

## **Distribution Agreement**

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Kejun Li

April 15, 2019

# Improving Decoding of Electrocorticographic Signals using Deep Learning

by

Kejun Li

Dr. Chethan Pandarinath  
Adviser

Neuroscience and Behavioral Biology

Dr. Chethan Pandarinath  
Adviser

Dr. Joseph Manns

Committee Member

Dr. Michael Crutcher

Committee Member

2019

# Improving Decoding of Electrocorticographic Signals using Deep Learning

by

Kejun Li

Dr. Chethan Pandarinath

Adviser

An abstract of  
a thesis submitted to the Faculty of Emory College of Arts and Sciences  
of Emory University in partial fulfillment  
of the requirements of the degree of  
Bachelor of Sciences with Honors

Neuroscience and Behavioral Biology

2019

## Abstract

### Improving Decoding of Electrocorticographic Signals using Deep Learning by Kejun Li

Brain-machine Interfaces (BMIs) are technologies that aim to assist people who suffer from loss of motor function due to accidents or neurodegenerative diseases, by providing a pathway between brain and external assistive devices. Electrocorticography (ECoG) is one of the recording methodologies that could be used to infer people's movement intention for BMI systems. ECoG covers large areas of the brain and has a demonstrated safety record. However, its low spatial resolution usually leads to worse accuracy in decoding movement intention than BMIs that use intracortical recordings. In this study, we found that ECoG decoding performance could be improved by modeling neural population dynamics using a deep learning technique: Latent Factor Analysis via Dynamical Systems (LFADS). LFADS attempts to uncover structure from neural population activity that is consistent with a low-dimensional dynamical system. In previous applications to intracortical recording from motor cortex, LFADS improved decoding accuracy by uncovering estimates of neural population dynamics on a single-trial, moment-by-moment basis. However, since LFADS was previously evaluated using intracortical recordings with high spatial resolution, it was unclear whether LFADS is appropriate for de-noising recordings with low spatial resolution, such as ECoG. To test this, we applied LFADS to ECoG recordings from seven human subjects who were being monitored as part of clinical treatment for epilepsy or glioma. Subjects performed a one-finger button-press task as their finger kinematics and kinetics were recorded. We compared our ability to decode behavioral states (i.e. pre-movement, movement or force) before and after application of LFADS. The accuracy of the discrete classifier was

improved significantly by the application of LFADS. Our results represent a new avenue toward improving the performance of ECoG-based BMI systems.

Improving Decoding of Electrocorticographic Signals using Deep Learning

by

Kejun Li

Dr. Chethan Pandarinath  
Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences  
of Emory University in partial fulfillment  
of the requirements of the degree of  
Bachelor of Sciences with Honors

Neuroscience and Behavioral Biology

2019

## Acknowledgements

I would like to sincerely thank my advisor Dr. Chethan Pandarinath for giving me the opportunity to complete my honors thesis project in his lab, for always offering me timely feedback and constant encouragement, and for being such an outstanding mentor. I would also like to thank Drs. Marc Slutzky and Robert Flint, without whom this project would not be possible, for collecting, preprocessing, and graciously sharing data and related figures with us. I would also like to thank Lahiru Wimalasena and Dr. Mohammad Reza Keshtkaran for being great teachers, and for helping me brainstorm ideas and troubleshoot when things went wrong.

## Table of Contents

<b>Terminology and Abbreviations</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>2 - 5</b>
<b>Methods</b> .....	<b>6 - 14</b>
1. Subjects and Recording	
2. Experimental Protocol	
3. Spectral Feature Extraction	
4. Spectral Feature Extraction and Power Amplitude (PA) Matrix	
5. LFADS Structure	
6. Hyperparameter Tuning with Population-based Training	
7. Principal Component Analysis	
8. Decoding Analysis	
<b>Results</b> .....	<b>15 - 17</b>
<b>Figures</b> .....	<b>18 - 25</b>
1. Demonstration of PCA analysis	
2. Decision tree diagram	
3. Experimental protocols and recording ECoG array placement for the subjects.	
4. Scheme for discrete states classifiers	
5. Percent variance explained by the first three principal components from PCA for PA matrix and LFADS-denoised ECoG	
6. Example trial for ECoG and PCA trajectories of PA matrix and LFADS-denoised ECoG	

7. Bagged trees decoding performance

**Discussions** ..... 26 - 29

**References** ..... 30 - 33

## **Terminology and Abbreviations**

BMI: brain-machine interface

ECoG: electrocorticography

EEG: electroencephalography

LFADS: latent factor analysis via dynamical systems

MEA: microelectrode arrays

PA: power amplitude

PC: principal component

PBT: population-based training

PCA: principal component analysis

RNN: recurrent neural network

SCI: spinal cord injury

## Introduction

Brain-machine interfaces (BMIs) are technologies that allow communication between brains and external devices, such as computers or prosthetics. BMIs are a promising approach to restore lost motor function in patients who suffer from motor or sensory impairments, for instance by using brain activity to control electronic stimulation of paralyzed muscles or to command robotic limbs (Ajiboye et al., 2017). One of the key components of BMI system is the decoder, which is used to translate brain activity into control signals for external devices. Developing decoders that could infer subject movement intention accurately remains a critical challenge for building BMIs.

The most accurate decoding to-date has come from intracortical BMIs that require surgical implantations of microelectrode arrays (MEA) (Hochberg et al., 2012; Ajiboye et al., 2017). Besides reactive activities such as gliosis and progressive meningeal fibrosis that may degrade the materials of the implants, signal longevity is a major concern for MEAs since the neurons recorded from these microelectrodes change considerably over time, which makes the decoding process more challenging (Nurse et al., 2016). As a result, researchers are considering alternative methods to monitor neural activity. Electrocorticography (ECoG) represents a promising interfacing option among current non-invasive and invasive BMI recording methods. Compared to scalp-recorded electroencephalography (EEG), ECoG has higher spatial resolution, a wider informative frequency range, higher characteristic signal amplitudes (50-100  $\mu\text{V}$  vs. 10-20  $\mu\text{V}$ ), and less vulnerability to ambient noise (Schalk et al., 2008; Liang and Bougrain, 2012). ECoG also covers a much larger area than intracortical recording, potentially enabling researchers to develop decoders that rely less on local neural

activity, and are thus less affected by slight changes in the implant position. Because ECoG is widely used in the treatment of epilepsy, it provides ample opportunities to record motor cortical activity from patients and test out any BMI prototypes.

Previous work has demonstrated that ECoG signals are useful for decoding movement intention to control external devices such as computer cursors (Vansteensel et al., 2016) and robotic arms (Wang et al., 2013). However, as previously stated, the performance of ECoG BMIs are not as impressive as those of intracortical BMIs (Pandarinath et al., 2017). In this project, we aimed to test whether by applying cutting-edge computational techniques (deep learning), we might be able to bypass the limitations imposed by the spatial resolution of ECoG.

We used a deep learning technique called Latent Factor Analysis via Dynamical Systems (LFADS) to uncover brain states from the recorded ECoG data. Briefly, LFADS attempts to denoise recorded neural activity based on the assumption that the complex firing patterns of motor cortical neurons reflect an underlying dynamical system. In such a dynamical system, the rules that decide how firing rates change over time remain consistent even if the external variables, such as movement directions change. LFADS aims to extract these rules to infer denoised representations of neural population activity on single trials.

We chose LFADS because it outperforms current best approaches for de-noising neural population activity, which leads to more accurate prediction of movement intention (Pandarinath et al., 2018). It accepts unlabeled neural population activity, attempts to infer the underlying dynamics, and rejects signals that are not consistent with the learned dynamics to generate a denoised version of the data. LFADS can also

be modified to combine data across many sessions to better model one underlying dynamical system. This would lead to more robust decoding across hours/days, since it creates a better estimate of underlying dynamics, while also compensating for potential instabilities in the recording process. Even though not used in this project, the ability of combining data across sessions is incredibly useful in terms of BMI study since a consistent dynamical model learned across different tasks and across days could provide critical information to guide future BMI design in terms of which recording parameters to be used, and where to place the electrodes, etc.

Because the performance of LFADS has only been evaluated on datasets with high spatial resolution, it remains unclear whether LFADS could model neural population dynamics using lower resolution data, such as ECoG. However, it has been shown that the underlying patterns for the dynamical system are broadly distributed across the cortex (Pandarinath et al., 2018), which suggests that the aggregated spiking activity ECoG recorded would still rely on the same set of rules and the location of the recording does not matter that much as long as it is still related to the task. As a result, it is likely that LFADS could perform relatively well since the decreasing in the spatial resolution of data should not significantly hinder the ability of LFADS to extract the latent factors that underlie these observed ECoG activity, which thus should not decrease that much of the overall performance of LFADS. Since ECoG has lower resolution than other invasive recording methods, if the decoding performance of ECoG data after going through LFADS is indeed compelling, it would provide strong evidence to demonstrate the potential of ECoG to be used as an alternative recording technology for BMI applications.

We applied LFADS on ECoG data collected from human patients being monitored as part of clinical treatment for epilepsy or glioma by researchers at Northwestern University. Patients performed a cursor control task by using their index fingers to press on a button (force sensor). Both finger kinematics (movement) and kinetics (force) are recorded. In order to quantify whether LFADS output was indeed meaningful, discrete decoding analysis was performed to evaluate how well we could differentiate neural representations related to kinematics from kinetics. We hypothesized that LFADS should be able to generate a denoised version of the ECoG data that results in higher decoding accuracy.

## **Methods**

### *Subjects Overview*

Seven human subjects were involved in this study (all males). One subject (S06) went through extraoperative intracranial monitoring prior to surgery as part of the treatment for epilepsy. Six other subjects had awake-intraoperative mapping prior to the resection of low-grade gliomas. The tumors of these six patients were located remotely to cortical areas that were related to hand grasp, and no upper limb sensorimotor deficits were found in their neurological testings.

All experiments were performed under protocols approved by institutional review board of Northwestern University and were conducted by researchers from Slutzky Neuroprosthetics Laboratory. Written consent was obtained prior to subjects' participation in the study. Subjects were recruited if the location of their craniotomy or the monitoring arrays were expected to cover primary motor cortex.

### *Recordings*

In subject S06, a 32-electrode (8x4) array with a 1.5mm exposed recording site diameter and a 4mm inter-electrode spacing (Integra, Inc.) were placed according to clinical need. For all subjects other than S06, a 64 electrode (8x8) high-density ECoG array with the same size of electrodes and spacing as the 32-electrode array was placed over hand motor areas. Anatomical landmarks, preoperative fMRI or transcranial magnetic stimulation, and direct electrocortical stimulation mapping were used to identify functional motor areas. Recordings were performed after the direct stimulation mapping. Intraoperative MRI was performed with Curve (BrainLab, Inc., Munich,

Germany). Primary motor cortex (M1), premotor cortex, and part of primary somatosensory cortex (S1) were covered in the recordings as demonstrated in figure 1.

ECoG was sampled at 2k Hz with Neuroport Neural Signal Processor (Blackrock Microsystems, Inc) with a bandpass filter of 0.3-500 Hz. Finger kinematics were recorded by a 22-sensor CyberGlove (Immersion). Force was recorded with a custom-built load cell sensor. Both finger kinematics and force were sampled at 2kHz.

### *Experimental Protocol*

The task was designed to only require movement from one's index finger and all the other fingers should remain motionless during the trials. As a result, only data from the sensors of the CyberGlove that were related to the movement of index finger was analyzed.

Subjects were instructed to hold their index finger at a neutral resting position. After a cue was presented, the subjects started a flexion movement, which allowed them to make contact with the force sensor using the palmar surface of their fingers. The subjects were then instructed to apply force to the force sensor to control a cursor on the monitor. The goal of their task was to make sure that the position of the cursor was matched with a force target presented on the monitor. The target force levels were randomly set for each trials. The completion of the trial is indicated by either a successful match or a 2s timeout. The subjects then extended their fingers to the previous neutral position. After a delay of 1s, the next trial began. BCI2000 software was used to present target and provide feedback for the cursor. The time resolution for the movement of cursor was 50ms.

### *Spectral Feature Extraction and Power Amplitude (PA) Matrix*

Only channels that covered primary motor cortex (M1) and premotor cortex (PM) were used for analysis. Spectral features of each ECoG electrode covering these area were extracted using a short-time Fourier transforms with a window width of 512 ms. The spectral power was log-normalized and averaged in 25 ms time bins. Another study demonstrated that ECoG activity in several specific frequency bands such as sub-bands (1-60Hz), gamma band (60-100 Hz) and fast gamma band (100-300 Hz) was related to the task (Liang and Bougrain, 2012). High gamma power (80-150 Hz) of the ECoG signals could be used to predict movement types as well as provide the best information that could be used to classify different movement types (Yanagisawa et al., 2011). Based on above evidence, we decided to use a low frequency feature (8-55Hz) and a high frequency feature (70-150 Hz) for each channel. These two features were rearranged into a power amplitude (PA) matrix through the concatenation of low frequency feature (8-55 Hz) and high frequency feature (70-150Hz) of each channel. Since there were two features for one electrode, this resulted in a matrix size of  $N \times T$ , where  $N$  equals  $2 * \text{number of electrodes}$  and  $T$  equals number of time bins. This PA matrix was later used as the input for LFADS.

### *Latent Factor Analysis via Dynamical System (LFADS)*

We used LFADS to analyze the PA matrix and inferred a denoised version. LFADS is designed based on the idea that neural population activity could be described by low-dimensional dynamics. It is important to understand what low-dimensional

system and dynamical system mean before one can understand how LFADS model neural population activity.

In cases where a large number of neurons are being monitored, if there are  $D$  neurons, it is often reasonable to consider the dimensionality of the system to be lower than  $D$ - because the recorded neurons are from the same common network, their responses are likely not completely independent of each other. The activity of this network is often considered to covary according to a small number ( $K$ ) of explanatory variables, where  $K < D$  (Cunningham & Yu, 2014). LFADS can extract these  $K$  explanatory variables, often termed latent variables or latent factors, from neural population activity. LFADS provides a denoised version of the original activity by rejecting any data variance that is not captured by these latent factors.

In a dynamical system view, these latent factors could be considered to be the factors that describe the “dynamics” or the rules that neural population activity follow. By definition, a dynamical system is “a physical system whose future state is a function of its current state, its input, and possibly some noise” (Shenoy et al., 2013 P339). Considering a general dynamical system, a dynamics system consists of the current state of the system  $\mathbf{x}(t)$ , the function  $\mathbf{F}$  that determines how  $\mathbf{x}(t)$  is being updated, and the optional input  $\mathbf{u}$  where:

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t), \mathbf{u}(t)).$$

To model such dynamical system, the encoder network (an recurrent neural network or RNN in LFADS) processes single-trial activity and infers the initial condition ( $\mathbf{g}_0$ ) of each trial. This initial condition serves as the initial state of the generator network

(another RNN in LFADS), whose task is modeling the underlying dynamics from the observed spiking activity (figure 4b).

LFADS assumes that the observed action potentials for each neuron are samples from an inhomogeneous Poisson process, and attempts to infer this process (the inferred firing rate) for each neuron. In this study, since the ECoG power signal represents aggregated spiking activity from thousands of neurons, the underlying distribution was switched to Gaussian instead of Poisson.

#### *Hyperparameter tuning with Population-based training*

To apply LFADS or other deep learning algorithms, it usually requires careful hand-tuning of the hyperparameters that define the model structure and how the model is trained to make sure optimal performance could be reached. This step could be extremely challenging and time-consuming even for experts in this field. Thus, in our project we decided to use an automatic hyperparameter tuning method (population-based training, PBT) developed by Mohammad Reza Keshtkaran in our lab that utilizes an evolutionary algorithm to search for optimal hyperparameters in combination with LFADS. This method trains a number of networks with different hyperparameters in parallel and evaluate and choose the best performing networks in the current generation to continue the training process. LFADS performance was trained on partially-sampled data and then cross-validated by the complementary samples to avoid overfitting during the training process. The performance of LFADS with PBT had been found either matched or outperformed the performance of original LFADS that was hand-tuned by experts consistently. By using such algorithm, we avoid the tedious hand-tuning process

to obtain the optimal hyperparameters, which allowed this project to be completed in time.

### *Principal Component Analysis*

Since LFADS denoises data by modeling the low-dimensional dynamics of the neural population activity, we would like to find a way to compare the dimensionality of the modeled data and PA matrix as well as visualize any dynamics LFADS might learn. A dimensionality reduction method, principal component analysis (PCA), was used for this purpose to generate a low-dimensional representation of the PA matrix and LFADS output. PCA finds a set of principal components (PCs), which are orthogonal components, i.e. linearly uncorrelated components. These PCs are able to reconstruct the original dataset and they the maximized the variance they explained in the given set of possibly correlated observed variables. After finding one PC, it then removes the variance explained by this component and seek for the next best component, and so on. For instance, considering a 2d example (Figure 1), the maximum number of PCs that could be found are two since PCA requires each PC to be linearly independent (perpendicular to each other in a 2d plane). To demonstrate how PCA find a low dimension representation, let's try to find a 1d representation (i.e. a line) of this 2d example. In order to maximize the variance explained, where the variance is defined as the average squared distance from the center (white circle) of all the data points (blue dots), the projection of each data point (red dots) on the first principal component (black line) need to be as spread out as possible. Maximizing the variance explained also minimize the the difference between the found feature (first PC) and the actual data

points, which are represented by the red lines. It is obvious in a 2d plane that the one on Figure 1b is clearly the best line of fit, but not so intuitively in higher dimension. For instance, if we are looking for neural population activity, since there are so many potential features available, we are not able to visualize the systems and determine the most important features as easy as this one. But PCA could identify the these most essential of features for us (i.e. the first several PCs) by the nature of this method. By limiting the number of PCs that are plotted, we reduce the dimension of the original dataset to a lower dimension and thus allow visualization. One way of quantifying how well PCA is able to capture the structures or variances in the original data or how different the captured low dimension representation is from the original data is by calculating the percent variance explained by each of the PCs.

In this project, the percentage of variance explained by the first three, i.e. the most important three, components were reported. The results of PCA analysis was also used to provide us with low-dimension visual representations of the original datasets.

### *Decoding analysis*

To test whether the LFADS output was meaningful, we attempted to classify behavioral states. Data were labeled according to the following scheme: time bins between the target onset and the movement onset were designated “background”; time bins between the movement onset to force onset were designated “movement”; and time bins from force onset to force onset + 0.5 s were designated “force”. The length of the “force” window here was limited by discarding all the data from 0.5s after the force

onset till the end of the trial to obtain more equal class sizes. Bootstrap aggregated (bagged) trees classifier was then used to classify the remained data.

Bootstrap aggregating, also called bagging, is a technique used to build a machine learning ensemble that consists of multiple models. Bagging is often used in combination with a set of decision trees to reduce the variance of individual decision tree. By definition, a decision tree is a model that predict the values of a target variable based on the available inputs. For instance, the diagram below is a flowchart for how decision tree makes predictions (Figure 2). The internal (non-leaf) nodes shown in the diagram are labeled according to the given input feature. (i.e. A, B, C, etc.). What's being learned during the training process is the cutoff threshold for each feature (i.e. what value A should equal to so that the tree should follow the left path).

In our case, the bagged trees classifier is initialized with several different decision trees just like this one but with input features being different features of the recording. Specifically, low-frequency or high frequency features of each channel are used instead of A, C, or D. During learning, each decision tree finds a best cutoff threshold for each feature. By comparing the learned cutoff values with the actual values of the features during each internal node, a prediction could be made when the terminal node or leaf is reached. The bagged trees take one vote from each of the individual decision tree to make the final prediction, which was usually more robust than a single decision tree.

The performance of the bagged classifiers was evaluated by 5-fold cross-validation. For each test fold, every 25 ms time bin was classified and percentage accuracy of the classifiers was reported. Either PA matrix or LFADS-denoised ECoG

was used as the input for classification. A chance level performance was also evaluated by shuffling the class labels for 1000 times.

## Results

As described, ECoG activity was recorded from seven human subjects who required intraoperative or extraoperative mapping as part of their treatment for glioma or epilepsy. Subjects performed a cursor control task, where they were instructed to press the force transducer to move the cursor to meet with the target (Figure 3a). Both finger kinematics (movement) and kinetics (force) were recorded at the same sample rate as the ECoG signals (2 kHz). The ECoG arrays covered primary motor cortex, premotor cortex and in some cases, prefrontal or postcentral cortices (Figure 3b).

To classify ECoG data into discrete behavioral states of background, movement, and force, we followed the following scheme (Figure 4a): ECoG data were first preprocessed with a short-time Fourier transform (STFT) and constructed into a power amplitude matrix, which then went through LFADS (Figure 4b). The LFADS-denoised ECoG data were then classified into the three discrete states mentioned above using a bagged trees classifier. In order to determine whether LFADS improve decoding performance, PA matrix was also used as input for a separate classifier to establish a baseline that LFADS results could be compared with.

In order to visualize the neural population activity, we applied PCA to both the PA matrix and LFADS-denoised ECoG to generate a low-dimensional representation for each subject's neural population activity. The first PC of LFADS-denoised ECoG explained far more variance of the data than that of PA matrix. This is true across all subjects, with an average of  $67.18\% \pm 9.5\%$  explained by the first PC for LFADS, and  $26.93\% \pm 13.25\%$  explained by the first PC of the PA matrix. The cumulative variance of first three PCs follow this trend as well. The average cumulative variance explained by

the first three PCs was  $46.1\% \pm 16.98\%$  and  $94.6\% \pm 2.91\%$  for PA and LFADS, respectively (Figure 5). These results suggested that the dimensionality of the dynamics LFADS extracted was incredibly low since 3 PCs were sufficient to explain most of the variance even though a much larger number of latent factors, ranging from 10-30 were allowed when applying LFADS. This also means that the visualization PCA provided, especially the 3-PC trajectories, captured majority of the variance of the original data.

The example single trial from Subject S04 clearly demonstrate that LFADS output was smoother and more structured than the PA data (Figure 6a, b), suggesting that LFADS may be rejecting noise. Visualization of PCA trajectories, both the progression of the first PC over time, and progression of 3 PCs in state space for subject 4 or subject 5, also suggested that LFADS output was denoised relative to the original PA data (Figure 6c & d, 6 e & f, 6 g & h).

While the visualization results were promising, it is unclear whether LFADS was actually removing noise from the data, or also removing essential features of the data. In order to test that, we performed a classification analysis to determine if the LFADS outputs contained enough information to classify the participant's behavioral states. We used bagged trees classifier here to classify ECoG data before (PA matrix) and after LFADS into three discrete states: background, movement, and force.

For PA matrix, mean decoding accuracies and standard error of mean ranged from  $77.2\% \pm 1.2\%$  to  $90.7\% \pm 0.7\%$  across the 7 subjects with an overall mean decoding accuracy and standard error of mean of  $83.47\% \pm 1.7\%$ . For LFADS-denoised ECoG, mean decoding accuracies and standard error of mean ranged from  $86.8\% \pm 3.8\%$  to  $93.9\% \pm 2.7\%$  across the 7 subjects. The overall mean decoding accuracy and

standard error of mean for all subjects was  $89.9\% \pm 1.15\%$ . LFADS significantly improved performance of the bagged trees classifier for 5 out of 7 subjects (Figure 6).

Overall, decoding using the LFADS-denoised ECoG resulted in better performance than the PA matrix ( $p=2^{-31}$ , Kruskal-Wallis test). The decoding accuracy for all subjects was statistically significantly higher than chance (Figure 7).

## Figures

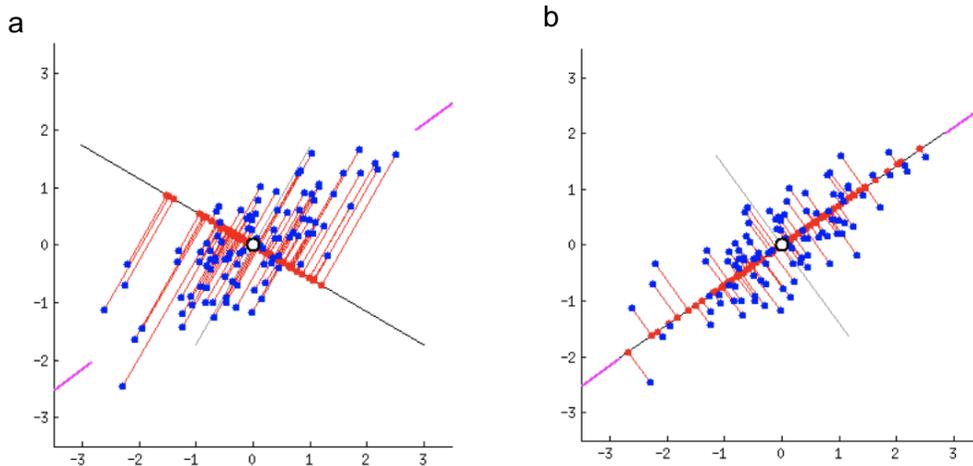


Figure 1: Demonstration of PCA analysis (plots generated using code written by amoeba, 2018). Each blue dot represents a data point in the original dataset. The black line represents potential direction of first principal component (PC). Each red point on the black line represents the projection of the data points onto the first PC. The grey line represents the direction of second PC, which is determined according to the direction of the first PC since it has to be perpendicular to it by the nature of PCA. The red line that connected blue dots and red dots are the reconstruction error (i.e. how different the low-dimension representation is from the original dataset). a) the black line does not maximize the explained variance. b) the black line maximize the explained variance

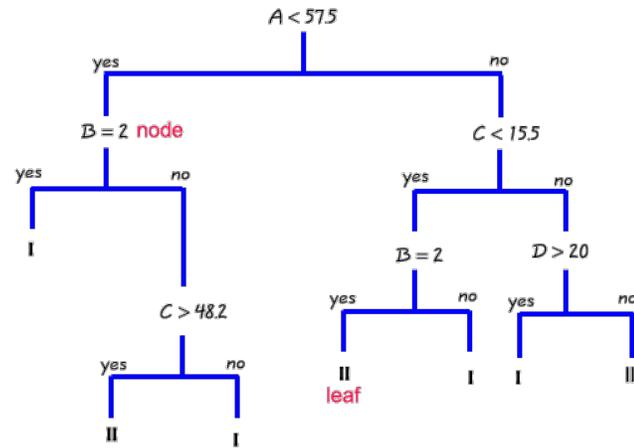


Figure 2: Decision tree diagram (Fielding, n.d.). This is an example decision tree diagram for a binary classification problem (groups I and II). The predictors, which are inputs, are either continuous (A, C & D) or categorical (B). Each internal node is associated with one of the features. By following along one of the paths for the tree, a classification result is determined if a terminal node (leaf) is reached.

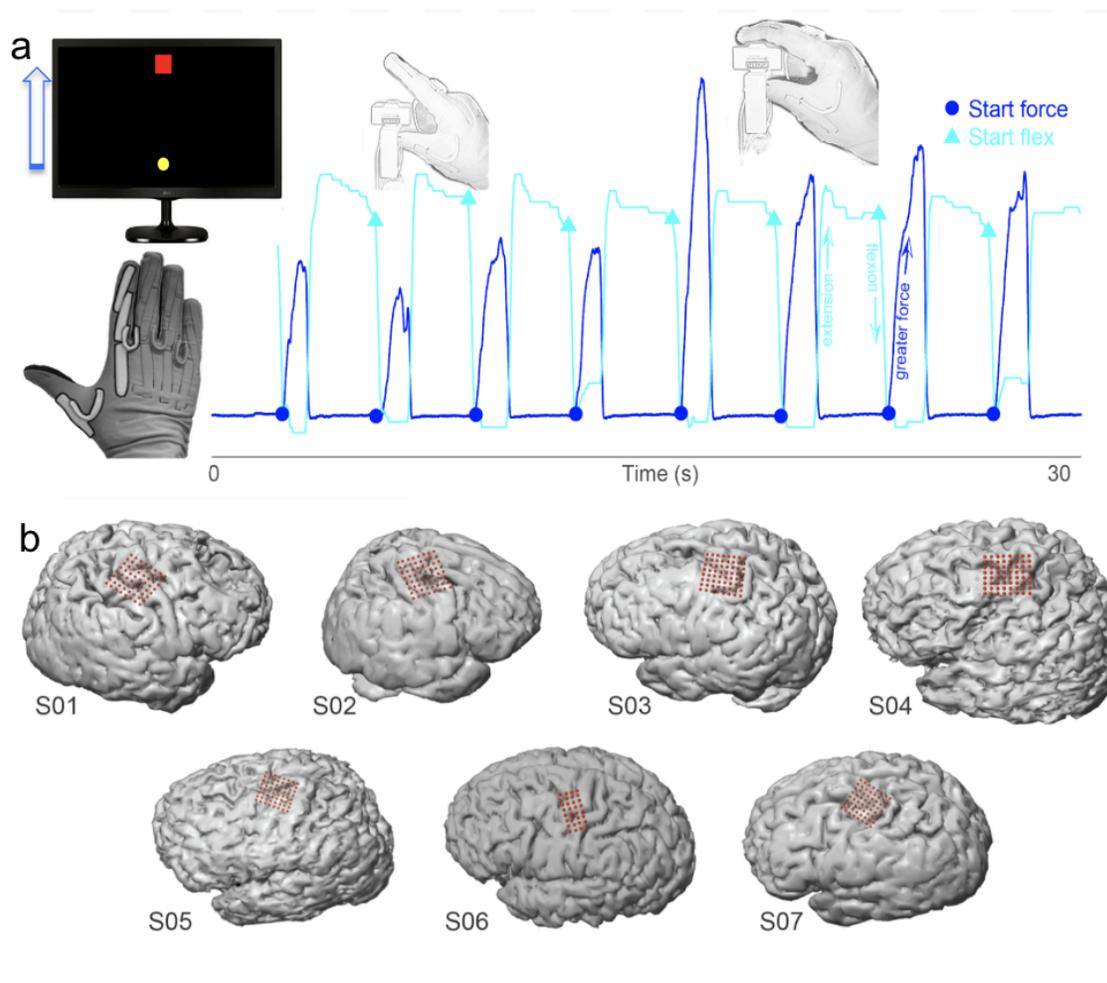


Figure 3: Experimental protocols and ECOG array placement. a) Example trials. The subject first held their index finger at a neutral resting position until the target showed up on the screen. They then started flexion movement (light blue traces) to make contact with the force transducer. In order to match the cursor on the screen with the target, they started to press the force transducer (dark blue traces) till the task was indicated as successful. They repeated the same procedure for each trial. The light blue triangle indicates the movement onset and the dark blue circle indicates the force onset. b) Electrode placement for all the subjects. S01 and S02 were recorded from the right hemisphere, while all other subjects were recorded from the left hemisphere.

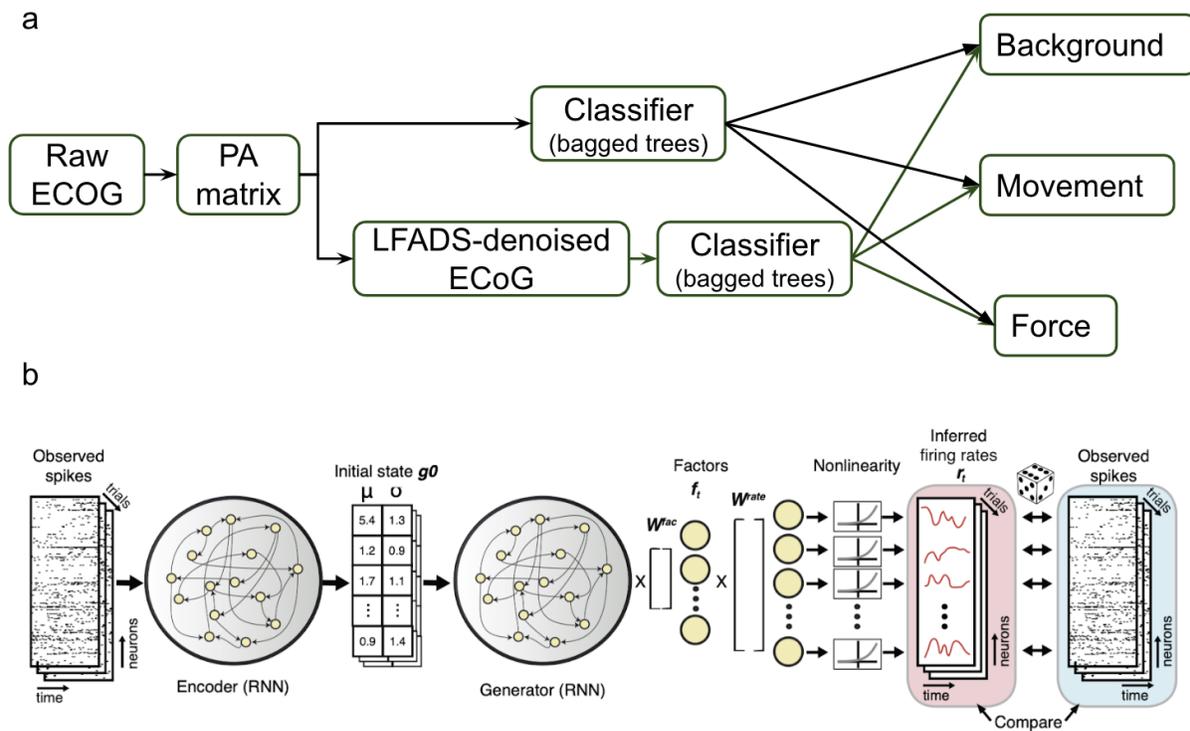


Figure 4: Scheme for discrete state classification. a) ECoG data were rearranged into the PA matrix and then denoised by LFADS. The LFADS-denoised ECoG was then used in the bagged trees classifier, to label each 25ms time bin as either “background”, “movement”, or “force”. In order to determine the effect of applying LFADS, PA matrix was also used as input for the bagged trees classifier. b) The detailed structure of LFADS. After LFADS takes in observed neural activity, the encoder RNN infers the initial state  $g_0$ . The generator RNN then uses  $g_0$  and the learned dynamics to infer the vector of activity, which are then later transformed into denoised rates.

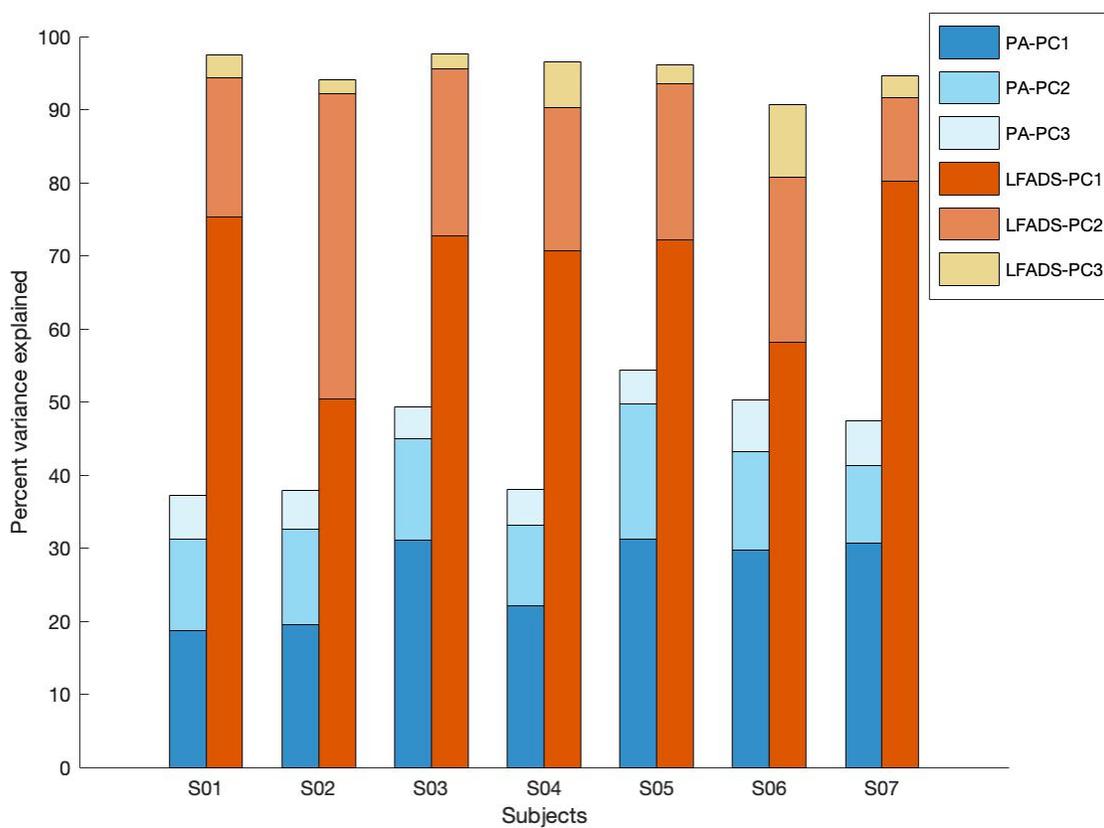


Figure 5: Percent variance explained by the first three principal components from PCA for PA matrix and LFADS. For each subject, PA matrix results are on the left (indicated by blues), LFADS results on the right (indicated by red). Each three PCs are indicated by different forms of blue or red. The intensity of the color reflected how important that PC was. For instance, the darkest blue or red means this is the most important PC (PC1) and the lightest segment would be the least important PC (PC3).

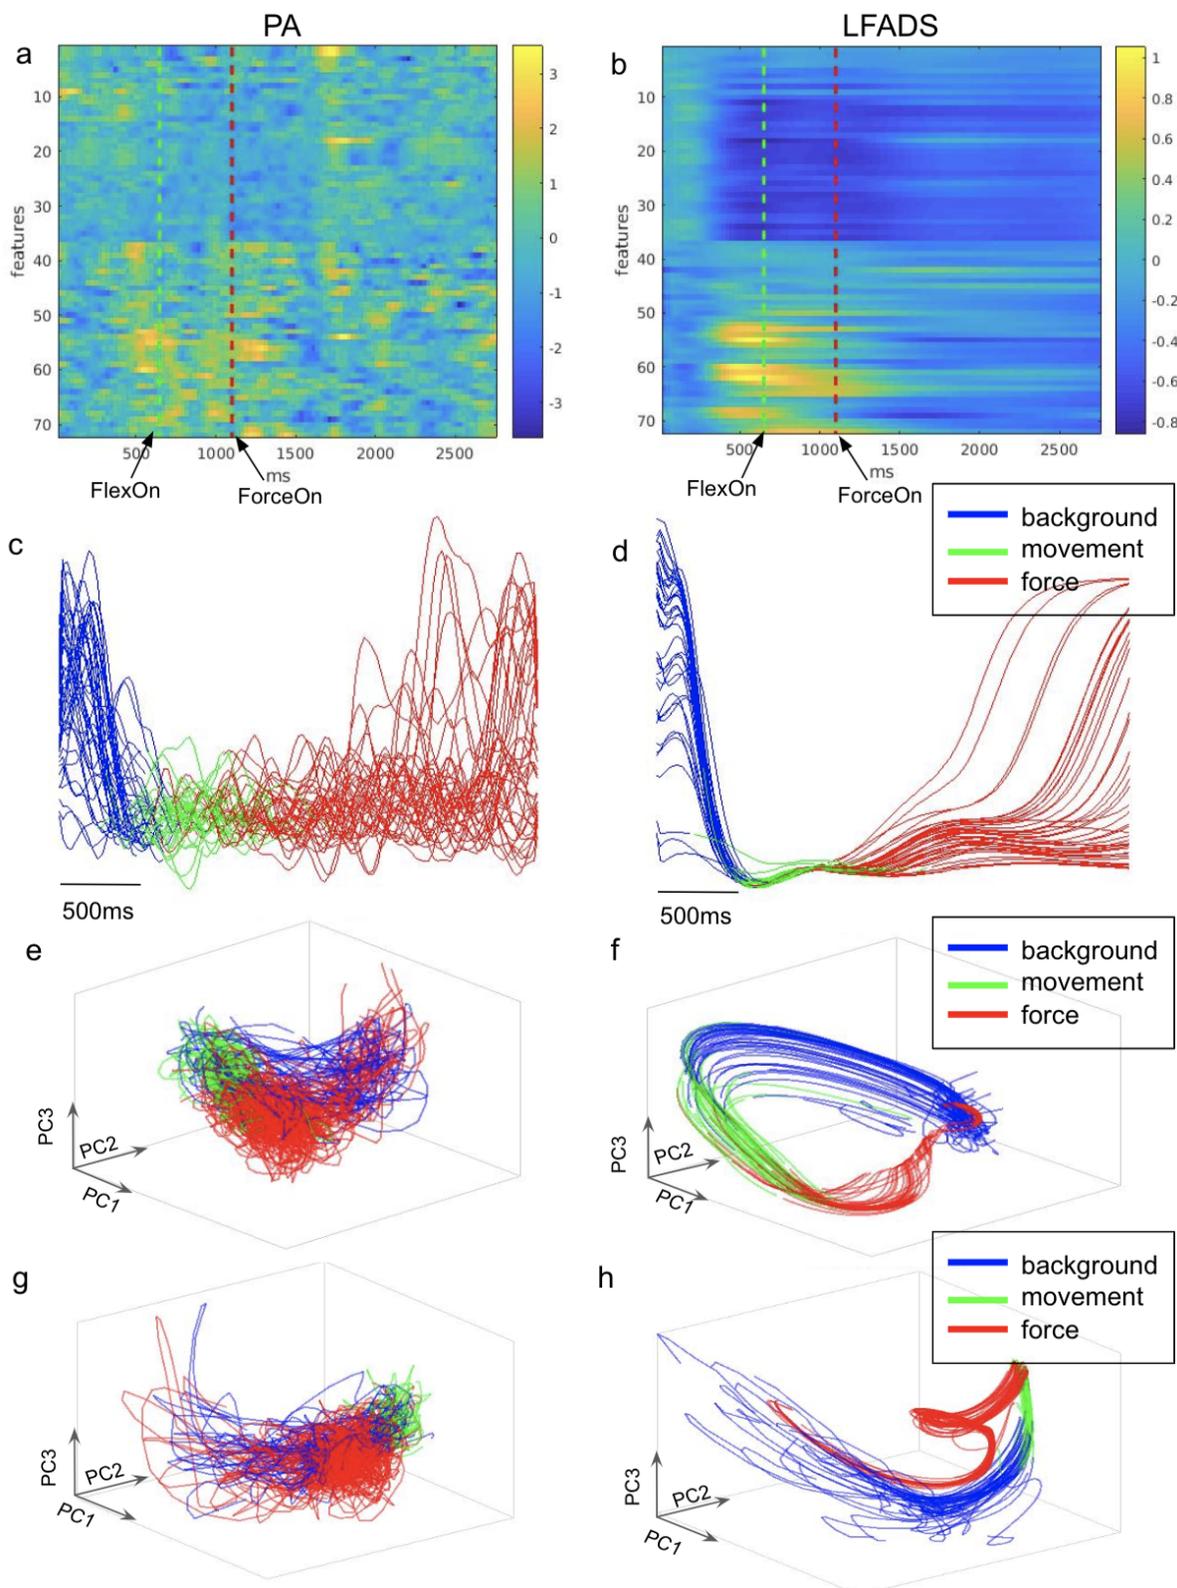


Figure 6: Example trial for ECoG and PCA trajectories for PA matrix and LFADS-denoised ECoG. a) PA matrix for one example trial subject 4 (S04)

performed. b) LFADS-denoised ECoG for one example trial subject 4 (S04) performed. Seventy-two features were used for this dataset. Features 1-36 and 37-72 are low frequency and high frequency features from 36 electrodes, respectively. For both panels a) and b), movement onset (FlexOn) is marked by green dash lines and force onset (ForceOn) is marked by the red dash lines. d) First principal component of PA matrix (c) and LFADS-denoised ECoG (d) for subject 4 (S04). All 69 trials were plotted. e, f) First three PCs of PA matrix (e) and first three PCs of LFADS output (f) for all 69 trials for subject S04. g, h) First three PCs of PA matrix (g) and first three PCs of LFADS output (h) for all 78 trials for subject S05. For the first three PCs plot, each point on the trajectory is representing the projection of original data onto the first three PCs at one specific time point.

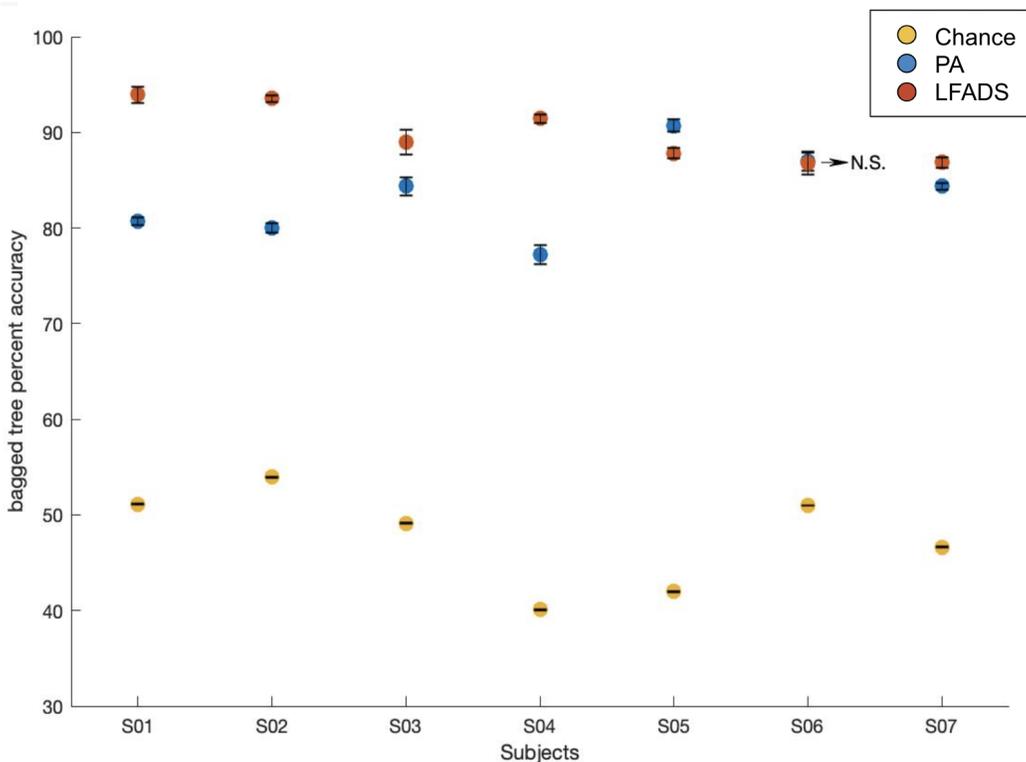


Figure 7: Bagged trees decoding performance. The chance level decoding, performed by using shuffled labels, and actual classification results using either PA matrix or LFADS as input was indicated by yellow, blue, red markers, respectively. Error bars represent standard error of the mean. For all subjects except S06, LFADS performed significantly better than PA. For subject S05, the LFADS decoding was significantly worse than that of PA. There was no statistical difference between LFADS and PA performance for subject S06. S06 is also the subject with epilepsy, who was monitored using an array with fewer contacts (32 channels) than for other subjects (64 channels). Decoding from both the PA matrix and LFADS output resulted in performance that exceeds chance levels.

## Discussion

We evaluated LFADS performance in extracting low-dimensional dynamics and de-noising ECoG through performing a discrete behavioral state classification. When using bagged trees classifiers with LFADS-denoised ECoG activity, the decoding performance of 5 out of 7 subjects improved significantly, with an average of  $89.9\% \pm 1.2\%$  that was better than PA matrix with an average of  $83.47\% \pm 1.7\%$ . Traditionally, ECoG studies rely on a selected range of features such as the sets of channels and range of spectral frequency bands to be used to reach optimal decoding performance on subject intention. These features that are selected usually vary across subjects, task, cortical areas and require experts in this area to evaluate them manually for optimal performance (Pailla et al., 2018). For instance, the channels to be included at the end were not simply selected based on the anatomical position but through massive evaluation of different combinations of frequency bands and channels (i.e. for one channel, its low frequency feature might be responsive to behavioral variables, but its high frequency feature is not). In our case, we selected channels purely according to the brain area they covered (M1, PM channels) and used the same general range of low frequency feature and high frequency feature across all channels. Our results suggest that LFADS can effectively denoise ECoG and would at least allow us to reduce the effort that are required for manual selection in traditional ECoG studies by modeling the neural population dynamics. This application is not limited to ECoG data only. Other recording methods with low spatial resolution, such as stereo-electroencephalography (SEEG), may also be modeled by LFADS and potentially lead to a higher performance in decoding analysis. Combining the advantage of a larger coverage for ECoG array

and extracting low-dimensional neural population dynamics of LFADS, ECoG also has the potential to be used to study other basic science questions for systems neuroscience, which could then in return benefit BMI development in the future.

Even though for one of the subjects (S05), LFADS-denoised ECoG has a lower decoding accuracy than PA matrix, modeling neural population dynamics using LFADS significantly improve the overall performance of ECoG ( $p=2^{-31}$ ) and the performance of LFADS-denoised ECoG could be further improved in several ways. First, the current recording was a relatively short one (~300s), it is possible that we can obtain higher decoding accuracy by training LFADS with longer recording durations. Second, as previously mentioned, we can modify LFADS to combine data from across days to generate a more consistent model over time for the neural population activity. Another thing to notice is that the chance level was higher than a purely random classification of three states, which should be 33% if each state is equally likely to occur. This was likely due to the disproportion of the three behavioral states because the force state occurs more often than movement even after we limited the amount of force data we used (Figure 3). The classification performance for a different experimental task with more balanced behavioral states may also be evaluated in the future to provide a thorough assessment.

As the ability to grasp and manipulate objects is essential to our daily lives, the restoration of hand and arm function was ranked first among other loss motor functions in the patients who have limited hand functions (Xie et al., 2018). Developing a BMI that allows patients to execute object interaction movement is thus definitely a priority goal to provide a better standard of life for these subjects. Even though ECoG generally

performed poorly in terms of decoding accuracy when compared to intracortical recording methods, our results demonstrated that LFADS could help ameliorate this disadvantage as the ability to differentiate the states of moving and force is essential to develop functional BMIs that could be used for daily activities (Downey et al., 2018). Because in our results, ECoG could be used to accurately determine the behavioral states, ECoG has the potential to be used for hand control, which was not likely to be achieved when only using traditional spectral feature extraction.

Although ECoG is invasive, survey results concluded that spinal cord injury (SCI) patients preferred wireless brain implants than wired EEG caps, indicating that in some cases, the convenience and cosmetic appearance of a system is more strongly weighed than surgical concerns (Blabe et al, 2015). Because non-invasive interfaces, like EEG, result in relatively poor decoding of neural activity, they are not suitable to restore behaviors like grasping, which requires fine control (Krucoff et al., 2016). The leading intracortical recording method, the Utah array, has high decoding performance, but only has been implanted chronically in a small set of patients. Most of these patients had no remaining motor function, and therefore, little clinical safety data is available to demonstrate that implanting intracortical arrays does not interfere with residual motor function. In contrast, ECoG recording is widely used during clinical treatment for epilepsy. On the timescale of weeks, ECoG has a well-proven safety record, which is ideal for rehabilitative applications, and may hold promise for patients that have some remaining motor function.

As mentioned early, hand grasping is so essential that developing a rehabilitative BMI that utilizes the unique opportunity ECoG and LFADS provided for hand should be

the priority and is considered as one of the future directions for this project. Because rehabilitative BMIs usually facilitate motor function restoration through pairing patients' movement intention with precisely-timed execution and feedback of the task to activate plasticity mechanisms, a real-time decoder for the BMI needs to be developed and tested first. Individuals who have chronic epilepsy are usually monitored with ECoG recording for days or weeks in an epilepsy monitoring unit, which will provide us with a unique opportunity to build a pre-trained LFADS model using the first days of data and then apply LFADS in real-time to evaluate its performance in closed-loop BMI setting. We can also evaluate the performance of LFADS in improving ECoG decoding performance for more complex behaviors.

## References

Fielding, A. (n.d.). Clustering and Classification methods for Biologists. Retrieved March 21, 2019 from <http://www.alanfielding.co.uk/multivar/cart.htm>

Amoeba (2018), Making sense of principal component analysis, eigenvectors & eigenvalues, Retrieved March 20, 2019 from URL (version: 2018-03-14):  
<https://stats.stackexchange.com/q/140579>

Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., Walter, B. L., Sweet, J. A., Hoyen, H. A., Keith, M. W., Peckham, P. H., Simeral, J. D., Donoghue, J. P., Hochberg, L. R., & Kirsch, R. F. (2017). Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: A proof-of-concept demonstration. *The Lancet*, 389(10081), 1821-1830.  
doi:10.1016/s0140-6736(17)30601-3

Blabe, C. H., Gilja, V., Chestek, C. A., Shenoy, K. V., Anderson, K. D., & Henderson, J. M. (2015). Assessment of brain-machine interfaces from the perspective of people with paralysis. *Journal of neural engineering*, 12(4), 043002.

Cunningham, J. P., & Byron, M. Y. (2014). Dimensionality reduction for large-scale neural recordings. *Nature neuroscience*, 17(11), 1500.

Downey, J. E., Weiss, J. M., Flesher, S. N., Thumser, Z. C., Marasco, P. D., Boninger, M. L., Gaunt, R. A. & Collinger, J. L. (2018). Implicit Grasp Force Representation in Human Motor Cortical Recordings. *Frontiers in neuroscience*, 12.

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., Van Der Smagt, P. & Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398), 372-375. doi:10.1038/nature11076

Krucoff, M. O., Rahimpour, S., Slutzky, M. W., Edgerton, V. R., & Turner, D. A. (2016). Enhancing nervous system recovery through neurobiologics, neural interface training, and neurorehabilitation. *Frontiers in neuroscience*, 10, 584.

Liang, N., & Bougrain, L. (2012). Decoding Finger Flexion from Band-Specific ECoG Signals in Humans. *Frontiers in Neuroscience*, 6. doi:10.3389/fnins.2012.00091

Nurse, E., Mashford, B. S., Yepes, A. J., Kiral-Kornek, I., Harrer, S., & Freestone, D. R. (2016). Decoding EEG and LFP signals using deep learning. *Proceedings of the ACM International Conference on Computing Frontiers - CF 16*.  
doi:10.1145/2903150.2903159

Pailla, T., Miller, K. J., & Gilja, V. (2018). Autoencoders for learning template spectrograms in electrocorticographic signals. *Journal of neural engineering*.

Pandarinath, C., Nuyujukian, P., Blabe, C. H., Sorice, B. L., Saab, J., Willett, F. R., Hochberg, L. R., Shenoy, K. V. & Henderson, J. M. (2017). High performance communication by people with paralysis using an intracortical brain-computer interface. *ELife*, 6. doi:10.7554/elifelife.18554

Pandarinath C, O'Shea D. J., Collins J., Jozefowicz R., Stavisky S. D., Kao J. C., Trautmann E. M., Kaufman M. T. , Ryu S. I., Hochberg L.R., Henderson J.M., Shenoy KV, Abbott LF, Sussillo D.(2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*.

Vansteensel, M. J., Pels, E. G., Bleichner, M. G., Branco, M. P., Denison, T., Freudenburg, Z. V., Gosselaar, P., Leinders, S., Ottens, T.H. & Van Rijen, P. C. (2016). Fully implanted brain–computer interface in a locked-in patient with ALS. *New England Journal of Medicine*, 375(21), 2060-2066.

Vogel, J., Haddadin, S., Jarosiewicz, B., Simeral, J., Bacher, D., Hochberg, L., Donoghue, J. P. & Smagt, P. V. (2015). An assistive decision-and-control architecture for force-sensitive hand–arm systems driven by human–machine interfaces. *The International Journal of Robotics Research*, 34(6), 763-780.  
doi:10.1177/0278364914561535

Wang, W., Collinger, J. L., Degenhart, A. D., Tyler-Kabara, E. C., Schwartz, A. B., Moran, D. W., Weber, D.J., Wodliner, B., Vinjamuri, R. K., Ashmore, R. C. & Kelly, J. W.

(2013). An electrocorticographic brain interface in an individual with tetraplegia. *PloS one*, 8(2), e55344.

Schalk, G., Miller, K. J., Anderson, N. R., Wilson, J. A., Smyth, M. D., Ojemann, J. G., Moran, D. W., Wolpan, J. R. & Leuthardt, E. C. (2008). Two-dimensional movement control using electrocorticographic signals in humans. *Journal of Neural Engineering*, 5(1), 75-84. doi:10.1088/1741-2560/5/1/008

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, 36, 337-359.

Suykens, J. A., & Vandewalle, J. (1999). Least squares support vector machine classifiers. *Neural processing letters*, 9(3), 293-300.

Xie, Z., Schwartz, O., & Prasad, A. (2018). Decoding of finger trajectory from ECoG using deep learning. *Journal of neural engineering*, 15(3), 036009.

Yanagisawa, T., Hirata, M., Saitoh, Y., Kishima, H., Matsushita, K., Goto, T., Fukuma, R., Yokoi, H., Kamitani, Y. & Yoshimine, T. (2011). Electrocorticographic control of a prosthetic arm in paralyzed patients. *Annals of Neurology*, 71(3), 353-361. doi:10.1002/ana.22613