Real-time Naive Learning of Neural Correlates in ECoG Electrophysiology

Zachary V. Freudenburg, Nicolas F. Ramsey, Mark Wronkiewicz, William D. Smart, Robert Pless, and Eric C. Leuthardt

***Abstract*—Brain Computer Interfaces (BCI) seek to measure brain signals in order to control computational or robotic devices, with important applications to motor disability. Electrocorticography (ECoG) is an emerging signal platform for long term implantation of a brain signal recording device, but current approaches rely heavily on screening tasks and trained technicians to find and specify repeatable features in the ECoG signal. Here we explore unsupervised approaches to reducing the ECoG signal stream into a few components that correspond most directly to neural patterns that correlate to subject task performance (neural correlates). We report on the development of a real-time feedback system we call the "Brain Mirror" which is based on the real time, incremental learning of a Deep Belief Network. On real patient data, we demonstrate that the components learned online with Deep Belief Networks have higher correlations with neural patterns than PCA.**

***Index Terms*—Brain Computer Interface, Deep Belief Networks, Electrocortiogram, neural correlates, Unsupervised Learning.**

I. Introduction

Brain Computer Interface (BCI) device control is an exciting and rapidly expanding field. BCI control has been demonstrated using many types of brain signal features recorded with several different recording technologies. In this paper we consider the analysis of electrocorticography (ECoG) signals. This recording technology captures higher frequency and more varied signals than the better known electroencephalography (EEG), and samples a larger spatial area than single cell recordings. Significant evidence now exists that supports the idea that information in the brain is represented as distributed oscillatory patterns [1-2] and ECoG signals are particularly well suited to record these types of patterns.

Compared to EEG recordings, ECoG captures a richer representation of the brain signals. Although this supports more robust and complex control, it also requires the analysis of data from more electrodes and across a much larger frequency range than EEG.

As a result, finding the best features of the brain signal to use as BCI control features becomes a significantly more complex problem. Traditional BCI methods rely on what the experimenter thinks should work for the patient. The real-time feedback provided by the Brain Mirror system allows an individual to explore which brain activity patterns they can control and how they can control them.

This paper explores methods to support online, real time data summary and analysis tools, to rapidly identify multiple user controllable brain patterns. In particular, we compare online PCA and Deep Belief Net approaches to develop patient specific representations of brain activity, and we characterize how well they correlate with cued behaviors, which serve as proxies for controllable brain states.

A. Background: Current Status of ECoG BCI

Current clinical use of BCI technology relies on explicit screening procedures, which ask the patients to systematically move different body parts repeatedly. The recorded data is then evaluated offline to see which motions had repeatable and unique neural signatures. This approach has led to successful one dimensional control [3] including playing a simplified game of space invaders [4] two dimensional control based on overt and imagined hand and tongue movements [5] and motions of either hand [6], and one patient who was able to control a computer mouse [7] with both up-down, left-right and a click. BCI subjects show improvement over time when engaged in closed loop control with feedback even when the control features are held constant[3, 7-10]. This indicates that they are altering their neural patterns to fit the signal detection system. In fact, subjects often report just thinking about ``moving the object on the screen", instead of trying to repeat the screening task that led to the repeatable signal.

The key limitation of current systems is the offline training process, where data from screening tasks is analyzed (overnight, in clinical settings) to find features that correlate with subject motions. The key to quickly discovering neural correlates is facilitating an informative feedback pathway to the BCI subject, which requires summarizing the high-dimensional brain signals into a small set of patterns. To this end we have developed a real-time ECoG signal

Manuscript received March 29, 2011. This work was supported in part by NSF Grant MDII2008103, the Children's Discovery Institute at St. Louis Children's Hospital and the Washington University School of Medicine

Zachary V. Freudenburg is with the Department of Computer Science & Engineering Washington University in St. Louis, St. Louis, Missouri, USA, and also with the UMC Utrecht, Utrecht, The Netherlands, voges78@gmail.com.

Nicolas F. Ramsey is with the Rudolf Magnus Institute, UMC Utrecht, Utrecht, The Netherlands, n.ramsey@umcutrecht.nl.

Mark Wronkiewicz is at the Department of Biomedical Engineering at Washing University in St. Louis, St. Louis, Missouri, USA, mdw4@cec.wustl.edu.

William D. Smart **and** Robert Pless are with the Department of Computer Science & Engineering at Washington University in St. Louis, St. Louis, Missouri, USA. wds@cse.wustl.edu, pless@cse.wustl.edu.

Eric C. Leuthardt is with the departments of Neurosurgery and Biomedical Engineering at Washington University in St. Louis, St. Louis, Missouri, USA, LeuthardtE@nsurg.wustl.edu.

processing and feedback system, which we have named 'Brain Mirror'.

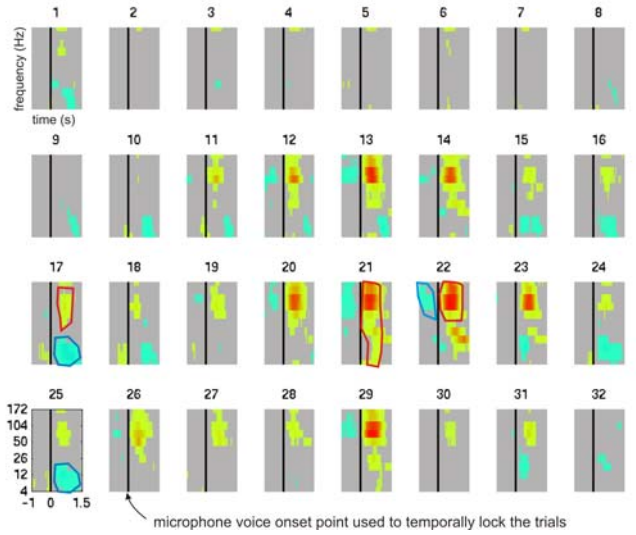

Fig. 1. Example of Time-locked Response (TLR) based Spectral Control Features (SCFs) for the 'oe'-'ah' activity in the FineSp task. Each plot represents the TLR for 1 of the 32 electrode channels for each frequency (y-axis) across time (x-axis). The non-parametrically statistic TLR points are in grey, the positive response. E values and in red, and the negative response value are in blue. Three positive SCFs are circled in red and three negative SCFs are circled in blue.

B. *Algorithmic and Experimental Contributions*

The best linear basis for reconstructing a signal drawn from a given distribution (measured as sum of squared reconstruction error) is given through Principal Component Analysis. But patterns of neural signals at the scale measured by ECoG arise from non-linear interactions of large groups of neurons.

Thus, our contribution in this work is to explore the representation of patterns in ECoG recordings using Deep Belief Networks (DBNets), a compelling class of algorithms capable of quickly learning a non-linear basis for patterns in high dimensional data sets. We evaluate the effectiveness of DBNets within the Brain Mirror system running with ECoG data from 4 subjects who performed a range of tasks. We quantify the relationship between the task a subject was performing and the signal patterns learned by the DBNets-based Brain Mirror system, and show that it is more effective than a linear signal decomposition approach based on PCA, across a wide variety of experimental conditions and for neural signals corresponding to many tasks.

II. Methods

A. *Human Subjects and Data Recording*

This work is based on Electrocorticography (ECoG) recordings from 4 intractable epilepsy patients. These patients ranged in age from 12 to 58, had implants of ECoG electrode grids of varying electrode sizes and spacing that covered a range of cortical locations. Subjects 1 and 3 had an 8x8 and 6x8 grids over left frontal and left tempero-parietal regions, respectively. Electrodes had 2.3mm diameter exposed contact surface and had 1cm inter-electrode spacing (center-to-center). Subject 2 was implanted with a 4x4 microgrid with 1mm spacing and 12 recording electrodes (the 4 corner electrodes were not used) with 75 micron exposed surfaces over the left hand motor cortex. Subject 4 was implanted with a high-density 4x8 electrode grid with 3mm spacing and with a 1.3mm exposed surface over the left mouth motor cortex. All subjects willingly participated in a range of research tasks. The standard procedure of cortical stimulation mapping [11] was followed to locate electrodes that covered functional speech areas for Subject 1.

Biosignal amplifiers were used to record the changing field potentials at a sampling frequency of 1.2 kHz with 24-bit resolution. BCI2000 software [12] was used to synchronize task cue stimulus presentation with the recorded ECoG signal during initial research task performance.

B. *Performed Research Tasks*

Data was recorded for five different tasks from four different patients. These tasks asked for different behaviors from each subject and use different approaches to labeling the behavior at each time point.

1. Subject 1 participated in a hear-and-repeat task (referred to as 'HS' in the rest of the text). The audio cues and recorded microphone signal were used to define activity labels --- the time points when the subject was performing the hearing and speaking activities of the task, respectively.
2. Subject 1 also performed an overt hand and tongue movement motor task (Motor). The timing of when the subject was visually cued to move is used as labels for the Motor task.
3. Subject 2 was performed a task involving three fine hand motor activities; moving a finger, a thumb, or both in a pinching motion (FineM). A 5DT hand glove recorded finger movements to use as activity labels.
4. Subject 3 participated in a read and repeat phonemes task (FineSp) consisting of two categories of spoken phonemes that involve distinct mouth movements. The recorded microphone signal was also used to define voice onset times used as task activity labels.
5. Subject 4 participated in a task (Multi) using the SIGFRIED cortical mapping system [13], in which the subject was visually cued to repeat overt hand movements, tongue movements or speak for longer time periods of 15 seconds for a total of 4 times. The activity cue time points were used as labels.

All tasks included periods of rest (Inter-trial-intervals (ITI)) between successive task activities. All ITI periods were visually cued and the subjects were asked to relax and remain still and silent, but continue to focus on the screen.

C. *Standard spectral control feature selection*

The standard method for choosing control features for BCI control forms the basis of our analysis of the Brain Mirror performance. There are multiple variations of the same basic approach used in choosing control features from spectral ECoG data. Here we define the methods that we used in our analysis in detail. All common control feature selection methods involve the use of prerecorded data from the performance of tasks with an activity-ITI-activity structure. The ECoG data is then converted to the frequency domain. Spectral features that have significant time-locked responses (TLR) to the task activity labels are subsequently used as BCI control features. In this paper spectral control feature

selection was performed on the data sets from the five research tasks described above. The method used in our analysis involved a three step process.

First the TLRs for each task activity were computed. The spectral responses of a range of frequency bands (spectral features) for each electrode channel (see section III) were logged for the duration of the tasks. The task activity labels where then used to align time segments of the spectral data from all repetitions (trials) of the same task activity. Next, the TLR for each spectral feature (SF) was computed by taking the average over the aligned time segments for each task activity.

Second, the time points of the TLRs for each SF that showed significant increases or decreases in amplitude induced by the task activities were used to define candidate channels and groups of SFs to be used as control features. Following the method of non-parametric statistics described by Maris and Oostenveld [14], each TLR value was compared to the distribution of values given by choosing the maximum value of 100 pseudo-TLR formed from randomly choosing N temporal points in the task data to time lock to. N was set to the number of trails for each of the task activities. All values in the actual TLR that were > 95% of the 1000 maximum pseudo-TLR values were considered to be significant. An example of the significant components of the TLRs for all SFs of all electrode channels for the 'oe-aa' activity of subject 4 performing the FineSp task is given in Fig. 1.

Third, groups of neighboring SFs from the same channel with conjoining spans of significant TLR components of the same sign were chosen as Spectral Control Features (SCFs). Six groups of significant TLR components, which define six SCFs, are circled in Fig. 1. For example the SCF chosen for the red circle in plot 17 (electrode 17) of Fig. 1 would be the frequency band 50 to 148 Hz with a positive weight. Likewise, the band 4 to 18 H with a negative weight could also be chosen from electrode 17. Electrode 22 gives an example of a frequency band (50 to 172 Hz) that could be chosen as a positive weighted SCF and a negatively weighted SCF. Generally, the functional importance of the SCF is based on both the frequency and temporal span of their significant TLR components, and the best SCFs are chosen as BCI control features. In this work, the SCFs were ordered according to the absolute value of the total weight in frequency and time and the top 20 SCFs of each task were chosen from the sets of SCFs from each task activity.

D. Deep Belief Network Implementation

Deep Belief Networks (DBNets) [15] are an adaptive, multilayer artificial neural network structure used to transform high-dimensional data to low-dimensional codes. The algorithm for training a DBNet treats each layer as a Restricted Boltzmann Machine with a set of logistic visible units V and a set of stochastic binary hidden units H with a full set of symmetrically weighted connections W. Given the activation states $v_i \in V$, $h_j \in H$, and $w_{ij} \in W$ the energy of the RBM is given by

$$E(V,H) = -\sum_{i \in V} b_i v_i - \sum_{j \in H} b_j h_j - \sum_{i \in V, j \in H} v_i h_j w_{ij}$$

where b_i and b_j are the biases of the visible and hidden units respectively[15]. The weight matrix W and biases that minimizes this energy for a given set of observed input activations of the visible units can be efficiently approximated using the contrastive divergence algorithm[16]. The base layer Restricted Boltzmann Machine in the network uses the input data as its visible units and the hidden units of each preceding layer are the visible units of all higher layer Restricted Boltzmann Machines. The weights and biases that minimize the energy of the entire network for a given set of observations can be minimized by treating each Restricted Boltzmann Machine layer separately.

In our implementation, the input spectral ECoG data is presented to the DBNet as an input vector (of dimension number of frequencies components x number of electrode channels), where each element of the vector is the power of the signal at a given time, for each frequency and electrode.

The activations of the top layer hidden units will be referred to as the responses of the DBNet encoded basis functions. The structure of these functions in the observed ECoG spectra feature space can be calculated by fixing each top layer hidden binary unit to 1 with the rest 0 and projecting down the network to get the corresponding activations of the logistic visible units of the bottom network layer.

III. The Brain Mirror System

The Brain Mirror system is built on top of the BCI2000 [12] system. This is a common platform used for systematic evaluations and comparisons of different brain signals, recording methods, and processing algorithms.

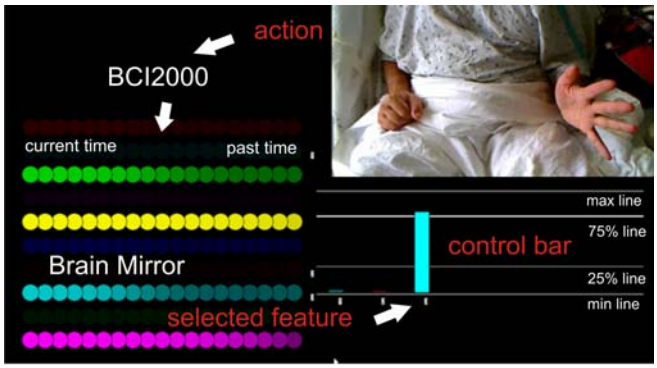

Fig. 2. Brain Mirror Screen Shot. Three control features have been selected as control features to drive the control bars. The third feature is currently activated by the subjects hand movement.

The Brain Mirror system (see Fig. 2) replaces the standard BCI2000 control feedback and instead shows the activation (probability) of the top layer units of the DBNet. The brain mirror summarizes the recent time history by showing features (the response to each top-layer DBNet unit) as either active "colorful", or inactive "dark", with the current time step on the left and the recent history going to the right. If the patient finds they have good conscious control over the brightness of one of the features, they, or the experimenter, can select that feature at which point it will drive a control bar on the right on the screen. When the control feature can reliably be moved above 75% of maximum strength or below 25% of maximum strength the subject is deemed to have demonstrated control, and the chosen feature would be considered suitable for controlling an external device.

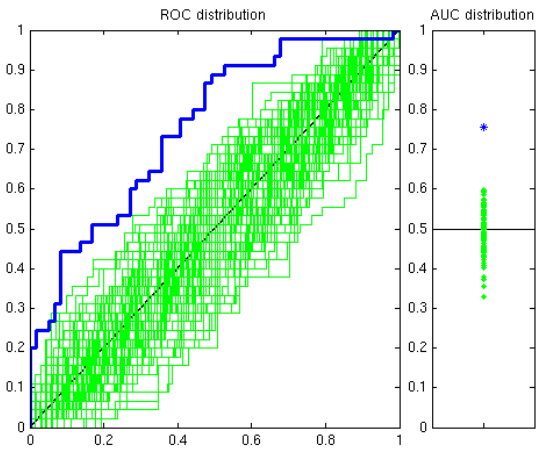

Fig. 3. An example of Receiver-Operator Curve (ROC) Non-parametric statistics. Left: ROC for the 'oe'–'ah' condition vs. ITI for the FineSp task(blue line) shown with 100 ROCs(green lines) calculated from the response of the 'oe'–'ah' control feature using and equal number of random time-locked trials as there were actual 'oe'–'ah' trials. Right: the corresponding Area Under the ROC partition distribution (green points) and 'oe'–'ah' control feature value (blue point).

Implementation details:

Because the Brain Mirror system is designed to give direct feedback from neural correlate patterns over all recorded channels and a broad spectral range, 11 frequency bands of 2 Hz width are defined for each channel. The results in this paper used the frequencies 4, 8, 12, 18, 26, 38, 50, 82, 104, 148, and 172 Hz based of the range of common ECoG features and noise bands. The BCI2000 system is configured to run for 10s to calculate a gain and mean for each frequency band and then output the normalized amplitude response for each band for all future time points. The normalized amplitude responses represent the spectral features (SFs) discussed in section II.C. After the initial 10s, the Brain Mirror system feeds these SFs into the DBNet implementation.

The Brain Mirror system used a 3 layer DBNet. The number of hidden units (NUM_HIDDEN) and output units (NUM_OUTPUT) parameter is defined by the user. The first Restricted Boltzmann Machine layer of the DBNet has as many scalar visible units as frequency features * recorded channels and NUM_HIDDEN binary hidden units. The second Restricted Boltzmann Machine layer is symmetric with NUM_HIDDEN visible and hidden binary units. The third Restricted Boltzmann Machine layer has NUM_HIDDEN visible binary units and NUM_OUTPUT hidden binary units.

The DBNet layers are updated each sample block (every 32 samples with a sampling rate of 1200 Hz) on a set of previous classifier inputs. The number of previous classifier inputs corresponds to the number of circles in each output row of the feedback (ie: 20 in the Fig. 2 example and the analysis in this paper) which can also be set by the user. Hence, the DBNet layers are trained on a small sliding window through the spectral feature data as apposed to the traditional method of exhaustively training DBNets on a representative sampling of the entire data set. Such training will only allow the top layer of the DBNet to converge to encode a pattern in the spectral feature space if it shows some degree of consistency over time. The learning rates for the scalar (s_lr) and binary (b_lr) nodes and the learning momentum (m) are also user defined parameters and will effect the need level of temporal consistency a spectral pattern must have to become encoded in the DBNet. For all Brain Mirror results in this paper, the parameters were as follows; NUM_HIDDEN 100, NUM_OUTPUT 20, s_lr 0.001, b_lr 0.01, and m 0.5. These parameters define a relatively small DBNet, with relatively slow learning parameters, which prove to be able to robustly encode interesting patterns in a diverse sampling of ECoG data.

IV. Analysis

A. Evaluating Brain Mirror neural correlates

We now present a systematic comparison that attempts to capture how well the DBNet is able to discover neural correlates --- brain signals that correlate with conscious activities of the subject. For each of the five research task data sets, we perform the following analysis:

1. Compute the spectral features from the ECoG data stream and incrementally train the DBNet.
2. Log the current responses of the DBNet output nodes, and all internal weights and biases every 37.5 ms.
3. Evaluate each DBNet output node as a potential neural correlate (see section IV.A.1)).
4. Consider the sets of significant neural correlates and evaluate their performance as classifiers of activity of the subject (see section IV.A.2)).

We compare these results of the live Brain Mirror system (BM), with 4 alternatives to the Deep Belief Net representation.

1. The final state of the DBNet to capture its performance if trained over a long time (finalBM)
2. Online PCA [16] (onlinePCA), because it is the natural alternative representation to DBNets.
3. Final state PCA (finalPCA), to capture online PCA performance when trained over long data.
4. Batch PCA (batchPCA), similar to final state PCA, but is independent of the order of the ECoG signals.

These 4 alternatives, along with the top SCFs and the BM features, give 6 methods to find feature representations that are neural correlates and distinguish user activities. For all six methods 20 features were considered.

1) Defining neural correlate sets

To compare the six representations, we first computed TLRs for the features and again use non-parametric statistics, as described in section II.C, to find neural correlates that were significantly and positively modulated by the task. Thus, neural correlates are considered to be features with a significant positive TRL with regard to at least one of the task activities.

2) Defining independent control features

We then consider how well each significant neural correlate serves as a classifier of that activity vs. the inter-trial interval, that activity vs. other activities, and that activity vs. all other time points. Different thresholds of the feature give different classifiers, and the receiver-operator curve summarizes the performance for all thresholds. We use the

standard "area under the receiver-operator-curve" (AUROC), as a summary statistic for how well each significant neural correlate would perform as a classifier.

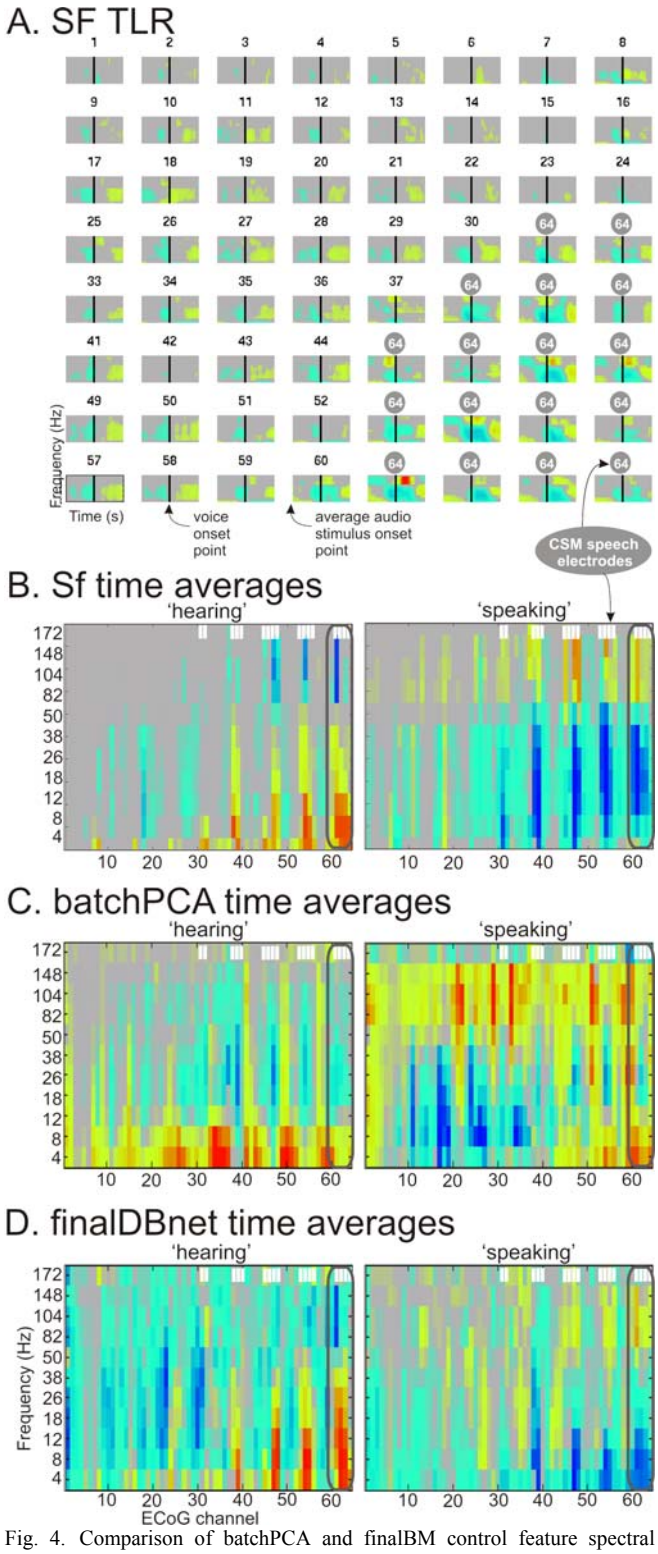

Fig. 4. Comparison of batchPCA and finalBM control feature spectral patterns. A. SF TLRs including the 'hearing' and 'speaking' activities of the HS task. B. The results of averaging the TRLs of each SF over the time period in A. C. The temporally averaged batchPCA control features. D. The temporally average finalBM control features. The electrodes indicated CSM as speech arrest electrodes are indicated with white numbers in A and white boxes in B-D.

We first compute the response value for each significant neural correlate for each activity trial. For each trial the dot product of the feature response over the trial with the TLR pattern for that feature is taken as the response of that feature for that trail. Then we compute the AUROC for the responses over all trials. We again use non-parametric statistics to define AUROC values that are significant. To do this, the per trial neural correlate response values were randomly partitioned 100 times and 100 pseudo-AUROC values were computed based on these partitions. An example of the receiver-operator curves for the 100 random partitions and their corresponding pseudo-AUROC values, along with the actual receiver-operator curve and AUROC value for the 'oe'-'ah' activity of the FineSp task, is given in Fig. 3.

3) Temporal Evolution of Control Features

In addition to evaluating the performance over the entire data set, a third evaluation of the online representations (BM, onlinePCA) was done to characterize the amount of time they needed to find useful control features in the data. To make a fair comparison, we specify the final segment of the ECoG data stream as test data. We use the fact that we logged that weights and biases of the DBNet and the components of the online PCA during the training --- giving us the state of the DBNet and online principal components for different amounts of training time.

For each logged state, we evaluate the testing data and choose the single best control that best discriminates the best activity from the other activities and from the inter-trial intervals for each task. The best activity for each task was chosen as the one with the most significant control features for the Brain Mirror and onlinePCA methods in the above analysis. For this control feature we again use the same two step method as above and use the AUROC to characterize the effectiveness of the best control feature available after a given amount of training time.

V. Results

The results show that the Brain Mirror system captures neural correlates that can be used to define control features that corresponding to the different activities the subjects performed. Furthermore, the DBNet representation used in the Brain Mirror outperformed the natural alternative representations. This is especially impressive considering that the SCF representation used labels during training and the batchPCA was trained on more data. The results of the three analysis methods defined in sections IV.A 1-3 are discussed in more detail below.

A. Naive learning of neural correlates

For all tasks the BM system naively learned patterns in the ECoG spectral data that reflected functionally meaningful neurophysiologic phenomenon induced by brain activity. In fact, the DBNet-based representations were the only ones that learned at least one neural correlate for every activity preformed by the subjects in the five tasks (see Fig. 5).

Further illustration of the ability of the online DBNet algorithm to learn relevant neural patterns is seen when the control features of the batchPCA and finalBM representations for the HS task are compared to the results of the cortical stimulation mapping done with subject 1.

Fig. 4A shows the TLRs of the SFs voice onset time (indicated by the black vertical line in each subplot) of the HS task. The electrodes at locations indicated as essential to speech production through cortical stimulation mapping

(white numbers) show the strongest TLRs to the voice onset times. However, the significant TLRs are not limited to these electrodes. In Fig. 4B, the TLR for each electrode (x-axis) and frequency (y-axis) for both the hearing and speaking activities are averaged over the hearing and speaking time periods respectively. In Fig. 4C and D the finalPCA and finalBM control features are shown without the temporal component of their TLRs. It can be seen that both the batchPCA and finalBM representations are matches to the patterns of significant SF responses to the hearing and speaking activities. However, the finalBM control feature patterns match the SF patterns remarkable well considering that no labels were used in the learning process. The following section demonstrates that the neural correlates found by the DBNet representations also lead to better control features.

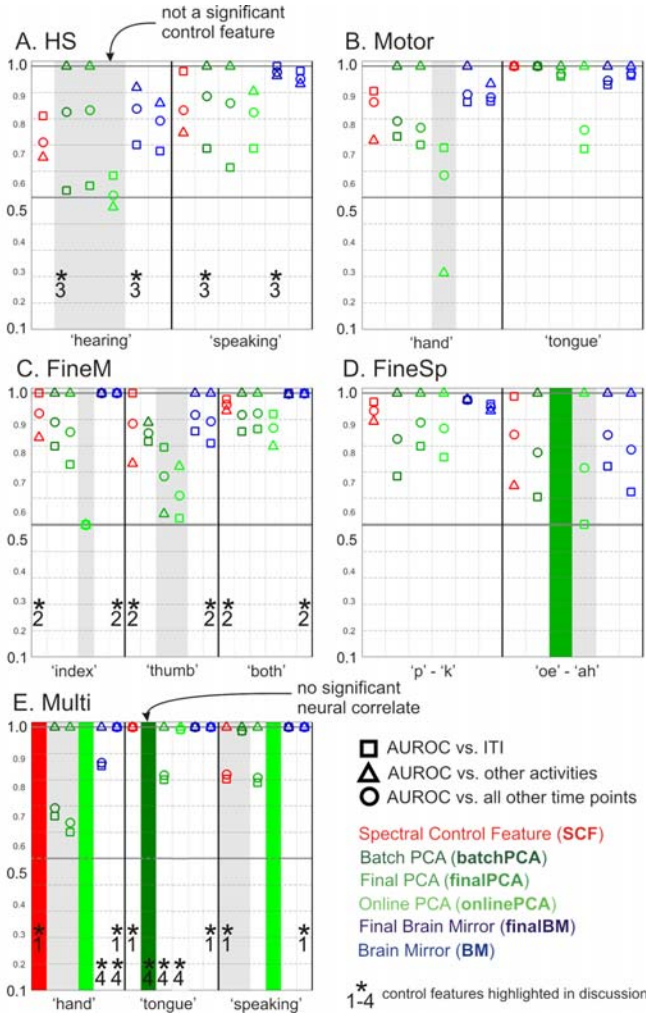

Fig. 5. Area under the Receiver-operator Curve (AUROC) for the selected control features for each activity in each task.

B. Finding Control Features

Fig. 5 summarizes the AUROC scores (squares, triangles, and circles) for all alternative representations and all different tasks and task activities. Task activities for which a control feature that significantly separated that activity from the other activities of the same task could not be found for a given a representation are indicated with light grey bars in the corresponding column. If a representation had no significant correlates for a given activity then the column is marked with a solid colored bar.

Fig. 5A focuses on the HS task. From left to right, we show the AUROC values for control features that arise from the hearing activity, based on single spectral features (in red), batch PCA (dark green), the final PCA basis (green), the online PCA, (light green), the Final Brain Mirror (dark blue), and the online Brain Mirror (blue). The right side of 3A shows the same evaluation of the control features for the speaking activity. In each case, squares represent the AUROC for classifying the task versus the ITI, the triangle is the classifier of one activity versus other activities, and the circle is the AUROC for distinguishing the activity from all other activities and ITI.

The most interesting parts of the result in 3A are that, (a) all representations encode for at least one significant neural correlate for each activity, (b) all representations led to statistically significant control features for the 'speaking' activity, (c) the BM does very well, with the highest total AUROC values (indicated with the circles) for both hearing and speaking, and finally, (d) while PCA based representations produce a control feature for the speaking activity, they do not give neural correlates that are independent control features for the hearing condition. These representations are marked in gray. A discussion of the fundamental differences in representations leading to the AUROC results in the columns highlighted with a '*3' in Fig. 5A is given in section VI.B.

Fig. 5B shows results for the Motor task. These results show similar features to those of the HS task. In particular; (a) the BM representations give consistently good results, (b) the onlinePCA representation performs well on the tongue task, and poorly in creating a classifier to capture the hand motion. The fact that the AUROC value for the hand activity control feature verses the tongue activity is considerable below 0.5 (witch indicates equal response value to both activities) arises due to the fact that there is an increase in response value for both activities but the response to the hand activity is consistently less the that of the tongue activity.

Fig. 5C shows results for the FineM task. For this task, it was again the case that not all representations created statistically significant control features. The most significant result here is that the (BM and finalBM) representations perform very well, and in general, outperform the SCF representation. These results, highlighted by the columns marked with a '*2,' are discussed further in section VI.A.2).

Fig. 5D shows results for the FineSp task. Again, the online PCA based representations perform poorly and the online DBNet based representation perform as good as or better than the SCF representation.

Fig. 5E shows results from a subject doing a task that involves a simple motor activity (hand movement), and mouth motor activity (tongue movement), and a 'speaking' activity. Again, BM features perform well across all categories and the online PCA approaches do particularly poorly. Also, notice that in this task the set of the top 20 SCFs did not include a significant neural correlate to the hand movement activity while the BM representations did. The columns displaying this result are highlighted with a '*1' and discussed in section VI.A.1). An additional result ('*4'), that is discussed in section VI.A, is that the BM representation outperforms the finalBM representation for the hand activity

and onlinePCA out performs finalPCA for the tongue activity.

C. Temporal Evolution of Control Features

The second set of experiments explores the time required for the Brain Mirror or Online PCA to compute useful features. Fig. 6A considers the HS task. The green area corresponds to the best control features defined by online PCA. The dark green line shows the AUROC of the classifier for speaking vs. ITI, and the bottom of the green region shows the AURUC of the classifier for speaking vs. hearing. The blue regions show the same plot for the BM features. The most relevant part of this plot is the indication that the BM features become good at the classification tasks within 45 seconds, and earlier than the online PCA control features. While we do not show these plots, for all tasks, the AUROC for the best control feature converges within a minute.

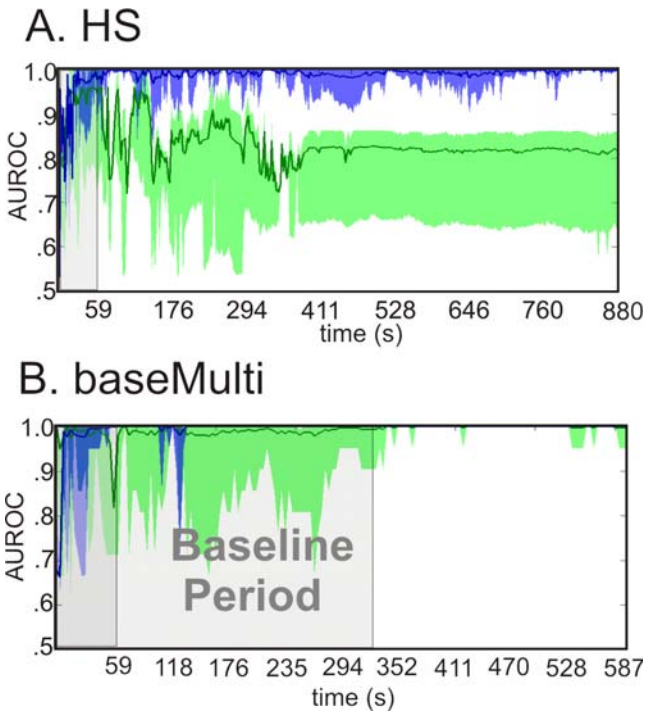

Fig. 6. Temporal evolution of the BM (blue) and onlinePCA (green) neural correlate basis. A. The AUROC scores given by the best feature responses of the onlinePCA and BM encoded at each time point of the first 9/10^{ths} of the HS task data to the last 1/10th of the data. B. The AUROC scores for the first 3/5^{ths} of the baseMulti task data to the last 2/5^{ths} of the data. The dark solid color lines show the AUROC of the classifier for the best activity vs. the ITI times and the other activities (corresponding to the circles in Fig. 4), and the bottom of the light shaded regions show the span between the AURUC of the classifier for the best activity vs. other activities and the best activity vs. ITI (corresponding to the triangles and squares in Fig. 4 respectively).

For many of the tasks, there is more than a minute of rest time before the task starts, so the patterns encoded at the beginning of the task can only reflect patterns in resting brain activity. However, the responses of those patterns capture variability in non resting data also. Those patterns classify data from the subject task activity with performance comparable to patterns defined using both rest and all the data through the end of the task. The Multi task highlights this because in this task a 5 minute period of 'resting' activity was recorded before the subject started to perform the task activities. Fig. 6B shows that after a short initialization period, the AUROC is high during the entire length of this 'resting' period and does not significantly increase once the subject starts to perform activities.

VI. Discussion

The most important result from this work is that the Brain Mirror was able to quickly learn the best control features, and our analysis highlights the potential for Deep Belief Networks to offer better representations of ECoG activity patterns. Online training of Deep Belief is possible, converges quickly to useful features, and we showed this working across multiple subjects in multiple tasks with grids of different electrode sizes. The ability to quickly learn patterns that represent meaningful changes in brain activity for a range of diverse tasks is crucial for an interactive system that aims to allow subjects to find controllable brain signal features for BCI. This is crucial for the goal of not relying on what the experimenter thinks should work for a subject and allowing the subject to explore which brain activity patterns they can control and how they can control them.

Particularly interesting are the Fine Motor Control and Fine speech task experiments (Fig. 5C-D). In these cases, the ECoG data is from smaller grids covering much less cortex. The Brain Mirror system found neural spectral patterns that could distinguish between 3 different fine motor tasks from a micro electrode grid and two different classes of mouth motor movements from a high-density grid. These results are on par with the best functional detection results reported to date, even though the Brain Mirror does not require pre-screening for control features. This is important because Microgrids and high-density grids are very new and promising technologies for BCI.

Four advantages of the Brain Mirror system that allow it achieve such good results are illustrated below.

A. BM features verses SCF

The BM DBNet based representation has an advantage over the standard SCF representation when there is significant overlap among the SCF from separate activities. Two specific problems arising from overlapping SCF and why the BM representation avoids these pitfalls are discussed below.

1) When the most significant SCFs are similar for several activities.

Multi task results serve as a good illustration of the case when the major SCF is shared by several activities. Fig. 7A shows the TLRs of the three activities. The electrodes 19 and 22 containing the control features found from the set of SCFs for the tongue movement and speaking activities respectively marked with white numbers. For each subplot, frequency is given on the y-axis and time relative to the time locking point is give on the x-axis. Fig. 7B shows the temporal spectral patterns of the BM control features formed by multiplying the DBNet encoded spectral patterns by the TLR for each activity.

The wideband spectral amplitude increase seen in Electrode 19 is the dominant SCF for all three of the task activities. However, the rules of finding distinct control features for each task dictate that it can only be chosen as the control feature for one of the activities. Since it has the largest response for the speaking activity, it is chosen as the control feature for that activity. While this still leaves the lower frequency decrease in electrode 22 as a good alternative for the tongue movement activity, there is no good alternative for

the hand movement activity. In fact, the higher frequency increase in electrode 19 for the hand movement activity did not even make it into the set of top 20 SCFs for this task. This resulted in no significant SCF neural correlates for the hand movement activity.

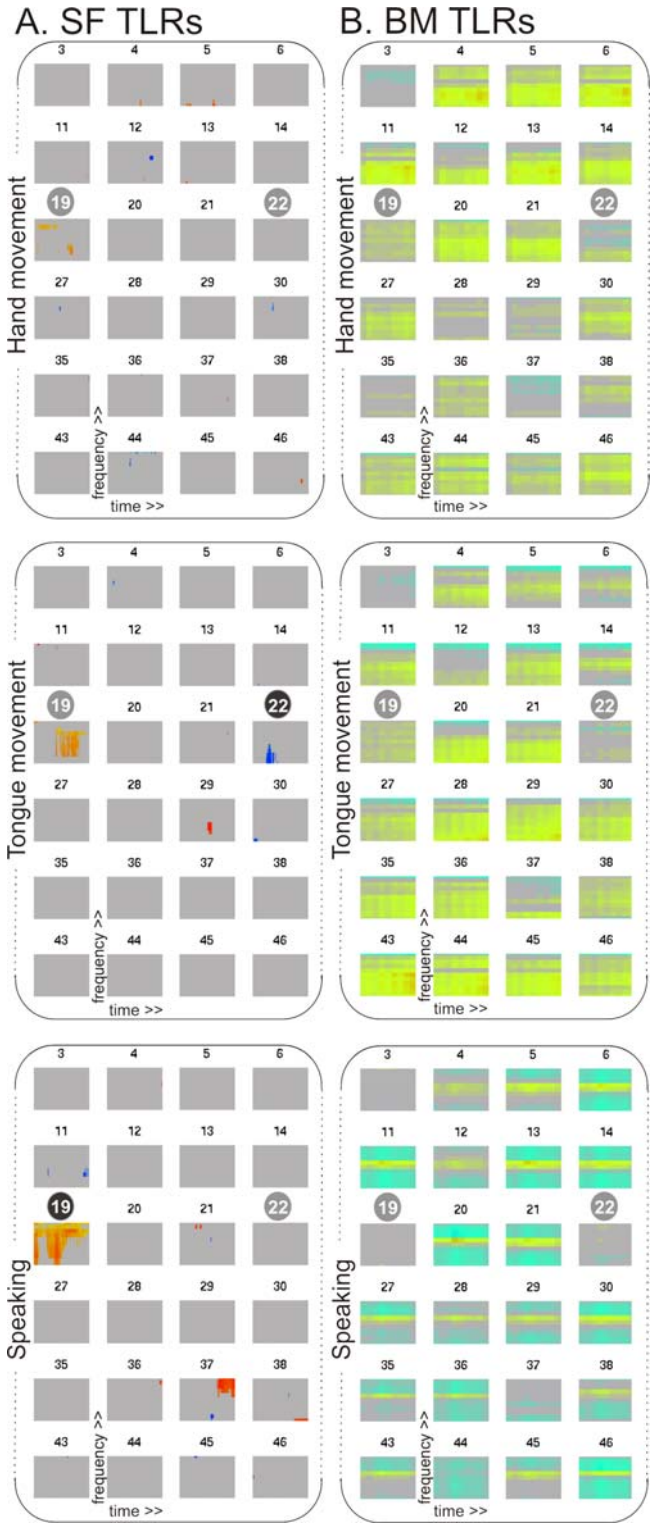

Fig. 7. SCF vs. BM control feature. A. TLR patterns for a subset of electrodes 48 electrodes from the Multi task for the 'hand movement', 'tongue movement', and 'speaking' activities that include the electrode with the largest TLR for all three activities (electrode 19). B. TLR patterns of the chosen BM control features for the 'hand movement', 'tongue movement', and 'speaking' activities over the same subset of electrodes. A subset of the total 48 electrodes was chosen due to space concerns.

In contrast, neither electrode 19 nor 22 are prominent components of any of the three BM control feature patterns. Because the BM features are anatomically much broader, collective small spectral changes outside these two electrodes can be used to distinguish the task activities. It is not uncommon that many of the top SCFs are shared among several activities and the global view of the changing spectral patterns given by the DBNet representations better distinguish such activities. As shown here, this can be the case even when the DBNet is not trained using activity labels.

2) When a single activity has SCFs with opposite signs from the same channel that overlap in frequency range.

The FineM task results serve as an example of when a SCF is modulated both positively and negatively by the same activity Fig. 8A shows the TLRs to the index, thumb, and pinch movement activities of the FineM task. Fig. 4B shows the TLRs of the SCF chosen as the control features for this task. The TLR for the index SCF control feature decrease before the onset of index figure movement (indicated by the black vertical line). Fig. 8A clearly shows that the chosen index control feature is also positively modulated by the pinch activity. The decrease in response value preceding the index movement discriminates the response values of this SCF to index and pinching movements when the TLR patterns are considered by the classifier. However, the decrease makes this feature less suited as a control feature since each activity driven increase is first preceded by a decrease.

On the other hand, the index activity BM control feature focuses in on this decrease, which is repeated across many of the electrodes, as a separate pattern that can itself be a control feature (Fig. 8C). The pattern of increases across the same electrodes in the same frequency ranges after movement onset can then be used as the pinch control feature. Such a pattern of coupled increases and decrease is also not an uncommon characteristic of SFs (see Fig. 1 for additional examples in the FineSp task).

B. BM features verses PCA based features

The BM representation also has an advantage over the PCA based methods when the global control feature patterns share significant spectral components. PCA finds the best linear basis for representing consistent variations across the data set. This offers feature responses over many electrodes and provides robustness to noise, but these large variations may not be best suited for classification tasks. The fundamental difference in the two types of pattern encoded with DBNet and PCA representations are well illustrated by the results of the HS task illustrated in Fig. 4.

The main spectral feature for distinguishing the speaking activity is both the increase in the 100-150 Hz and the decrease in the 4-25 Hz bands, as highlighted by the dark grey ovals in the left columns of Fig. 4B. The dark grey ovals in the left columns of Fig. 4B point to the fact that the hearing activity demonstrates the opposite pattern. Fig. 4C and D show that both the batchPCA and finalBM encode for this opposite hearing pattern in their respective hearing control features. However, since PCA creates an orthogonal basis, the opposite of this pattern can not be encoded as a separate component. This is illustrated by the fact that the increase in the 100-150 Hz bands in the batchPCA speaking control feature is not coupled with a decrease in the 4-25 Hz bands. In contrast, the non-linear finalBM representation is able represent the opposite patterns as two separate features.

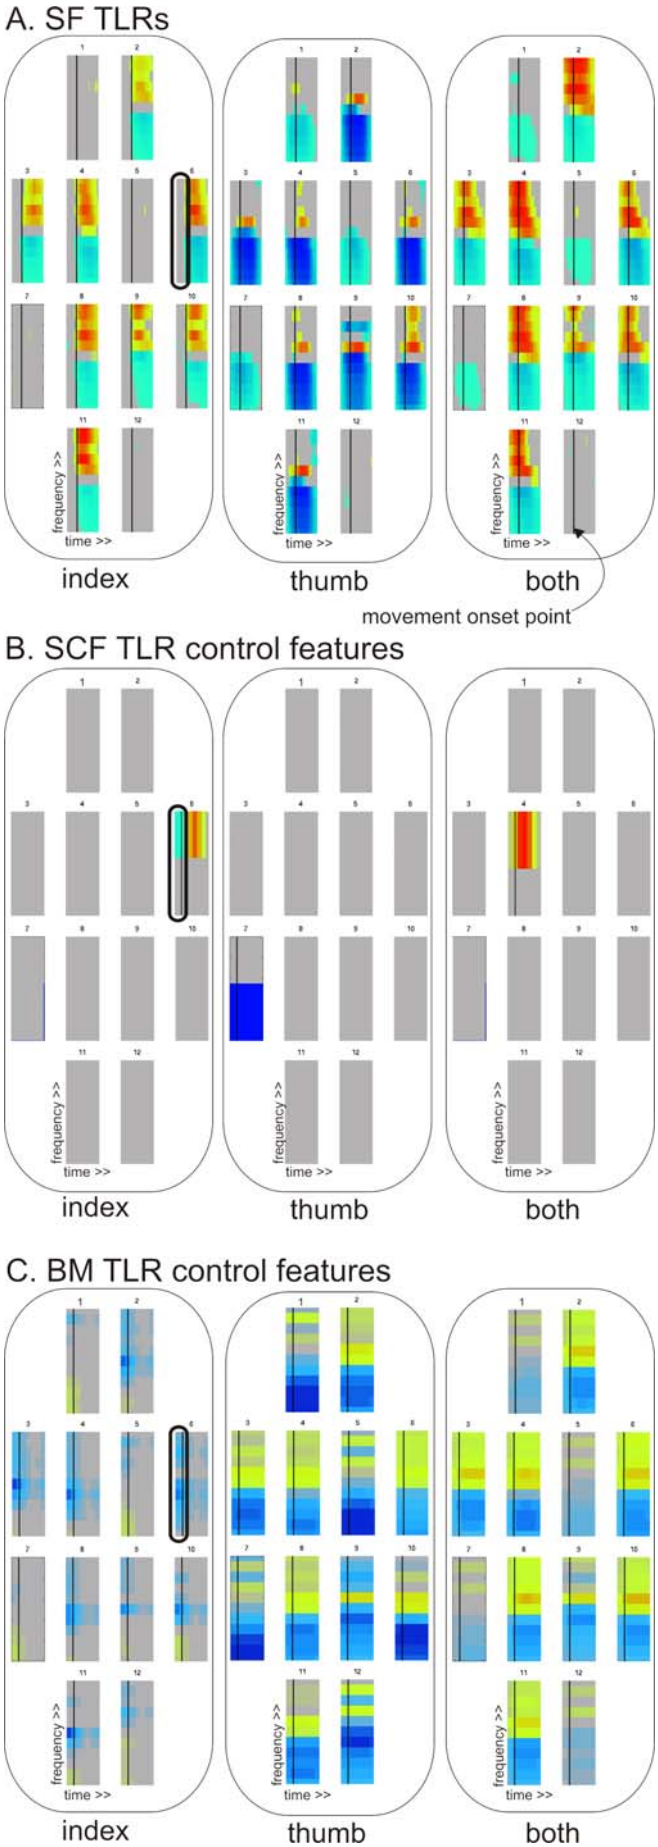

Fig. 8. FineM control feature patterns. A. SF TLRs for the activities of index movement, thumb movement, and pinching movement from left to right respectively. B. SCF control features for each activity weighted over time by their TLRs. C. BM control features for each activity weighted over time by their TLRs.

A. Online feature adaptation verses static features

It is also significant that the online BM control features were rarely worse and sometimes better than those of the finalBM. This shows that the DBNet in the Brain Mirror system continues to adapt to changes in the ECoG signal characteristics as the task is performed can lead to more consistent control features. The advantage of continued adaptation is also seen in the results of the onlinePCA and finalPCA for tongue movement activity of this task. Continued adaptation to the spectral features in the ECoG signal can be expected to be more important in tasks such as the Multi task where activities are performed for longer periods of time and are spaced further apart due to the known phenomena of habituation and resting signal drift in ECoG. However, the generally poor results of the onlinePCA representation and impressive results of the DBNet representation indicate that the DBNet algorithm is much better suited for online training.

In addition, the fact that the online methods converge quickly to significant neural correlates that are significant control features for data from the end of the task, in some cases even before task activities are performed, indicates that online adaptation is able to cope with the changes in neural electrophysiological spectral properties while retaining the functionally informative aspects of the signal.

VII. Summary

We believe the Brain Mirror offers a compelling machine learning tool to capture patterns of brain signals for BCI. First, we demonstrated that we can quickly and reliable find neural patterns that correlate with activity across several subjects and multiple tasks using ECoG, high-density ECoG, and Micro-ECoG data streams. Second, we showed that the Brain Mirror system can find patterns that correlate with future activities because it is able to find a rich representative basis from just observing rest activity. These results show that we can improve the current clinical approaches that require multiple days for a patient to undergo screening tasks before finding a control feature and replace it with the Brain Mirror interactive tool which allows patients to quickly find controllable features in their own brain signals.

Finally, we have argued in the Discussion section that these results were achieved because the Brain Mirror uses a Deep Belief Network algorithm. The DBNet approach showed great promise in both online training and distinguishing many different functional responses.

References

[1] A. K. Engel, *et al.*, "Dynamic predictions: oscillations and synchrony in top-down processing," *Nat Rev Neurosci,* vol. 2, pp. 704-16, Oct 2001.

[2] W. Singer, "Understanding the brain. How can our intuition fail so fundamentally when it comes to studying the organ to which it owes its existence?," *EMBO Rep,* vol. 8 Spec No, pp. S16-9, Jul 2007.

[3] E. C. Leuthardt, *et al.*, "A brain-computer interface using electrocorticographic signals in humans," *J Neural Eng,* vol. 1, pp. 63-71, Jun 2004.

[4] K. J. Miller, Schalk, G., Leuthardt, E.C., Shenoy, P., Rao, R.P.N., Ojemann, J.G. , "Correlation in Paired One-Dimensional, Closed Loop, Overt, Motor Controlled BCI," *Journal of Technical University of Graz,* 2007.

[5] G. Schalk, *et al.*, "Decoding two-dimensional movement trajectories using electrocorticographic signals in humans," *J Neural Eng,* vol. 4, pp. 264-75, Sep 2007.

[6] K. J. Wisneski, *et al.*, "Unique cortical physiology associated with ipsilateral hand movements and neuroprosthetic implications," *Stroke,* vol. 39, pp. 3351-9, Dec 2008.
[7] L. R. Hochberg, *et al.*, "Neuronal ensemble control of prosthetic devices by a human with tetraplegia," *Nature,* vol. 442, pp. 164-71, Jul 13 2006.
[8] D. J. McFarland and J. R. Wolpaw, "EEG-based communication and control: speed-accuracy relationships," *Appl Psychophysiol Biofeedback,* vol. 28, pp. 217-31, Sep 2003.
[9] C. Guger, *et al.*, "How many people are able to operate an EEG-based brain-computer interface (BCI)?," *IEEE Trans Neural Syst Rehabil Eng,* vol. 11, pp. 145-7, Jun 2003.
[10] B. Blankertz, *et al.*, "The non-invasive Berlin Brain-Computer Interface: fast acquisition of effective performance in untrained subjects," *Neuroimage,* vol. 37, pp. 539-50, Aug 15 2007.
[11] D. M. Branco, *et al.*, "Functional variability of the human cortical motor map: electrical stimulation findings in perirolandic epilepsy surgery," *J Clin Neurophysiol,* vol. 20, pp. 17-25, Feb 2003.
[12] G. Schalk, *et al.*, "BCI2000: a general-purpose brain-computer interface (BCI) system," *IEEE Trans Biomed Eng,* vol. 51, pp. 1034-43, Jun 2004.
[13] P. Brunner, *et al.*, "A practical procedure for real-time functional mapping of eloquent cortex using electrocorticographic signals in humans," *Epilepsy Behav,* vol. 15, pp. 278-86, Jul 2009.
[14] E. Maris, Oostenveld, R., "Nonparametric statistical testing of EEG- and MEG-data," *J. of Neuro. Sci. Methods,* pp. 177-190, 2007.
[15] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science,* vol. 313, pp. 504-7, Jul 28 2006.
[16] G. E. Hinton, *et al.*, "A fast learning algorithm for deep belief nets," *Neural Comput,* vol. 18, pp. 1527-54, Jul 2006.

Zacahry V. Freudenburg is currently working as researcher at the Rudolf Magnus Institute, UMC Utrecht, Utrecht, The Netherlands while finishing the PhD degree in Computer Science at the Department of Computer Science and Engineering at Washington University, St. Louis, MO. Mr. Freudenburg received the M.S. degree in Computer Science from the University of Groningen in the Netherlands and his B.S. degree in Physics from Beloit Collage, Beloit, WI.

His research interests include complex neurological system modeling, brain computer interface development, and real time high dimensional data clustering techniques.

Nicolas F. Ramsey is a full Professor at the Dept of Neurosurgery at the University of Utrecht. He obtained his masters degree in Psychology in Utrecht in 1987 and completed his PhD thesis on brain mechanisms involved in the rewarding effects of cocaine in rats in 1992. With a Fogarty Fellowship he then moved to the United States to continue his addiction research in humans at the National Institute on Alcohol Abuse and Alcoholism, where he become involved in the development of functional Magnetic Resonance Imaging technology for non-invasive measurement of brain functions. In 1993 he moved to the National Insitute on Mental Health to continue this work focussing on schizophrenia and cognition. In 1995 he returned to Utrecht to start an fMRI program in the Dept of Psychiatry. After several years of also being closely involved in functional brain research in epilepsy patients with electrocorticography, he moved to where he became a full Professor in 2007.

He has started a brain research program focussed on Brain-Computer Interfacing and new functional imaging methods including high-field MRI (human 7 Tesla), pharmacological imaging and clinical applications.

Mark D. Wronkiewiczis an undergraduate senior in the Biomedical Engineering program at Washington University in St. Louis. Mark is minoring in Electrical Engineering and plans on attending graduate school after 2012 graduation to continue studying Biomedical Engineering.

He is a part of Dr. Eric Leuthardt's research laboratory working on controlling external devices with brainwaves acquired through EEG and ECoG. Connected with this, he is a part of an undergraduate research team developing an EEG controlled hand orthosis for severely disabled stroke patients. For five years, he has also been part of a research team under Dr. William Stoecker at the Missouri University of Science and Technology (MS&T) with the goal of developing an automatic skin cancer detection program. During the summer of 2011, he will participate in a medical device design Fellowship directed by the Center for Innovation in Neuroscience and Technology (CINT) and sponsored by the Stryker Corporation. Currently, he is developing both an EEG brain-controlled game and an automatic pill identifying application for the iPhone and continuing work on the stroke orthosis.

Mr. Wronkiewicz is a part of Alpha Eta Mu Beta, the undergraduate BME honor society as well as the newly founded WUSTL BCI team, which evolved out of the stroke hand orthotic project.

William D. Smart is an Associate Professor of Computer Science and Engineering at Washington University in St. Louis, where he co-directs the Media and Machines Laboratory. He holds a Ph.D. and M.S. in Computer Science from Brown University, an M.S.C in Intelligent Robotics from the University of Edinburgh, and a B.Sc. (Hons) in Computer Science from the University of Dundee.

His research interests span the areas of mobile robotics, machine learning, human-robot interaction, and brain-computer interfaces.

Robert Pless is an Associate Professor of Computer Science and founder of the Media and Machines Laboratory at Washington University in St. Louis. Dr. Pless has a Bachelors Degree in Computer Science from Cornell University in 1994 and a PhD from the University of Maryland, College Park in 2000.

He is active on the program committees of the International Conference on Computer Vision (ICCV) and the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), and chaired the IEEE workshop on Omnidirectional Vision and Camera Networks (Omnivis 2003). His research focus is the statistics and geometry of video, including anomaly detection and motion pattern analysis of surveillance video and MR-imagery.

Dr. Pless he received the NSF CAREER award in 2006.

Eric C. Leuthardt is a neurosurgeon who is currently an assistant professor with the Department of Neurological Surgery and the Department of Biomedical Engineering at Washington University in St. Louis. Dr. Leuthardt is Director of the Center for Innovation in Neuroscience and Technology. He received his B.S. in Biology and Theology at St. Louis University in 1995 and received his M.D. at the University of Pennsylvania's School of Medicine in 1999. He went on to complete his training at Barnes Jewish Hospital and Washington University in St. Louis in 2005 and went on to pursue a combined fellowship in epilepsy and spinal surgery at the University of Washington in Seattle in 2006.

His research has focused on neuroprosthetics.– devices linked to the brain that may lead to cures for paralysis, allow patients to move artificial limbs, or restore other neurological functions. His work in the field of neuroprosthetics and neurosurgical devices has yielded him numerous accolades as a scientist, a neurosurgeon, and an inventor. In 2004, for his work "A Brain-Computer Interface Using Electrocorticographic Signals in Humans" he was awarded the James O'Leary Prize for Outstanding Neuroscience Research at Washington University in St Louis. Recently, the Academy of Science in St Louis awarded him the Innovator Award for his research and translation efforts. He was given one of the highest acknowledgments in his field by being presented with the Annual Award of the American Academy of Neurological Surgery in Berlin, Germany. On a national level, he was named one of the Top Young Innovators by MIT's magazine Technology Review. The magazine names 100 individuals under the age of 35 each year whose work in technology has global impact..

Dr. Leuthardt uses an integrated approach by employing multiple domains of expertise ranging from biomedical engineering, clinical neurosurgery, mathematical modeling, complex signal analysis, and computer programming. In addition to numerous peer reviewed publications, Leuthardt has numerous patents on file with the U.S. Patent and Trademark Office for medical devices and brain computer interface technologies.