A potentially faster way to leverage the human labelers' strengths is to monitor their visual cortex directly via Brain Computer Interface (BCI). Although BCIs are classically researched for clinical applications [2], researchers have developed BCIs specifically for this image triage problem, in addition to a growing number of applications for healthy users [3,4,5].

These image labeling BCIs are commonly based on the P300 neural response, which is a quick reflexive reaction the brain has to relevant visual stimuli [6]. It is a robust signal that can be elicited via rapid serial visual presentation (RSVP), where a large number of images are rapidly displayed to a user one-at-a-time. It can also be non-invasively recorded via electroencephalogram (EEG). This presents a unique opportunity to detect, at a rapid rate, images that are relevant even in the face of evolving task demands. Due to growing interest in RSVP-BCI for image labeling, the benchmark Neurally Augmented Image Labeling Strategies (NAILS) dataset was made available as part of a collaborative evaluation of P300 detection methods over a range of image search tasks [7].

The detection of any neural event via BCI follows the same general process regardless of the application. The EEG signals must first be acquired from a human via sensors placed on the scalp. These signals are then pre-processed to improve the signal-to-noise ratio (SNR) prior to feature extraction, where distinguishing characteristics are pulled from the signal. These pre-processing and feature extraction methods often depend on the type of expected neural activity. These extracted features are passed off to a machine learning algorithm, often a binary classifier previously trained on similar features to discriminate between the neural activities of interest. The classifier decision is then sent to the application interface, where, if necessary, feedback is given to the user. When analyzing prerecorded data, such as the NAILS competition data, effort is generally focused on signal processing, feature extraction and classification.

In this paper, we elaborate on the methods and motivations behind our five submissions to the NTCIR NAILS evaluation task. Our submissions focus on the feature extraction and classification portions of the BCI process, and are based primarily on our previous work on EEGNet [8], a generalized DNN architecture for EEG-based BCIs. We compare this model to an approach that uses xDAWN spatial filtering together with Riemannian Geometric analyses, an approach that has enjoyed success in several EEG classification challenges [9, 10, 11]. From our results over our five submissions, we discuss the effect of training set composition on final test set performance.

2. MATERIALS AND METHODS

2.1 EEG Data

The NAILS dataset contains 32 channel EEG recordings, sampled at 50Hz, from 10 subjects performing a several 6 Hz RSVP target detection tasks. There are a total of 6 tasks with varying difficulty

(described in [7]) drawn from three popular image databases: Places 365[12], ImageNet[13] and VEDAI[14]. Each task was structured into 9 blocks, each block containing 180 images (9 targets/171 standards), with self-paced breaks in between. In total, each subject was shown 116,640 images, and was instructed to count the occurrence of each target as defined by the task. As electrooculogram (EOG) can vary with class in target detection paradigms, the evaluation organizers removed epochs containing EOG by thresholding the peak-to-peak amplitude at 70 μV and confirmed their removal using ICA and wavelet analysis. For each subject and search task, 15/285 target/non-target trials were held out as the unlabeled test set for the evaluation.

The NAILS dataset provided epochs of [-0.5, 1.5]s around image onset. We further subsampled this epoch to [0, 1]s, consistent with previous literature on P300 classification [15]. No further signal processing of the data was performed.

2.2 Classification Models

2.2.1 EEGNet & EEGNet2

Deep Learning uses multiple layers of artificial neural networks to extract complex features from multidimensional datasets [16]. Convolutional Neural Networks (CNNs) are deep neural networks in which each neuron is a convolutional kernel, similar to those used in classic image processing methods. During learning, the weights in these kernels are optimized with respect to training loss, usually computed with the categorical cross-entropy loss, which is suitable for multi-class classification. A hold-out validation set is generally used to stop model training in order to prevent overfitting to the training set. CNNs typically use a combination of convolutional and pooling layers for feature extraction before a fully connected neural network is used for classification [17, 18]. CNNs have generated large performance improvements in the image processing community, as they are relatively robust to the rotational, translational, and minor distortion effects that accompany multiple instances of the same object.

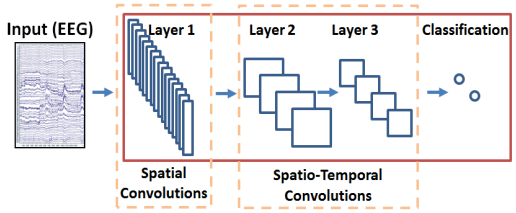

Figure 1. EEGNet architecture: Layer 1 learns spatial convolutions kernels while Layers 2 and 3 learn spatiotemporal convolution kernels.

Previously, we developed EEGNet, a compact convolutional network architecture for EEG-based BCIs [8]. We showed that EEGNet enabled cross-subject transfer performance equal to or better than conventional approaches for several BCI paradigms Our previous work focused mainly on cross-subject (sometimes called subject-independent) classification, as this remains a difficult problem for developing practical BCI systems, especially for oscillatory-based BCIs [19, 20] where the signal-to-noise ratio (SNR) of the feature of interest can vary significantly across subjects. We have also developed an extension of EEGNet, called EEGNet2, which uses Filter Bank Blocks to facilitate localization of frequency-specific activity in the EEG signal.

The basic architectures of EEGNet and EEGNet2 are given in Figures 1 and 2, respectively; here we provide a brief description of the models (the reader is referred to [8] for more details). Given an input EEG trial of size $C \times T$, for C channels and T time points, we first learn a set of spatial filters, implemented as a convolution across the channel dimension (Layer 1), which primarily reduces the dimensionality of the data while potentially improving the SNR of the signal of interest. Layers 2 and 3 of the network use spatio-temporal convolutions to learn correlations in time and across spatial filters. All convolutional layers used Dropout [21] and L_2 regularization to mitigate overfitting. These features are then used in a logistic regression to perform the classification.

EEGNet2 is a variant of the EEGNet architecture, where we replace Layer 1 of EEGNet with multiple Filter Bank Blocks, each of which comprises one temporal convolution of length 20 samples plus two

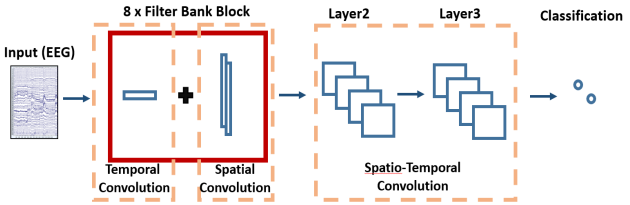

Figure 2. EEGNet2 architecture, replacing the spatial filter layer in EEGNet with a Filter Bank Block.

spatial convolutions. The temporal convolution allows the network to adaptively learn a frequency filter while the spatial convolutions help localize spatial activity along the scalp in frequency-specific bands. We learn 8 of these blocks in total. This operation closely resembles that of the Filter-Bank Common Spatial Pattern (FBCSP) [22], an algorithm specializing in classification in oscillatory-based BCIs. The rest of the network remains the same as EEGNet.

2.2.2 xDAWN Spatial Filtering + Riemannian Geometry

We also investigated using an approach described by Alexandre Barachant, which won the Kaggle BCI Challenge (https://github.com/alexandrebarachant/bci-challenge-ner-2015). This approach uses a combination of xDAWN spatial filtering [23], Information Geometry using Riemannian manifolds [9,10,11], electrode selection, and L_1 feature regularization. This model will be referred to as xDAWN + RG for the rest of the manuscript.

2.3 NAILS Submissions

Here we describe the five submissions we made for the NAILS competition. A summary of the approaches, as well as the data partitioning procedures per submission, can be found in Table 1.

Table 1. Summary of labeled training data partitioning across submissions.

			Labeled Training Data		
		Classification Type	**Train Set**	**Validation Set**	**CV Test Set***
Submission	**1**	Across	5 subj.	4 subj.	1 subj.
	2	Within	1/3	1/3	1/3
	3	Within	1/3	1/3	1/3
	4	Within	2/3	1/3	N/A
	5	Within	4/5	1/5	N/A

Table 2. Maximum BA observed across all folds per subject for each submission. The average validation BA column also has numbers in parentheses, which denote one standard error of the mean. Δ denotes the difference between the Validation BA and the Test BA. The 5th submission, with a Test BA of 0.884, was the winning solution to the NAILS task.

		Train Set	Validation BA										Test BA	Δ
			S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	Average		
Submission	1	Across	0.680	0.819	0.815	0.795	0.842	0.754	0.813	0.789	0.735	0.776	0.782 (0.015)	0.772	0.010
	2	Within	0.779	0.885	0.875	0.887	0.926	0.838	0.902	0.880	0.837	0.856	0.867 (0.013)	0.846	0.021
	3	Within	0.783	0.945	0.858	0.914	0.915	0.854	0.911	0.913	0.853	0.852	0.880 (0.015)	0.853	0.027
	4	Within	0.829	0.951	0.943	0.918	0.967	0.890	0.944	0.918	0.891	0.883	0.913 (0.013)	0.872	0.041
	5	Within	0.829	0.965	0.940	0.932	0.964	0.891	0.954	0.941	0.898	0.901	0.921 (0.013)	**0.884**	0.037

Please note that for the first three submissions, we hold out a CV (cross validation) test set within the labeled training data. In Table 2, the Validation BAs for these three submissions are computed on this CV Test set.

2.3.1 Submission 1: Cross-Subject EEGNet

Our first submission uses EEGNet trained in a cross-subject fashion: We select at random (from the labeled training data), 5 subjects to be the training set, 4 subjects to be the validation set, and 1 subject to be the CV test set. This process was repeated 100 times, creating 100 different "folds" of the data. Note that, due to random sampling effects, not all subjects were the CV test set an equal number of times. The EEGNet model configuration used 8 spatial filters in the first layer, followed by 4 filters of size (2, 12) and (4, 4) for Layers 2 and 3, respectively. We then used an ensemble voting procedure to make the final prediction on the unlabeled test set by finding a threshold such that each subject had on average around 90 predicted targets. The EEGNet model was trained until minimum cross-entropy loss was achieved on the validation set.. Because of the severe class imbalance in the data (only 5% of which are the target class), we trained the EEGNet model with a class weight of 50 to 1 for the target class, essentially assigning higher loss when a target was misclassified.

Though the ensemble is effectively trained on all subjects' training data, the validation BA in Table 2 is computed on the CV test set, which is an entirely separate subject. Thus, the cross-validation approach is representative of a "calibration-free" or subject-independent BCI [19, 20]. This approach represents the situation where knowledge of the test subject may not be known *a priori*; the classification approach used should be able to learn a general representation of the signal from other subjects.

2.3.2 Submission 2: Within-Subject EEGNet2

Since we know the subject IDs for each sample in the test set, it is possible to train a within-subject classification model for this task. Although this model has less data to train on, it can exploit subject-specific EEG dynamics. We also used EEGNet2 here, as this model might better exploit within-subject EEG frequency bands for improved classification. We used filters of size (4, 4) with (2, 2) strides for Layers 2 and 3, respectively. Similar to Submission 1, EEGNet2 was trained by minimizing the cross-entropy loss function on the validation set, using a class weight of 20 to 1 for the target class.

The main challenge here is determining the amount of data that is required to train and validate the EEGNet2 model. For this submission we used a cross-validation procedure consisting of selecting 1/3 of each subjects' data at random to form the training, validation and CV testing sets, respectively. We repeated this procedure approximately 10 times per subject, and used the run which produced the highest cross-validation BA (computed on the CV test set) to make a prediction on that subject's portion of the unlabeled test set.

2.3.3 Submission 3: Within-Subject xDAWN + RG

This submission used the same data partitioning as Submission 2, but instead uses xDAWN + RG for the classification model. We trained 4 xDAWN filters per class, and subsampled the 32 channels to 12 channels using the Riemannian metric (see Section 2.2.2 for more details). This submission was used to compare and contrast the performance of Deep Learning-based models (EEGNet and EEGNet2) with more traditional BCI classification algorithms. The validation BA in Table 2 is from the CV test set with the highest BA. We used this run to make a prediction on that subject's portion of the unlabeled test set.

2.3.4 Submission 4: Within-Subject EEGNet2 with 2/3 Training Data

This submission is the same as that of Submission 2, but we increased the training set size to 2/3 of the available data per subject. We trained the model to minimize cross-entropy loss but logged the validation sets' Balanced Accuracy (BA) at each training iteration. We then saved the model with the highest validation set BA.

2.3.5 Submission 5: Within-Subject EEGNet2 with 4/5 Training Data

This submission is the same as Submission 4, but we further increased the training set size to 4/5 of the available data per subject, and validating on the remaining 1/5.

3. RESULTS

We present our results in Balanced Accuracy (BA), shown in Table 2. For submission 1, we see variation across subjects, with subject 1 performing the worst (0.680 BA) and subject 5 performing the best (0.842 BA). The internal Validation BA achieved a BA of 0.782. When the ensemble of classifiers was applied to the test set, the performance fell 0.010 points to 0.772 BA.

The second and third submissions were used to test potential performance differences between EEGNet2 and xDAWN + RG. We observed minimal differences between the two approaches, both in Validation BA (0.867 and 0.880 for EEGNet2 and xDAWN + RG, respectively) and in Test BA (0.846 and 0.853,

respectively). These submissions also strongly suggested that our cross-validation procedure closely approximated the test set, with relatively minor differences in performance observed (Δ=0.021 and 0.027).

For our fourth submission, we increased the training set size to 2/3 of the available data per subject at the expense of having a CV test set. We observed improved performance in the Validation BA (0.913) and test BA (0.872). Although these are improvements upon our previous approaches, the discrepancies in BA between the validation and test sets grew to 0.041, suggesting that our approach was over-fitting more than in previous approaches.

The fifth and final approach, which further increased the training set size to 4/5 of the available data per subject, improves upon previous attempts, achieving a validation BA of 0.921 and a test BA of 0.8838, and was the winning submission to the NAILS task.

4. DISCUSSION

Although the results increased with each submission, the discrepancy between the validation BA and test BA tended to grow. This could be for a few reasons: first, the validation set size decreased across the submissions in order to maximize training set size. This would have caused the validation predictions to be less stable across learning, yielding a more over-fit stopping point. We also notice that it is greater for the within-subject trained models than for the cross subject trained models. This is potentially due to the overlap between trials that exists in the training set: since the presentation is at 6 Hz, trials have the capacity to overlap up to 0.83s with each other. Since the validation set was randomly sampled from the training data, there could have been trials with significant overlap across the training and validation folds.

When we switched to within-subject training we saw improvements in BA. We did not initially train within subjects because we assumed more data would outweigh the benefits of more specific data when training effective DL models. However, we see that between approaches 1 and 2 that more data is not necessarily better: data that better matches (i.e. within subject data) outperformed the models trained on additional subjects' data. We found in independent tests that various methods of augmenting the within-subject data with other subjects did not necessarily improve results, indicating that there is a tradeoff between data size and data 'content'.

In our previous and current work with our in-house datasets, we've seen that this issue of training set content holds true for within- vs. cross-task training as well: within-task training outperforms cross-task transfer in target detection BCI paradigms [24,25]. Although our previous work uses experiments with greater variations in presentation rate, subject response, and gaze settings, it is possible that tuning to task in the NAILS dataset could generate further improvements. Thus, it should be explored. We did not explore task-specific models in this paper for three reasons: 1) we assumed, as in the within-subjects case, that the benefit of more training data would override the benefit of more specific training data, 2) training within-subject and cross-task would present a challenge for DL model validation and early stopping, and 3) we continued to see improvements in the within-subject cross-task approach.

5. CONCLUSION

In conclusion, we applied an ensemble of EEGNet convolutional neural networks to the NAILS data for feature extraction and classification, resulting in a 0.8838 BA on the test set. The NAILS task is one of the first BCI competitions for image analysis, which is an extremely valuable contribution to the field, and we are grateful for the opportunity to participate. We wish to thank the organizers of the 2017 NAILS competition for their efforts and dedication, which have allowed us to showcase our research.

We believe that the structure of the majority of current BCI competitions, where offline feature extraction and classification approaches are applied to preprocessed data, omits the difficulty of evaluating the approach in an online system. The structure of current competitions privileges computationally intensive methods that can often overfit to specific data sets, and in some cases, can accidentally violate the causal nature of the EEG data (by using future data to make a prediction in the past). Thus, we believe that future BCI competitions should have an online validation component, although we realize that this can be a significant burden for competition organizers.

We note that research on RSVP-based BCI for image analysis stretches back well over a decade, and there has been significant progress in that time. With the advent of image-based DNN approaches (and their ability to perform at or even above human-level labeling performance, see [1]) there is a belief that image labeling BCIs are no longer relevant. However, we believe that this result is only possible when large labeled image datasets are collected and made available [12,13], and thus it remains unclear how to use these approaches when limited labeled data is available and/or when subject-matter expertise is required to analyze the image (i.e.: satellite and medical imagery). We believe that a multi-agent human-computer interface (HCI) for image analysis can potentially be used in these scenarios [26].

6. ACKNOWLEDGEMENTS

This project was sponsored by the U.S. Army Research Laboratory under CAST 076910227001, ARL-H70-HR52, ARL-74A-HRCYB and through Cooperative Agreement Number W911NF-10-2-0022. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

7. REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition." CoRR, vol. abs/1512.03385, 2015. [Online]. Available: http://arxiv.org/abs/1512.03385

[2] J. D. R. Millán, R. Rupp, G. R. Müller-Putz, R. Murray-Smith, C. Giugliemma, M. Tangermann, C. Vidaurre et al. "Combining brain–computer interfaces and assistive technologies: state-of-the-art and challenges." Frontiers in neuroscience 4 (2010).

[3] J. van Erp, F. Lotte, and M. Tangermann, "Brain-Computer Interfaces: Beyond Medical Applications." Computer, vol. 45, no. 4, pp. 26–34, Apr. 2012

[4] S. Saproo, J. Faller, V. Shih, P. Sajda, N. R. Waytowich, A. Bohannon, V. J. Lawhern, B. J. Lance, and D. Jangraw, "Cortically coupled computing: A new paradigm for synergistic human-machine interaction." Computer, vol. 49, no. 9, pp. 60–68, Sept 2016.

[5] P. Sajda, E. Pohlmeyer, E. Wang, L.C. Parra, C. Cristoforou, J. Dmochowski, B. Hanna, C. Bahlmann, M.K. Singh and

S.F. Chang. “In a blink of an eye and a switch of a transistor: cortically coupled computer vision.” Proceedings of the IEEE, vol. 98(3), 2010.

[6] J. Polich, “Updating p300: An integrative theory of P3a and P3b.” Clinical Neurophysiology, vol. 118, no. 10, pp. 2128 – 2148, 2007.

[7] G. Healy, T. Ward, C. Gurrin and A.F. Smeaton. Overview of NTCIR-13 NAILS task. In Proceedings of the 13th NTCIR Conference on Evaluation of Information Access Technologies, NTCIR-13, Tokyo, Japan, 5-8 December 2017.

[8] V.J. Lawhern, A.J. Solon, N.R. Waytowich, S.M. Gordon, C.P. Hung, and B.J. Lance. “EEGNet: A Compact Convolutional Network for EEG-based Brain-Computer Interfaces.” arXiv:1611.08024. [Online]. Available: http://arxiv.org/abs/1611.08024

[9] B. Rivet, A. Souloumiac, V. Attina, and G. Gibert, “xDAWN algorithm to enhance evoked potentials: Application to brain-computer interface.” IEEE Transactions on Biomedical Engineering, vol. 56, no. 8, pp. 2035–2043, Aug 2009

[10] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, “Multiclass Brain-Computer Interface Classification by Riemannian Geometry.” IEEE Transactions on Biomedical Engineering, vol. 59, no. 4, pp. 920–928, Apr. 2012.

[11] A. Barachant and M. Congedo, “A Plug & Play P300 BCI Using Information Geometry.” arXiv:1409.0107 [cs, stat], Aug. 2014, arXiv: 1409.0107. [Online]. Available: http://arxiv.org/abs/1409.0107

[12] B. Zhou, A. Khosla, A. Lapedriza, A. Torralba and A. Oliva. “Places: An Image Database for Deep Scene Understanding”. CoRR abs/1610.02055 (2016). http://arxiv.org/abs/1610.02055

[13] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV) 115, 3 (2015), 211–252.

[14] S. Razakarivony and F. Jurie. “Vehicle detection in aerial imagery: A small target detection benchmark”. Journal of Visual Communication and Image Representation (34), 187 – 203, 2016.

[15] A. R. Marathe, V. J. Lawhern, D. Wu, D. Slayback, and B. J. Lance. "Improved neural signal classification in a rapid serial visual presentation task using active learning." IEEE Transactions on Neural Systems and Rehabilitation Engineering 24, no. 3, 2016.

[16] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, pp. 436–444, 2015.

[17] A. Krizhevsky, I. Sutskever and G. Hinton. “ImageNet Classification with Deep Convolutional Neural Networks.” Advances in Neural Information Processing Systems, pp. 1097 – 1105, 2012.

[18] K Simonyan and A. Zisserman. “Very Deep Convolutional Networks for Large-Scale Image Recognition.” arXiv preprint arXiv:1409.1556, 2014.

[19] F. Lotte, “Signal Processing Approaches to Minimize or Suppress Calibration Time in Oscillatory Activity-Based Brain-Computer Interfaces,” Proceedings of the IEEE, vol. 103, no. 6, pp. 871–890, Jun. 2015

[20] N. R. Waytowich, V. J. Lawhern, A. W. Bohannon, K. R. Ball, and B. J. Lance. "Spectral transfer learning using information geometry for a user-independent brain-computer interface." Frontiers in neuroscience (10), 2016.

[21] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov. "Dropout: a simple way to prevent neural networks from overfitting." *Journal of machine learning research* 15.1 (2014): 1929-1958

[22] K. K. Ang, Z. Y. Chin, C. Wang, C. Guan, and H. Zhang, “Filter bank common spatial pattern algorithm on bci competition iv datasets 2a and 2b,” Frontiers in Neuroscience, vol. 6, p. 39, 2012. [Online]. Available: http://journal.frontiersin.org/article/10.3389/fnins.2012.00039

[23] B. Rivet, A. Souloumiac, V. Attina, and G. Gibert. "xDAWN algorithm to enhance evoked potentials: application to brain–computer interface." IEEE Transactions on Biomedical Engineering 56, no. 8, 2009.

[24] S. M. Gordon, M. Jaswa, A. J. Solon, and V. J. Lawhern. "Real World BCI: Cross-Domain Learning and Practical Applications." In Proceedings of the 2017 ACM Workshop on An Application-oriented Approach to BCI out of the laboratory, pp. 25-28. ACM, 2017.

[25] A.J. Solon, S.M. Gordon, V.J. Lawhern and B.J. Lance. “A Generalized Deep Learning Framework for Cross-Domain Learning in Brain-Computer Interfaces.” The First Biannual Neuroadaptive Technology Conference, pp. 34-35, 2017. http://neuroadaptive.org/files/NAT17_Berlin_Conference_Programme.pdf#page=38

[26] A.W. Bohannon, N.R. Waytowich, V.J. Lawhern, B.M. Sadler and B.J. Lance. “Collaborative Image Triage with Humans and Computer Vision.” IEEE International Conference on Systems, Man and Cybernetics (SMC), pp. 004046-004051, October 2016.