

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Changes in neural activity during brain-machine interface control: from information encoding and connectivity to local field potentials

Permalink

<https://escholarship.org/uc/item/1tn846g9>

Author

So, Kelvin

Publication Date

2013

Peer reviewed|Thesis/dissertation

**Changes in Neural Activity during Brain-Machine Interface Control:
from Information Encoding and Connectivity to Local Field Potentials**

by

Kelvin So

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences
and the Designated Emphasis

in

Communication, Computation and Statistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jose Carmena, Chair
Professor Michel Maharbiz
Professor Bruno Olshausen

Fall 2013

**Changes in Neural Activity during Brain-Machine Interface Control:
from Information Encoding and Connectivity to Local Field Potentials**

Copyright 2013
by
Kelvin So

Abstract

Changes in Neural Activity during Brain-Machine Interface Control:
from Information Encoding and Connectivity to Local Field Potentials

by

Kelvin So

Doctor of Philosophy in Electrical Engineering and Computer Sciences
and the Designated Emphasis
in
Communication, Computation and Statistics

University of California, Berkeley

Professor Jose Carmena, Chair

Brain-machine interfaces (BMIs) are an emerging field of research that seeks to interface the brain directly with prosthetic devices. Researchers have exhibited compelling proof-of-concept experiments for BMIs, with multiple groups showing impressive demonstrations of non-human primates and humans using neural activity to control external devices. However, significant improvements in both reliability and performance (achieving control and dexterity comparable to natural movements) are still needed for BMIs to be clinically viable.

One of the fundamental challenges in developing BMIs is designing the decoding algorithm (or decoder), which translates the recorded neural activity into a control signal for an external device. Traditionally, decoders were often designed to mimic the natural motor system; decoders would be fitted from neural activity recorded while the subject performed natural movements, so that they could predict or decode the movement parameters from the neural activity. However, such decoders do not take into account an important factor that comes into play when the decoder is used in a real-time BMI system: the feedback that the user receives. BMIs are inherently “closed-loop” systems; users generate neural activity that is transformed into a control signal (via the decoder), and the user receives feedback (e.g. visually),

which in turn will affect the neural activity the user generates next. This feedback cycle that occurs during BMI operation can have profound impact on the neurons and the neural activity that is generated.

The first two chapters of this thesis explore differences at the neuronal level between BMI operation and natural reaching. We leverage techniques and insights from network information theory to characterize information encoding in neural ensembles and neural network structure in the two conditions, BMI control versus natural reaching. Using measures of redundancy derived from information theory, we found that neurons directly incorporated in the control of the BMI contained higher and more redundant information regarding the target, compared to surrounding neurons. Using a novel information-theoretic approach for inferring functional connectivity, we found a consolidation of the network connections between the control neurons of the BMI that parallels the learning curve, suggesting the formation of neural circuitry specific to BMI operation. Together, these findings indicate that fundamental changes are occurring at the neural level, which may cause the observed neural activity during BMI control to differ dramatically from during natural reaching.

The subsequent chapter of this thesis describes our progress in building an adaptive BMI system that takes into account the changes in neural activity that occurs in closed-loop, as well as our efforts in improving the long-term reliability of the system. We demonstrate 2-D BMI control using local field potentials as input with closed-loop decoder adaptation (CLDA) to fit the decoder parameters. CLDA is a method of adapting the decoder's parameters that aims to make the decoder's output accurately reflect the subject's intended movements during closed-loop BMI operation. Local field potentials (LFPs) are measurements of the extracellular electric field that reflect contributions from a population of cells. LFPs are an attractive alternative to single- and multi-unit activity in BMIs because they are more robust to signal degradation over time. In addition, our study sheds light on the features of LFP that can potentially be used as input to BMI, with implications for channel and feature selection.

Acknowledgments

First and foremost, I would like to thank my two advisors, Jose Carmena and Michael Gastpar, for their continual motivation, guidance and wisdom throughout my graduate career. Without them, my Ph.D would not be the same.

To Jose, thank you for believing in me from the start, convincing me to persevere when I was tempted to leave and giving me the freedom to choose what I want to do.

To Michael, thank you for your guidance both within research and on my career as a whole, even though that meant staying up late to entertain Skype calls.

While most graduate students would engage in non-research related activities to keep them “sane”, for me it is actually the reverse. Research work was always something I could be reliably confident on, and often, the reason to be excited to wake up and get to work in the morning. And so it is the motivation and encouragement in my work that I am truly grateful to Jose and Michael for providing throughout the years.

I also want to thank Bruno Olshausen and Michel Maharbiz for serving on my committee.

Secondly, I am grateful to my fellow lab-mates, both past and present, for their helpful input. In particular, I want to acknowledge Simon Overduin for training me on primate work, Siddharth Dangi for the smoothest collaboration I’ve ever had, and Vivek Athalye, for making our office interesting. In addition, I would like to thank my friends outside of research for making my time at Berkeley enjoyable, and in particular, Alan Malek, for giving me reasons to not do work.

Last but not least, I want to thank my parents and my brother for their continual support in everything else.

Contents

Acknowledgments	i
1 Introduction	1
2 Changes in information encoding	5
2.1 Redundant information encoding during BMI control	5
2.2 Materials and methods	6
2.2.1 Electrophysiology and surgery	6
2.2.2 Experimental setup and training	6
2.2.3 BMI paradigm	7
2.2.4 Data analysis	7
2.3 Results	10
2.4 Discussion	12
2.4.1 Redundancy in information encoded in M1	14
2.4.2 Implications for brain-machine interfaces	14
3 Changes in connectivity	16
3.1 Functional connectivity of neural ensembles during BMI control . . .	16
3.2 Directed information	18
3.2.1 Estimating directed information via point process GLM	18
3.2.2 Significance and confidence intervals	20
3.3 Simulation setup and results	21
3.3.1 Metrics for comparison	21

3.3.2	Two-neuron simulations	22
3.3.3	Three-neuron simulations	23
3.4	Rodent experimental data	25
3.4.1	Experimental setup	25
3.4.2	Data analysis	26
3.5	Primate experimental data	27
3.5.1	Experimental setup	27
3.5.2	Degree of connectivity	28
3.5.3	Stability of connections	32
3.6	Discussion	32
4	Tackling the closed-loop BMI problem	37
4.1	Brain-machine interface control using local field potentials	37
4.2	Materials & methods	39
4.2.1	Electrophysiology and behavioral task	39
4.2.2	Decoder design	39
4.2.3	Closed-loop decoder adaptation	41
4.2.4	Performance evaluation	42
4.3	Results	43
4.3.1	Testing different frequency bands	47
4.3.2	Testing different numbers of LFP channels	49
4.4	Discussion	52
4.4.1	Comparison of BMI performance	54
4.4.2	Practical considerations for neuroprosthetics	56
4.5	Multi-tasking using LFP BMI	56
4.5.1	Task setup and training	56
4.5.2	Results	58
4.5.3	Discussion	61
5	Conclusion	62
5.1	Thesis contributions	62
5.2	Future work	63

Bibliography **63**

Appendix **71**

 A.1 Raised cosine basis functions 71

 A.2 Sparse generalized Laguerre-Volterra model (SGLVM) 71

 A.3 Calculating modulation depth 72

Chapter 1

Introduction

The principle behind brain-machine interfaces (BMIs) is to interface neural activity directly from the brain to an external prosthetic device such as a computer cursor or robotic arm. It has the potential to restore communication and motor control to patients with severe motor disabilities such as amyotrophic lateral sclerosis, spinal cord injuries, and brainstem stroke, by establishing new output pathways for the brain. Although the concept of BMI has existed for decades, only recently have various technical improvements in hardware allowed development in BMI feasible. Advances in electrode arrays have enabled simultaneous stable recordings of neuron ensembles over extensive periods of time. As a result, a number of recent studies have demonstrated control of an external device via a BMI system (see [1] for review).

A BMI system (Figure 1.1) generally consists of a decoding algorithm, or *decoder*, which translates the recorded neural activity into a control signal, and an external actuator, which is typically visible to the user in order to provide visual feedback to the user. One of the biggest challenges in designing a BMI system is determining a good decoder that is intuitive to the user and accurately reflects the user's intentions. In past studies, decoders were often fitted in an "open-loop" fashion, whereby the user was first asked to perform a series of motor actions while neural activity was simultaneously recorded. A decoder was then designed to best replicate the motor action based on the neural activity recorded. Since the neural activity generated is unaffected by the output of the decoder, there is no feedback, and so this approach is considered open-loop. While the problem of decoder design in an open-loop setting can reveal insights on the how natural motor actions are represented in the neural activity, it is a fundamentally different problem than the decoder design problem in a BMI system. During BMI usage, the user's neural activity is affected by the real-time feedback of the decoder's output, forming a closed-loop, and resulting in neural activity that is different than activity generated during natural movement. This thesis will examine the changes in the user's neural activity during BMI operation,

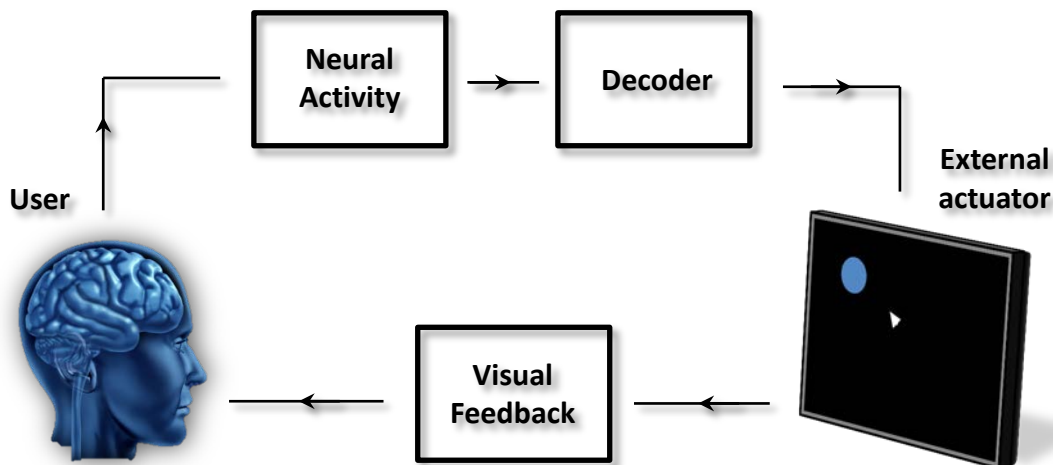


Figure 1.1: General architecture of a brain-machine interface. Neural activity is recorded directly from the brain, which is then translated by the decoder into a control signal for an external actuator. Some form of feedback, typically visual, is given the user, forming closed-loop control.

followed by recent progress made on the engineering front in implementing algorithms that can address these differences.

Overview

The first chapter of this thesis sheds light on the differences in the information encoded in the neural activity and circuitry during BMI operation as compared to natural motor control. We examine the amount of information carried by an ensemble of neurons and quantify the redundancy of movement-related information encoded in the macaque primary motor cortex (M1) by using concepts from information theory. In communication networks, information is often encoded with a certain degree of redundancy across channels to provide robustness against noise during transmission. For a network of neurons, it is conceivable that information is redundantly encoded since many neurons share common afferents [2]. Indeed, past studies have suggested that correlations (which is an indirect measure of redundancy) in neural activity can be important for encoding information [3]. During natural arm movement, we found neurons in contralateral M1 to contain higher and more redundant information about target direction than ipsilateral M1 neurons, consistent with the laterality of movement control. However, during BMI control, we found that the M1 neurons directly incorporated into the BMI (‘direct’ neurons) contained the highest and most redundant target information, even if the set of direct neurons were from the ipsilateral side. This result suggests a change in the information encoding during

BMI operation.

Next, we investigate the changes in connectivity of the neural circuitry that arises during BMI operation and how that differs from natural motor control. Neurons in the brain form highly complex networks through synaptic connections, and these connections are important in BMI operation as new synaptic connections are likely to form during learning. Traditionally, the connectivity between neurons has been explored using methods such as correlations. However, such methods do not contain any notion of directionality. Recently, an information-theoretic approach based on directed information theory has been proposed as a way to infer the direction of influence for intracortical recordings [4]. Applying this method to neural activity recorded during BMI operation, we found striking differences in the structure of the neural circuitry as compared to the structure seen in natural motor control. In addition, we found that the connection strengths in the neuronal ensemble consolidated as the subject became proficient at using BMI. Overall, these results demonstrate the changes that occur during BMI operation, and also the potential of using directed information as a means for inferring functional connectivity.

The subsequent chapter discusses new experimental results on applying closed-loop decoder adaptation (CLDA) to a BMI system utilizing local field potentials (LFPs). Closed-loop decoder adaptation is a procedure for updating the decoder during BMI operation in order to address the changes that occur in the neural activity while using BMI. This method has been used effectively in updating decoders utilizing spike activity as the control signal [5–8]. However, an increasing number of studies have shown the utility of LFPs for decoding motor related signals [9–11]. Local field potentials are measurements of the extracellular electric field that reflect contributions from a population of cells, and are considered to be a more robust and stable signal than single unit spiking activity. Currently, it is unclear how well different LFP frequencies can serve as features for continuous, closed-loop BMI control. Here, we demonstrate 2-D continuous LFP-based BMI control using CLDA, which adapts decoder parameters to subject-specific LFP feature modulations during BMI control. We trained two macaque monkeys to control a 2-D cursor in a center-out task by modulating LFP power in the 0–150 Hz range. While both monkeys attained control, they used different strategies involving different frequency bands. One monkey primarily utilized the low-frequency spectrum (0–80 Hz), which was highly correlated between channels, and obtained proficient performance even with a single channel. In contrast, the other monkey relied more on higher frequencies (80–150 Hz), which were less correlated between channels, and had greater difficulty with control as the number of channels decreased. We then restricted the monkeys to use only various sub-ranges (0–40 Hz, 40–80 Hz, and 80–150 Hz) of the 0–150 Hz band. Interestingly, although both monkeys performed better with some sub-ranges than others, they were able to achieve BMI control with all sub-ranges after decoder adaptation,

demonstrating broad flexibility in the frequencies that could potentially be used for LFP-based BMI control. These results demonstrate proficient, continuous BMI control using LFPs and provide insight into the subject-specific spectral patterns of LFP activity modulated during control.

Chapter 2

Changes in information encoding

In this chapter, we explore differences in the information encoding of the neurons in the primary motor cortex (M1) as the subject learns to control the BMI.

2.1 Redundant information encoding during BMI control

Redundancy is an integral part of reliable distributed information processing by providing robustness against component failure and noisy communication channels. For a neuronal ensemble, redundant information encoding refers to a coding scheme in which each neuron is encoding a similar signal. Using this scheme, firing rate information can be reliably transmitted in a short time interval despite variability in interspike intervals [2, 12, 13]. In the motor system where the firing rate of neurons carries information about the motor action, a redundant encoding scheme at the ensemble level may be a possible solution to reliably transmit firing rate information to produce consistent motor action.

Redundant encoding in the motor system has mainly been studied in the context of correlations. Experimental work in primates has shown that correlations in firing activity can be important for encoding movement parameters [14–16]. This is supported by theoretical work demonstrating that correlations in the variability of firing rates can improve the accuracy of information encoding [3, 17–19]. More recently, Narayanan et al. [20] extended the analysis to larger ensemble sizes and found largely redundant information encoding in the rat motor cortex during a reaction time task.

In this study, we examine the redundancy of movement-related information in the macaque primary motor cortex (M1) during natural (MC) and neuroprosthetic control (BC) via a brain-machine interface (BMI). A BMI paradigm enables us to exam-

ine the ensemble that directly controls movement, since for natural motor control, it is difficult to identify the corticomotoneuronal cells that connect monosynaptically to motoneurons [21]. In both experiments, we found that ensembles most associated to movement control showed highest redundancy in encoding target information. For natural arm control, this corresponds to the contralateral M1 (left M1). Interestingly, when we retrained the BMI to use ipsilateral M1 activity, we found that these neurons were more redundant and contained higher information than contralateral M1 neurons, even though ensembles from this hemisphere were previously less redundant during natural arm movement. This result suggests that information encoding in the neurons are different during BC than MC. Moreover, these results demonstrate the ubiquitous nature of redundant information encoding in the motor cortex and provide evidence of the contribution of redundancy to proficient motor control.

2.2 Materials and methods

2.2.1 Electrophysiology and surgery

The subjects were implanted with arrays of 8x8 Teflon-coated tungsten microelectrodes ($35\mu\text{m}$ in diameter, $500\mu\text{m}$ separation between microwires). Both subjects were implanted bilaterally in the arm area of the primary motor cortex (M1), positioned at a depth of 3mm targeting layer 5 pyramidal neurons. Localization of target areas were performed using stereotactic coordinates of the rhesus brain. Intraoperative monitoring of spike activity was used to guide the depth of electrode placement. All procedures were conducted in compliance with the regulations of the Animal Care and Use Committee at the University of California, Berkeley. Single unit activity was recorded with a Multichannel Acquisition Processor and sorted using an online sorting application (Plexon Inc., Dallas, TX). A subset of stable units was found several months postsurgery whose waveform shape, amplitude and interspike interval distribution exhibited little change from day to day. These stable units were then chosen as input to the BMI decoder.

2.2.2 Experimental setup and training

The monkeys were trained to perform a delayed center-out reaching task using his right arm (manual control, MC) or with a cursor controlled through a brain-machine interface (BC-Contra and BC-Ipsi, see Figure 2.1). The task involved cursor movements from the center towards one of eight targets distributed evenly on a 14cm diameter circle. Target radius was set at 0.75cm. Each trial began with a brief hold period at the center target, followed by a GO cue (center changed color) to

signal the reach towards the target. The monkey was then required to reach and hold briefly (0.2-0.5s) at the target in order to receive a liquid reward. In manual control, reaching was performed using a Kinarm (BKIN Technologies, Kingston, ON) exoskeleton where the monkey’s shoulder and elbow were constrained to move the device on a 2-D plane. During recording, the subjects were placed in a primate chair that permits limb movements and postural adjustments, with a headpost for head restraint. In BC-Contra and BC-Ipsi, spiking activity from the selected (direct) neurons were mapped to kinematic parameters via the BMI “decoder” (see below) to provide continuous control to the cursor. In these sessions, the right arm was restrained.

2.2.3 BMI paradigm

The BMI decoder is a linear regression model where the inputs $X(t)$ are vectors containing the spike counts of each the direct neurons over a number of time bin [22]. The output $Y(t)$ is a vector that contains the estimated shoulder and elbow joint angles. Formally, this linear model can be expressed as:

$$Y(t) = b + \sum_{u=0}^n a(u)X(t - u) + \varepsilon(t) \quad (2.1)$$

where the elements of $X(t)$ represents the spike count of each direct neuron at time t and the columns of $Y(t)$ represents the shoulder and elbow joint angles. The vectors $a(u)$ for $u = 1, \dots, n$ and b are the causal linear filter parameters, and $\varepsilon(t)$ are the residual errors. Here, we used 10 time lags ($n = 10$), with spike counts binned at 100 ms. The linear filter parameters are estimated by using linear regression on 10 min of training data recorded while the monkey performed right arm movements. Cursor coordinate position is then computed from the joint angles to provide real-time cursor control.

2.2.4 Data analysis

We adopted a decoding approach using linear discriminant analysis (LDA) to estimate the information contained in an ensemble [23]. The information considered in this analysis is the target direction contained in M1 neurons during the movement period (~100-800ms for MC and ~100-6000ms for BC) of each successful trial. To compute the information contained in an ensemble about the target direction, we apply the concept of mutual information from information theory. The mutual information between two discrete random variables X and Y is formally defined as:

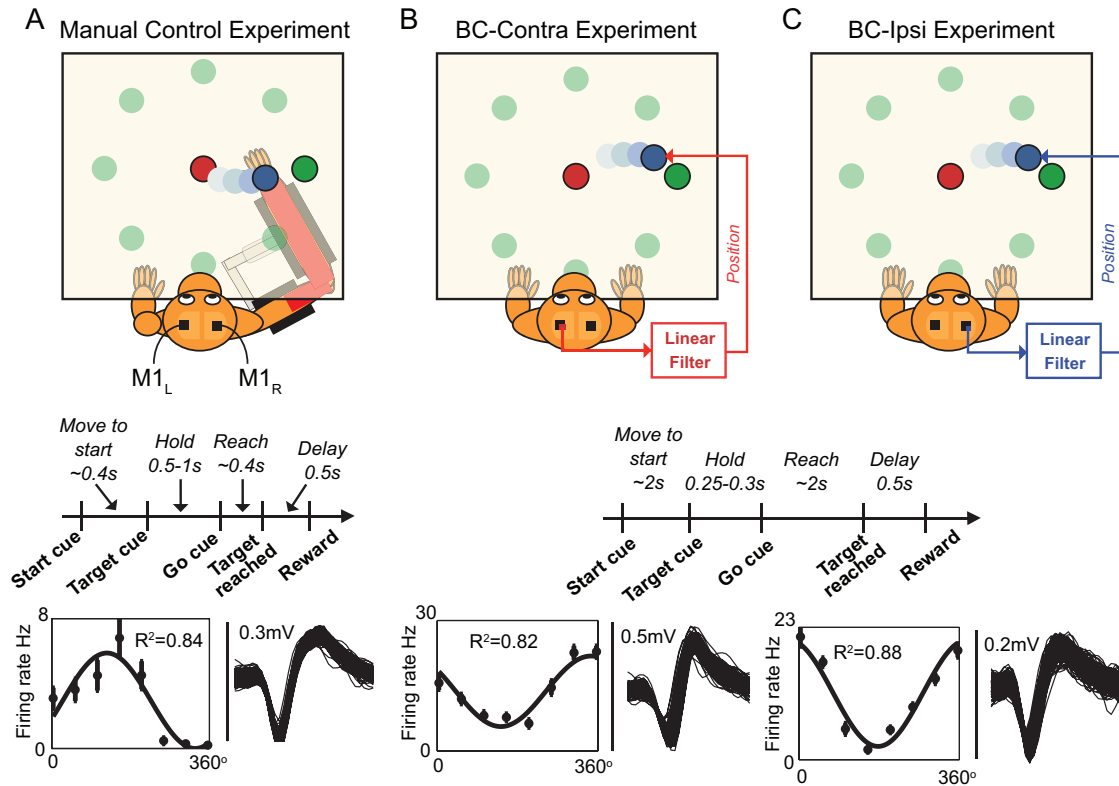


Figure 2.1: Experimental setup for (A) MC, (B) BC-Contra, and (C) BC-Ipsi experiments. The monkey performs a center-out reaching task constrained on a two-dimensional plane. One of eight targets is presented each trial, and liquid reward is given after a successful reach. In BC-Contra/BC-Ipsi, the position of the cursor is determined by filtering the output of the $M1_{L-Direct}/M1_{R-Direct}$ units. (D) Waveform and directional tuning of a typical unit with cosine tuning during each experiment.

$$I(X; Y) = \sum_x \sum_y p_{X,Y}(x, y) \log_2 \left(\frac{p_{X,Y}(x, y)}{p_X(x)p_Y(y)} \right) \quad (2.2)$$

where $p_{X,Y}(x, y)$, $p_X(x)$ and $p_Y(y)$ represent the joint and marginal distributions of X and Y . Conceptually, this measure quantifies the information contained in X about Y . Ideally, we would like to compute the mutual information $I(R; T)$ between the neural responses R and target direction T . However, the distribution of R can be very high dimensional and hard to estimate. Hence, we adopted a decoding approach wherein we would first “decode” the target direction \hat{T} , and then compute the mutual information $I(T; \hat{T})$ from the confusion matrix between the actual and predicted target. In general, the decoding approach can be applied using any statistical classifier. Here, we used linear discriminant analysis, which models the input distribution as a mixture of Gaussians and performs the classification by finding the target that maximizes the posterior distribution. Since LDA performs optimal classification for Gaussian distributed data, we follow the standard approach of first taking square roots of the firing rates such that the resulting quantities can be reasonably well modeled as Gaussians. The mean is estimated from training data for each target separately, with the covariance assumed fixed across targets. We used up to 10 neurons as input since an n -dimensional input requires approximately n^2 number of parameters to be estimated in the covariance matrix. Therefore, the model would not be very accurate for larger ensembles, since the number of trials to train our model is limited (typically <200 trials per session). Once the model is estimated, a separate test data is used to compute the target predictions and mutual information. We employed 5-fold cross-validation where the classifier is trained on 80% of the data and predictions are tested on the remaining 20%, repeated 5 times so that each trial is part of a test set once. Only late-learning sessions in which animals achieves over 90% of successful trials are used for the analysis.

From the data-processing inequality, $I(R; T)$ obtained using this decoding approach represents a lower bound of the true mutual information $I(R; T)$ between the neural responses and the targets. Conceptually, this results from the inability of the classifier to completely extract all target relevant features from the neural responses. Hence, the gap between $I(T; \hat{T})$ and the true value $I(R; T)$ depends on how well the decoder can capture all the relevant features. The particular classifier used here may not be the optimal decoding procedure, nor represent the actual method in which information is encoded in a neural ensemble [24]. However, the decoding approach has been successfully used to quantify relationships between ensemble responses and behavior despite the drawbacks [20, 25]. To assess the effect due to the choice of statistical classifier, we also performed the analysis using learning vector quantization (LVQ) as the statistical classifier. The results obtained from LVQ and LDA were highly correlated ($R^2 = 0.75$, $p < 0.001$).

To quantify the amount of redundancy in an ensemble, we compare the ensemble information with the sum of the individual information of each neuron in the ensemble [26]. Formally, $P_{Ensemble}$ is defined as [20]:

$$P_{Ensemble} = I_{Ensemble} - \sum_i I_{Neuron} \quad (2.3)$$

where $I_{Ensemble}$ is the information contained in the ensemble and I_{Neuron} is the information contained in a single neuron in the ensemble. If $P_{Ensemble}$ is negative, then the total information of the ensemble as a whole is less than the sum of the information of the individual neurons in the ensemble, which implies that the information encoding in individual neurons are redundantly encoded in other neurons. If $P_{Ensemble}$ is positive, the total information of the ensemble is greater than the sum of the individual, and such interactions can be considered synergistic.

2.3 Results

We recorded neural ensemble data from two macaque monkeys performing a center-out reaching task using natural arm control (“Manual Control”, MC, experiment) and brain control (“BC-Contra” and “BC-Ipsi” experiments). The recorded ensembles are primarily classified based on location. In MC, units are grouped as contralateral M1 ($M1_L$) or ipsilateral M1 ($M1_R$). In BC-Contra, we also separated the control units ($M1_{L-Direct}$, since they were chosen from the left hemisphere). Note that this label only applied to the sessions where those units explicitly served as control units for BMI. The remaining units are grouped according to their location: left M1 ($M1_{L-Indirect}$) or right M1 ($M1_{R-Indirect}$). For Monkey P, an additional BC experiment (BC-Ipsi) was performed using control units from the right hemisphere ($M1_{R-Direct}$). The number of units and percentages of directionally tuned units for each neuron group are shown in Table 2.1.

The average information for ensembles of ten units in each neuron group for each experiment is shown in Figure 2.2. For the MC experiment, $M1_L$ ensembles contained the highest mutual information for both animals (t-test, $p = 1 \times 10^{-6}$). For BC-Contra, the mutual information of the $M1_{L-Direct}$ units were found to be the highest (t-test, $p = 3 \times 10^{-4}$), while the $M1_{R-Indirect}$ units contained the least amount of information. Interestingly, for the BC-Ipsi experiment, the $M1_{R-Direct}$ and $M1_{R-Indirect}$ units contained significantly higher information than $M1_{L-Indirect}$ units (t-test, $p = 5 \times 10^{-3}$).

The defining characteristic of redundancy is that information encoded in each neuron is partially encoded in another neuron. Hence, the total information encoded

Experiment	Sessions	Number of units			Percentage of directionally tuned units			Average ensemble information (bits) and accuracy (%) for ensembles of size 10					
		M1 _{Direct}	M1 _L	M1 _R	M1 _{Direct}	M1 _L	M1 _R	M1 _{Direct}	M1 _L	M1 _R			
Monkey P	Manual Control	5	-	83±5	59±4	-	77±5	49±6	Acc	-	53±4	33±3	
									Info	-	1.24±0.09	0.65±0.05	
		9		-	41±5	42±4	-	76±3	56±4	Acc	-	50±4	38±2
										Info	-	1.1±0.1	0.77±0.06
	BC-Contra	5	40±1	43±5	59±5	85±3	77±9	59±8	Acc	48±3	42±4	37±2	
										Info	1.2±0.1	0.9±0.1	0.8±0.1
BC-Ipsi	4	14±1	31±7	26±4	90±6	82±7	90±3	Acc	60±3	44±3	51±3		
									Info	1.6±0.1	1.1±0.1	1.35±0.05	
Monkey R	Manual Control	4	-	54±2	85±1	-	75±4	69±9	Acc	-	45±1	37±4	
									Info	-	1.06±0.03	0.85±0.1	
	BC-Contra	3	10	125±9	-	63±11	49±18	-	Acc	42±3	32±5	-	
									Info	1.2±0.2	0.9±0.2	-	

Table 2.1: Number of units, percentage of directionally tuned units, and ensemble information (Info) and predictive accuracy (Acc) for ensembles of size 10 are listed for each of the different experiments and subjects. The two separate rows for MC experiments refer to different sets of MC sessions performed five months apart. The ensemble information and accuracy remain consistent despite the time lapse.

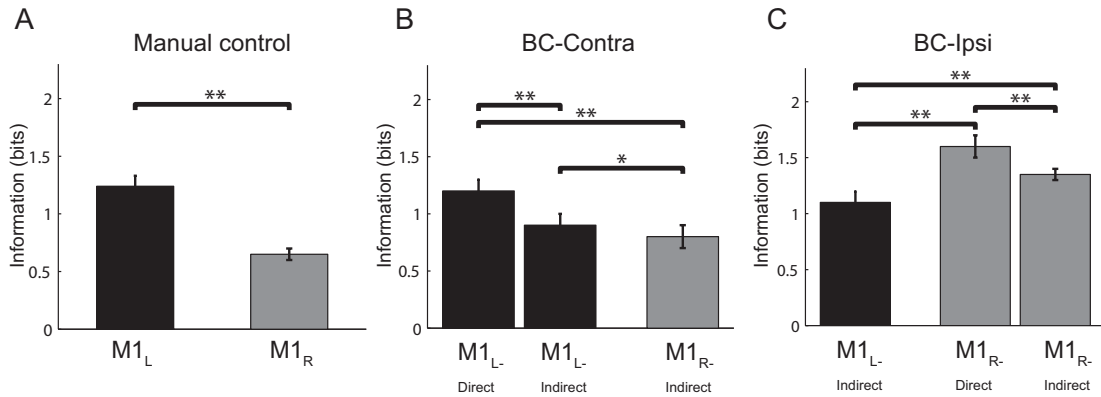


Figure 2.2: Average information contained in ensembles of 10 neurons in each area for (A) MC, (B) BC-Contra and (C) BC-Ipsi experiments. (*) implies $p < 0.05$ and (**) implies $p < 0.01$ using t-test to compare the difference. Error bars represent standard error of the mean across sessions.

in an ensemble should eventually saturate as more neurons are added since the information in additional neurons is already encoded in the ensemble. Conversely, if the information encoded in each neuron is not redundant, namely, if the information encoded in each neuron is independent of the other neurons in the ensemble, the total ensemble information should increase proportionally to the ensemble size. To test this notion of redundancy, we explored the increase in information as the ensemble size grew. For each experiment, we computed the average mutual information for ensemble sizes up to 10 within each of the neuron groups (Figure 2.3A-C). For small ensemble sizes (under 3), all possible combinations of units within each group were considered. For larger ensemble sizes we computed the mutual information for 1000 randomly sampled unique combinations due to computation limitations. For all neuron groups in both subjects, the information encoded in each neuron is not independent, as the extra information provided by each additional unit becomes less at larger ensemble sizes, suggesting redundancy. Secondly, the relative amount of mutual information for the neuron groups remains consistent for all ensemble sizes, which suggests that these results can be extrapolated to larger ensembles.

We quantified the degree of redundancy by calculating the difference between the predictive information of the ensemble and the sum of the individual information of its constituents ($P_{Ensemble}$, see Section 2.2.4, Data analysis) for each ensemble size (Figure 2.3D-F). The information carried in all ensemble sizes across all neuron groups was generally redundant. For MC, M1_L units were most redundant (t-test, $p = 6 \times 10^{-9}$). For the BC-Contra experiment, M1_{L-Direct} units were found to have highest redundancy (t-test, $p = 10^{-3}$). These results were consistent for both subjects. Interestingly, for BC-Ipsi, the M1_{R-Direct} units exhibited highest redundancy, even though during MC, M1_R ensembles had lower redundancy than M1_L ensembles. This was surprising given that cursor control in BC-Ipsi sessions was generated via a decoder originally fitted using ipsilateral M1 activity during natural arm movements. The remaining M1_{R-Indirect} units during BC-Ipsi also showed a large increase in redundancy relative to the M1_{L-Indirect} units when compared to the results from MC.

2.4 Discussion

We examined the redundancy of information encoded in M1 ensembles while monkeys performed a center-out reaching task under two different types of control: natural arm control and neuroprosthetic control. Ensembles most associated to movement control showed highest redundancy and information about the target direction. For natural control, this corresponds to contralateral M1 ensembles. For neuroprosthetic control, this corresponds to the laterality of the neural ensemble used in the decoder.




Figure 2.3: Ensemble information for different ensemble sizes for (A) MC, (B) BC-Contra and (C) BC-Ipsi experiments. P_{Ensemble} for different ensemble sizes for (D) MC, (E) BC-Contra and (F) BC-Ipsi experiments. The plots shown represent data from Monkey P averaged across all sessions (9 sessions in MC, 5 in BC-Contra and 4 in BC-Ipsi). Error bars represent standard error of the mean across sessions. Both analyses indicate a redundant encoding of information across all neuron groups.

2.4.1 Redundancy in information encoded in M1

Our results revealed highly redundant levels of target information in M1 ensembles. This result is consistent with the wealth of correlations that M1 neurons have with motor parameters that have been observed in past studies. These studies have found M1 neurons that encode hand position, velocity, acceleration, target direction, movement trajectory, shoulder/elbow angles, muscle coactivation and visual target [27–31]. For constrained reaching tasks (e.g. across the horizontal plane), many of these parameters are tightly coupled. Hence, it is reasonable for M1 neuron ensembles to contain redundant information about the target, even if the individual neurons encode different motor parameters. What are the implications of redundant information encoding? For a communication system, redundancy in the information transmitted across multiple channels provides robustness against degradations in the transmission channel, allowing for more reliable decoding of the transmitted message at a cost of reduced channel capacity. For a neural system, Barlow [32] argued that a non-redundant, distributed encoding scheme was less efficient than a redundant, sparse encoding for learning sensory input probabilities. In the motor system, a number of studies have demonstrated evidence of increased redundancy during learning. Nudo et al. [33] found a larger area of digit representation in contralateral M1 of squirrel monkeys after learning a small object retrieval task. An increase in digit representation indicated that more neurons were mapped to digit movement, suggesting higher redundancy in the motor map. A similar increase in wrist representation associated with learning a novel reaching task was also found in the rodent motor cortex [34]. These results corroborate with our finding that proficient motor control is associated with higher redundancy encoding target direction. The results of the BMI experiments show that redundancy in information encoding in M1 occurs even in the absence of arm movement. Previous studies have shown that M1 neuron firing activity can be modulated by higher cortical processes such as decision making and visual direction of target [35], unrelated to muscle activation. Hence, redundancy in information encoding in the absence of overt arm movement suggests independence from the nature of the actuator being controlled, i.e. musculoskeletal plant vs. computer cursor.

2.4.2 Implications for brain-machine interfaces

One prominent feature of a redundant information encoding scheme is the robustness to channel quality degradations. Therefore, redundancy in neural encoding suggests that the brain can potentially adapt to control a BMI using fewer direct neurons since the information is encoded across many neurons. Indeed, different BMI studies have reported successful BMI operation using vastly different number of units. Here, 10-41

units had been used; Jarosiewicz et al. [36] reported using 23-60 units, while Serruya et al. [37] used 7-30 units. Although these numbers can depend on task difficulty, the large range reported by each group points toward a redundant information encoding scheme that facilitates the brain's ability to adapt to BMI operation using a large range of neurons. Ultimately, the optimal number of neurons needed for proficient control of a prosthetic device will also depend on a number of other factors such as reliability of the neural interface. However, the redundancy of information found in our results suggests that in the event of neuron loss, the brain can adapt to using fewer neurons over time.

Chapter 3

Changes in connectivity

We continue our analysis of the differences at the neuronal level by investigating the changes in the connectivity properties of the underlying neural network using a novel approach known as directed information.

3.1 Functional connectivity of neural ensembles during BMI control

Changes in the circuitry and connectivity of neural ensembles are known to occur after learning new motor skills [38]. Traditionally, connectivity in the brain has been explored from two different angles: anatomical connectivity and functional connectivity. Anatomical connectivity refers to the structural layout of synapses, measured using staining and fiber tracing techniques [39]. Functional connectivity refers to the temporal correlations between neurophysiological events, such as action potentials or hemodynamic response, regardless of the actual anatomical connections, even though both measures of connectivity are closely related [40]. For in situ recordings, only functional connectivity can be measured feasibly due to the difficulty in measuring anatomical connectivity in vivo. In the past few decades, a number of statistical techniques have been developed for inferring functional connectivity for intracortical recordings, electroencephalography (EEG) and functional magnetic resonance imaging (fMRI). These include cross-correlation [41–43], spectral coherence [44, 45] and mutual information [46]. However, these methods do not characterize the direction that one neuron or ensemble exerts on another. Methods that infer directionality, such as dynamic causal modeling [47, 48], directed transfer function [49] and structural equation modeling [50], have largely been limited to EEG and fMRI studies. More generally, directionality can also be inferred by performing cross-correlation at

various time lags and determining causality depending on the time lag that results in the highest correlation, although few studies to date have employed this method. This problem was recently addressed by Quinn et al [4] using an information-theoretic measure known as directed information, which is a promising approach that is consistent with the concept of Granger causality. Other recent methods [51, 52] assess connectivity with likelihood ratios, which is based on the same concept as directed information. However, it is still unclear whether any additional insight can be gained beyond conventional nonparametric methods (i.e. correlations) that simply look at second-order statistics. Here, we first assess directed information under both simulated and experimental data sets. We describe a practical approach to estimating the directed information and show that this procedure results in higher accuracy than correlations in detecting simulated connections between spike trains. We further test the generality of our estimation method in rodent and primate experimental data. Applying directed information to stimulation experiments in rats, we demonstrate the method’s ability to infer causal influences in experimental data and illustrate how directed information can also provide accurate estimates of the conduction delay in connections between the primary motor cortex and the dorsolateral striatum.

In recordings from a macaque monkey performing a center-out reaching task, the procedure reveals striking relationships between properties of the neural network and behaviour. Specifically, we observe a strong relationship between a neuron’s modulation depth and the number of connections from the neuron when the primate performs the task under neuro-prosthetic control (i.e. movement controlled directly via neural activity in the absence of physical movement using a BMI), but a much weaker relationship when the task is performed under manual control (i.e. using his natural arm). The number of connections a node in a network makes (also known as degree centrality) is a commonly used measure in graph theory for evaluating the relative importance of the nodes to the network. For example, in computer server networks, nodes with many connections indicate network hubs that are important for routing network traffic. For a biological neural network, neurons with high degree centrality are presumably important to the network since their activity is related to many other neurons. Our analysis using directed information reveals an interesting relation between a neuron’s importance to the *internal network* (i.e. neural network) and its importance to *external behaviour* (i.e. target reaching). The procedure also identified a consolidation of the network connectivity within the ensemble that closely parallels the performance of neuro-prosthetic control and known changes in the ensemble’s tuning properties [22].

3.2 Directed information

Directed information is a modified version of the well-known mutual information that better captures the capacity of a communication channel in the presence of feedback [53]. Conceptually, directed information is based on the same principle as Granger causality by measuring how much better we can predict a process Y if past information of a process X is given. This metric is particularly attractive for analyzing connectivity in neural ensemble data for two reasons. Firstly, directed information is designed to characterize channels with feedback and since a neuron’s spiking activity is strongly influenced by its past spiking activity, a neuron can be interpreted as a feedback system transmitting to another neuron. Hence, directed information is well suited for capturing connectivity between neurons. Secondly, directed information measures the effect of having past information of another process. This means that if the connection between two neurons X and Y is unidirectional, as in the activity of neuron Y is affected by the past activity of neuron X , but X is not affected in a causal fashion by Y , then directed information will only capture the effect of X on to Y , but not the reverse. Moreover, unlike the definition of Granger causality proposed by Granger [54], directed information does not restrict the effect of past information to follow a linear, autoregressive model. Formally, the directed information between two stochastic processes X and Y is defined as:

$$I(X^N \rightarrow Y^N) \triangleq E \left[\sum_{i=1}^N \log \frac{P_{Y_i|Y_1^{i-1}, X_1^i}(Y_i|Y_1^{i-1}, X_1^i)}{P_{Y_i|Y_1^{i-1}}(Y_i|Y_1^{i-1})} \right] \quad (3.1)$$

where Y_1^n represents the process Y from time points 1 to n , $[Y_1, \dots, Y_n]$, and X_1^n represents the process X from 1 to n , $[X_1, \dots, X_n]$. Since this quantity scales with N , we define a normalized quantity:

$$\mathcal{I}(X \rightarrow Y) \triangleq \lim_{N \rightarrow \infty} \frac{1}{N} I(X^N \rightarrow Y^N) \quad (3.2)$$

Assuming that the processes X and Y are finite-memory Markov chains that are stationary and ergodic, it can be shown that a sample average of the log-likelihood ratio will converge almost surely to the normalized directed information. All subsequent computations will use the normalized directed information.

3.2.1 Estimating directed information via point process GLM

We adopted a parametric approach to estimate directed information from data. Specifically, we suppose that spike trains are generated via a point process generalized linear model (GLM). The point process GLM is an increasingly popular

approach to modeling the neuron spiking probabilities [4, 51, 55, 56]. Under this model, the conditional intensity function is explicitly modeled as a function of the past spiking activity (intrinsic history), the spiking history of another neuron (extrinsic history) and any other relevant covariates (not shown in (3.3)). Specifically, we model the spike train as a Poisson process where the conditional intensity function is related linearly to the covariates via the log function:

$$\log \lambda_i = \alpha_0 + \sum_{j=1}^J \alpha_j Y_{i-j} + \sum_{k=1}^K \beta_k X_{i-(k-1)} \quad (3.3)$$

$$P_{Y_i|Y_1^{i-1}, X_1^i} (Y_i|Y_1^{i-1}, X_1^i) = \frac{\lambda_i^{Y_i}}{Y_i!} \exp(-\lambda_i) \quad (3.4)$$

where X and Y are as defined above, $\alpha_0, \dots, \alpha_J, \beta_1, \dots, \beta_K$ are estimated parameters, and J and K are the lengths of the intrinsic and extrinsic histories, respectively. For a given J and K , the parameters can be readily estimated using maximum likelihood. One way to estimate J and K is by using model selection techniques such as minimum description length (MDL)[4]. However, in practice, for neurons with low firing rates ($< 5\text{Hz}$) there may not be sufficient data to provide reliable parameter estimates. Furthermore, exhaustively searching over all combinations for J and K to minimize the MDL-penalized negative log-likelihood can quickly become infeasible. To address this issue, we impose a smoothness constraint on neighboring parameters to reduce the number of estimated parameters. Since the parameter sets $\alpha_0, \dots, \alpha_J, \beta_1, \dots, \beta_K$ represent the effect of spiking activity in the past consecutively, it is likely that the effect at a time t in the past does not differ widely from the effect at time $t - 1$. Specifically, we restrict the parameter sets to lie in a subspace spanned by a set of raised cosine basis functions (see Appendix). These basis functions have been successfully used previously to model the spike response of retinal ganglion cells [57]. Since the functions use a log-scaling of time, they allow for a fine resolution to model the effect of past history near the time of a spike and a coarser resolution farther back in time. Under this restriction, the modified conditional intensity function is written as:

$$\log \lambda_i = \hat{\alpha}_0 + \sum_{l=1}^L \hat{\alpha}_l \langle b_l^{int} \cdot Y_{i-R}^{i-1} \rangle + \sum_{m=1}^M \hat{\beta}_m \langle b_m^{ext} \cdot X_{i-R}^{i-1} \rangle \quad (3.5)$$

where b_l^{int} and b_m^{ext} are the sets of raised cosine basis functions of length R modeling the intrinsic and extrinsic history effects respectively. $\langle \cdot \rangle$ denotes dot product. With this constraint, the number of parameters that need to be estimated is simply $L + M + 1$. As in [57], we use 10 raised cosine basis functions to represent the intrinsic history, and 4 basis functions for the extrinsic history to provide sufficient flexibility

without too many parameters. These functions cover spiking histories up to 100ms in the past. Conversely, to estimate the conditional distribution without knowledge of X , we fix the set of $\hat{\beta}_m$ to be 0.

3.2.2 Significance and confidence intervals

Since the log-likelihood with knowledge of X contains all the parameters of the log-likelihood without X , the directed information as estimated using the method described in Section 2.2 always results in a non-negative value. Hence, to determine whether including X *significantly* improves the maximized log-likelihood, we adopt two intuitive approaches: 5-fold cross-validation and confidence intervals. Using cross-validation, we test how well the model generalizes to new data. For new data, the log-likelihood containing X is no longer necessarily larger than the log-likelihood without X since the model could overfit. Therefore, if the log-likelihood of the new data decreases using the model including X , the directed information is set to 0.

Secondly, we compute 95% confidence intervals on the directed information by estimating the largest deviation the maximized log-likelihood takes on when evaluated at the confidence interval limits on the estimated parameters. Since the parameters are maximum likelihood estimates, 95% confidence intervals on each of the estimated parameters can be approximated using the inverse of the Fisher information matrix, which is known to represent the covariance matrix of the asymptotic distribution for maximum likelihood estimates [58]. Hence, for each parameter $\hat{\beta}_j$, the 95% confidence interval is approximated as:

$$\hat{\beta}_j \pm \frac{1.96}{\sqrt{\mathbf{I}(\hat{\beta})_{jj}}} \quad (3.6)$$

$$\mathbf{I}(\hat{\beta})_{jj} = -\left. \frac{\partial^2 L}{\partial \hat{\beta}_j^2} \right|_{\alpha_0, \dots, \alpha_L, \hat{\beta}_1, \dots, \hat{\beta}_M} \quad (3.7)$$

where L is the log-likelihood function. Since the key question is whether knowledge of X significantly improves the log-likelihood, we focus on the confidence intervals as a result of uncertainty in the extrinsic history parameters $\hat{\beta}_m$. Using this approach, we obtain a confidence interval on the directed information. If this interval includes 0, we again set the directed information to 0.

3.3 Simulation setup and results

We generated simulated spike trains to systematically compare the effectiveness of directed information against standard correlation procedures under various conditions and topologies. We simulated spike train data using the sparse generalized Laguerre-Volterra model (SGLVM), which is a nonlinear model used to characterize spike trains in the hippocampus [59] (see Appendix for more details). We specifically chose a different model to test how well our estimation procedure for directed information (which has an a priori model) can perform for a mismatched model, especially with non-linearities that are not specified in the a priori model. The preset parameters were determined from fitting the SGLVM on actual experimental data in order to generate physiologically realistic simulations. The average firing rates were chosen to reflect realistic firing rates seen in experimental data (~5-10Hz). Six different sets of linear/non-linear coupling parameters were randomly chosen from experimental data, and for each parameter set, we varied the strength from 0.5 to 3 times the original value. For each configuration, 100 sets of 200,000ms spike trains were generated. In terms of coupling, we tested both two-neuron and three-neuron coupling topologies (see figure 3.1A). For the three-neuron topology, we randomly chose a different set of coupling parameters for each of the two connections but varied the strength together. Under the proxy topology, spike train B is first generated independently, which is then used as the extrinsic history (the variable X in (A.2)) to generate spike train C, and in turn used in simulating spike train A. For the cascading topology, spike train C is generated independently, which is then used as the extrinsic history for simulating both spike trains A and B. We also generated independent (“null” case) spike trains to test for false positives.

3.3.1 Metrics for comparison

We computed directed information from the simulated spike trains using two methods: our proposed method, “CrossVal-DI”, described above and the method originally used by Quinn et al. [4] (hereafter referred to as “MDL-DI”). We compared the directed information to commonly used correlation methods for assessing function connectivity: standard correlation of firing rates (with no time shift) and cross-correlation. Other directed measures like Granger causality and its variants (e.g. directed transfer method (DTF)) have mainly been adapted for analyzing fMRI and EEG data. Since the statistical properties of spike trains are quite different, such methods do not readily apply. Although Kaminski et al. [49] showed that DTF can be extended for spike trains after suitable low-pass filtering of the signal, this method have yet to be applied extensively in experimental data to our knowledge. Hence, we chose to use more general, non-parametric, correlation measures for comparison.




Figure 3.1: (A) Two-neuron and three-neuron topologies simulated. Arrows indicate preset influences, i.e. an arrow from neuron B to A means that neuron B’s spiking activity is used as input to generating neuron A’s spiking activity. (B) Simulated spike trains. (C) Distribution of directed information obtained using CrossVal-DI for the two-neuron topology for one sample set of parameters.

The firing rate used in the correlation analysis is estimated by binning (“Bin/Corr”) or smoothed with a sliding window (“Conv/Corr”). We varied the binning and window size over 5, 10, 50 and 100ms. To determine if a correlation was significant, we followed the standard procedure and transformed the correlation coefficient to a t -statistic in which we obtained a p -value on the level of significance. For cross-correlation (“XCorr”), we shifted the spike trains at 1ms time steps from -100ms to 100ms with respect to each other. To determine significance, we computed the t -statistic and associated p -value for the largest correlation coefficient obtained within the -200ms to 200ms time interval. In addition, the resulting time lag can be used to infer directionality. For example, if $X(n+m)$ is found to be most correlated with $Y(n)$ for a positive value m , this would suggest that Y influences X .

3.3.2 Two-neuron simulations

For each simulation configuration, we measured the overall accuracy (true positive plus true negative rate) for each metric. For directed information, we also computed the measure for the opposite direction, wherein the accuracy was simply the number of true negatives observed. The results for one of the parameter sets at varying connection strengths are shown in figure 3.2A. Table 3.1 shows the accuracy averaged across all 6 parameter sets for all metrics. For CrossVal-DI and MDL-DI, only the accuracy from B to A was included. The binning/window size shown for Bin/Corr and Conv/Corr corresponds to the bin size with the highest accuracy across the




Figure 3.2: (A) Accuracy of each metric at different connection strengths for one set of parameters. Accuracy for (B) proxy topology and (C) for cascading topology. Connection analyzed is marked in orange.

different connection strengths. For Bin/Corr, we set significance at $p < 0.05$, which resulted in $\sim 5\%$ false positive rate under the null connectivity case. For Conv/Corr and XCorr, a 5% confidence interval on the p -value resulted in a high false positive rate, so a stricter 0.1% confidence interval was used. Despite the stricter p -value threshold, XCorr still showed a 64% false positive rate. Out of all the positive detections using XCorr, on average 77% provided the correct directionality. At all connection strengths, CrossVal-DI consistently showed the highest accuracy.

3.3.3 Three-neuron simulations

Proxy topology

We tested each metric's ability to detect indirect connections, namely in the case where the connection between neurons B and A is through a third neuron C. The accuracy for each metric at each connection strength is shown in figure 3.2B. As the connection strength (from B to C and from C to A) is increased, all metrics showed increased number of significant detections between B and A, with CrossVal-DI showing highest accuracy when also considering false positive rates for all the metrics. As expected, the directed information in the opposite direction (from A to B) measured using both directed information methods aptly remained around the null case level of significance at all connection strengths.

Table 3.1: Accuracy in two-neuron topology averaged across all 6 parameter sets

<i>Strength</i>	<i>Bin/Corr</i> (10ms bin)	<i>Conv/Corr</i> (5ms bin)	<i>XCorr</i>	<i>CrossVal-DI</i>	<i>MDL-DI</i>
0.5	0.51±0.03	0.56±0.07	0.49±0.04	0.63±0.1	0.5±0.01
1	0.65±0.1	0.69±0.1	0.55±0.09	0.79±0.1	0.6±0.2
2	0.84±0.2	0.88±0.2	0.62±0.2	0.92±0.1	0.83±0.2
3	0.89±0.09	0.91±0.07	0.64±0.07	0.96±0.02	0.95±0.2

Cascading topology

We also considered the topology where two neuron A and B are influenced by a common parent neuron C. Again, we used different coupling parameters for the two connections but varied the strengths together. The accuracy for each metric at each connection strength is shown in figure 3.2C. Unlike the proxy topology, a significant detection between A and B was considered a false positive. For all metrics, the number of significant connections detected between A and B increased with higher connection strength (leading to decreased accuracy), including the number of significant connections in the opposite direction (from B to A) measured by directed information. The figure also suggested exceptional accuracy for the MDL-DI method. However, this was to some extent an artifact due to the fact that this topology simulation only tested for true negatives and no true positives. From the 2-neuron and proxy topology simulations, we found that MDL-DI in general provided a low true positive rate, but high true negative rate, which led to high accuracies when true positives were not tested for. In contrast, XCorr showed particularly poor accuracy since this method tended to have low true negative rate.

Using directed information, this dependence between A and B can be accounted for by conditioning the distributions in (3.1) on the activity of C. Hence conditioned on C, the number of significant connections from A to B and from B to A, measured using CrossVal-DI decreased to 17% and 18% respectively at the highest connection strength, which would improve the accuracy of CrossVal-DI by 20% at that connection strength.

Under the cascading topology, the connectivity in both directions between B to A (measured using DI methods) appeared to be equally strong, compared to the proxy topology where the directionality was clearly one-sided. Comparing the strength of the coupling in both directions could also provide information about the topology between the two neurons.




Figure 3.3: (A) Experimental setup. (B) Distribution of first DLS spike time after ICMS across all units. (*right*) Distribution of first DLS spike for two individual units. The peak at ~ 6 ms reflects the conduction delay between M1 and DLS. (C) Effect on unit firing rate due to ICMS, estimated using CrossVal-DI, for an individual unit. (D) KS statistic measuring deviation of first-spike latency distribution from uniform for each unit in descending order. (E) Directed information measuring the effect of ICMS on each unit in the same order as in (D). (F) Correlation between KS statistic and directed information.

3.4 Rodent experimental data

3.4.1 Experimental setup

Next, we tested the ability of CrossVal-DI to detect connectivity in experimental data. We applied CrossVal-DI to measure the influence of intracortical microstimulation (ICMS) applied to the primary motor cortex (M1) on spiking activity in the dorsolateral striatum (DLS) of the rat (figure 3.3A), which was known to receive projections from the sensorimotor cortex [60]. A Long-Evans rat weighing approximately 250 grams was chronically implanted with tungsten microwire arrays in both M1 and DLS. Each array contained 16 microelectrodes ($35\mu\text{m}$ diameter, $200\mu\text{m}$ electrode spacing, $300\mu\text{m}$ row spacing in 8 by 2 configuration; Innovative Neurophysiology, Durham, NC). Stereotactic coordinates relative to bregma were used to center the arrays (anteroposterior 2 mm, mediolateral 2 mm, and dorsoventral 1.5 mm for M1; and anteroposterior 0.5 mm, mediolateral 4 mm, dorsoventral 4.5 mm for DS). M1 implants were targeted for layer 5 pyramidal neurons. Monophasic ICMS pulses

at $300\mu\text{A}$ and $250\mu\text{s}$ duration were applied at random intervals between 0.5s to 3s in M1. Single unit activity from the DLS was recorded with a Multichannel Acquisition Processor (MAP; Plexon Inc., Dallas, TX) and sorted using an online sorting application (Plexon Inc., Dallas, TX). All procedures were performed in compliance with the regulations of the Animal Care and Use Committee at the University of California, Berkeley.

Based on the distribution of the first DLS spike latency following ICMS (“first-spike latency” distribution, aggregated over all 15 units recorded, figure 3.3B), there was a sharp excitatory increase in spike probability at approximately 6ms after a brief inhibitory period. This 6ms latency was consistent with conduction delay estimates (~ 1.8 to 10ms) between M1 and DLS reported in [61]. Individually, the effect of ICMS varied from unit to unit. If the effect of ICMS on a unit was weak, we expected the first-spike latency distribution for this unit to resemble a uniform distribution. Hence, we used the deviation from uniform as a rough estimate of the influence of ICMS on each unit. To quantify this deviation, we computed the Kolmogorov-Smirnov (KS) statistic which measured the maximum deviation between the empirical cumulative distribution function (CDF) and the uniform CDF. The KS statistic for each unit’s first-spike latency distribution is shown in figure 3.3D.

3.4.2 Data analysis

We used CrossVal-DI to estimate the directed information from ICMS to each unit recorded in the DLS. The data set consisted of 321s of spike train data, binned at 1ms resolution, for 15 units. All parameters (e.g. basis function set) were kept the same as those used throughout all the simulations. Out of 15 units, 12 showed significant influence from ICMS (figure 3.3E), while only 1 unit out 15 showed significant connection in the opposite direction (i.e. the effect of DLS spiking on ICMS). Aside from significance of connection, directed information also provided information on the temporal profile of the effect of ICMS based on the estimated set of parameters $\hat{\beta}_m$. From (3.3) and (3.5), the exponential of the extrinsic history factor $\exp(\beta_k) = \exp\left(\sum_{m=1}^M \hat{\beta}_m b_m^{ext}\right)$ represented a multiplicative gain on the unit’s firing rate λ due to an ICMS pulse at k time lags in the past. This factor as a function of time lags for a representative unit is shown in figure 3.3C. On average, a sharp increase in gain occurred at $5.7 \pm 0.5\text{ms}$, which corresponded well with the conduction delay estimate from figure 3.3B. In addition, we tested CrossVal-DI’s ability to quantify the degree of influence. Comparing the directed information obtained using CrossVal-DI with the KS statistic obtained earlier, we found that directed information was strongly correlated with the KS statistic ($R^2 = 0.83$), which suggested that directed information was a strong indicator of the degree of influence.

We performed the same analysis using the correlation method with the highest overall performance in simulations (Conv/Corr with 5ms sliding window and $p < 0.001$ threshold). Using this procedure, only 1 out of 15 units showed significant correlations. Secondly, the R^2 values obtained from Conv/Corr did not correlate with the KS statistic ($R^2 = 0.04$). Moreover, the 6ms conduction delay cannot be estimated from this method. The low correlation likely resulted from the delayed influence from ICMS on DLS spiking due to the conduction delay. In this case, Conv/Corr with a larger smoothing window (i.e. > 6 ms) or XCorr (which accounts for time shifts) were much more suitable approaches, even though the accuracies were lower during simulations. Using a 10ms window, 11 out of 15 units were significantly correlated, but the R^2 values did not correlate well with the KS statistic ($R^2 = 0.17$). Using XCorr, 12 out of 15 units showed significant correlations, and the R^2 values correlated well with the KS statistic ($R^2 = 0.76$). Furthermore, the conduction delay could be estimated from the time shift that gave maximal correlation. On average, we found the conduction delay to be 4.5 ± 2 ms, consistent with the previous estimate of 6ms. Hence, choosing the appropriate method and binning parameters to do correlations may require some prior knowledge on the data, such as the conduction delay in this case.

3.5 Primate experimental data

3.5.1 Experimental setup

To further assess the applicability of directed information, we applied the analysis to spike trains recorded from M1 of an adult male rhesus monkey (*Macaca mulatta*). Data was collected over 16 days from 64 teflon-coated tungsten microelectrodes (35 μ m diameter, 500 μ m electrode spacing, in 8 by 8 configuration; CD Neural Engineering, Durham, NC). The arrays were implanted bilaterally in the hand/arm area of M1, positioned at a depth of 3mm targeting layer 5 pyramidal neurons. Localization of target areas were performed using stereotactic coordinates of the rhesus brain. Single unit activity was recorded with a Multichannel Acquisition Processor and sorted using an online sorting application (Plexon Inc., Dallas, TX). Isolated units were then verified off-line. All procedures were performed in compliance with the regulations of the Animal Care and Use Committee at the University of California, Berkeley.

The monkey was trained to perform a delayed center-out reaching task (figure 3.4A) using his right arm (“manual control”) or with a cursor controlled through a brain-machine interface (“neuro-prosthetic control”). The task involved cursor movements from the center towards one of eight targets distributed evenly on a 14cm diameter

circle. Target radius was set at 0.75cm. Each trial began with a brief hold period at the center target, followed by a GO cue (center changed colour) to signal the reach towards the target. The monkey was then required to reach and hold briefly (0.2-0.5s) at the target in order to receive a liquid reward. In manual control, reaching was performed using a Kinarm (BKIN Technologies, Kingston, ON) exoskeleton where the monkey’s shoulder and elbow were constrained to move the device on a 2-D plane. In neuro-prosthetic control, spiking activity from a select number of neurons were mapped to kinematic parameters that determined cursor position. Specifically, units with stable action potential waveforms and interspike interval distributions over days (“direct” units) were chosen as input to a linear regression model, which was then used to produce cursor position (see Section 2.2.3 for more details). Stability of the unit was the only criteria used for selecting direct units; directional tuning or modulation information was not considered. Stability of unit recording was determined by visual comparison of the current waveform shape and amplitude to the previously stored template. The interspike interval (ISI) distributions were also examined to ensure that the recordings were from stable single units. Finally, we confirmed the stability of those units using a commercially available software (Wavetracker; Plexon Inc., Dallas, TX), which decomposed the waveforms into two- and three-dimensional principal components such that comparisons of the waveforms over time can be readily made. Other units that were also recorded, but not directly linked to the BMI are referred to as “indirect” units. Neuro-prosthetic control sessions were performed on the same day as manual control sessions for 7 out of the 16 days.

3.5.2 Degree of connectivity

We applied CrossVal-DI to every pair of units to analyze the degree centrality for the network of M1 neurons while performing the center-out reaching task under manual and neuro-prosthetic control. We used 7 recording sessions each for manual and neuro-prosthetic control. Spike trains were binned at 1 ms resolution. For neuro-prosthetic control, we analyzed the neural data during the interval as the cursor moved from the center to the target. For manual control, since the reach times were shorter, we also included the interval as the monkey returned to the center target to initiate the next trial. Units with an average firing rate below 1Hz were excluded due to the difficulty in estimating the parameters for sparse spike trains. An average of ~50 units from contralateral/ipsilateral M1 were included in the analysis. All parameters (e.g. basis function set) were exactly the same as those used in simulations with the addition of the reaching target angle as another covariate to the GLM. This was used to account for dependencies that may arise purely from similar directional tuning.

The directed information method proposed here provides direction in connections




Figure 3.4: (A) Experimental setup. The monkey performs a center-out reaching task using neuro-prosthetic and manual control. (*bottom*) Timeline of task. (B) Sample waveform and directional tuning of a typical unit with cosine tuning. (C) Connectivity “map” obtained from applying CrossVal-DI to all possible pairs of units. Each pixel represents the connection strength between the two units. (D) Number of connections for each unit sorted by directionality.

between units. For each unit, we defined two separate degree centrality measures: out-degree, which measured the number of outgoing connections, and in-degree which was the number of incoming connections (figure 3.4B). For some pairs, we found the directed information to be significant in both directions. This may be indicative of a cascade topology between these two neurons (i.e. they are both influenced by a common neuron) and not a true functional connection. Hence, for these connections, we applied the following convention: if the connection in one direction was much stronger (specifically, twice as strong) than the other, we considered the overall connection to be unidirectional. Overall, the unidirectional connections constituted approximately 70% of all the connections (figure 3.4C). All subsequent analysis was focused on the unidirectional connections.

To quantify the relevance to external behaviour for each unit, we computed the commonly used modulation depth, which was the firing rate difference between the maximum and minimum of the fitted cosine direction tuning curve (see Appendix). Figure 3.5 shows the modulation depth plotted against out-degree and in-degree across all days for the two modes of control, separated by neuron population. The out- and in-degree were normalized by dividing by the total number of units recorded for each day. Interestingly, we found that modulation depth was most strongly correlated with out-degree ($R^2 = 0.5$, $p < 0.01$) during neuro-prosthetic control, but only weakly correlated during manual control ($R^2 = 0.14$, $p < 0.01$). On the other hand, modulation depth was uncorrelated with in-degree in both control modes ($R^2 = 0.002$, $p > 0.05$ and $R^2 = 0.001$, $p > 0.05$ for neuro-prosthetic and manual control, respectively). We further separated our analysis between the direct and indirect neurons. Under neuro-prosthetic control, direct neurons showed higher out-degree/modulation depth correlation ($R^2 = 0.74$, $p < 0.01$) than indirect neurons ($R^2 = 0.46$, $p < 0.01$). Under manual control, both populations were weakly correlated ($R^2 = 0.04$, $p > 0.05$, and $R^2 = 0.18$, $p < 0.01$ for direct and indirect neurons, respectively).

We also used an alternate metric for quantifying the relevance to external behaviour for each neuron: the mutual information between the neuron’s activity and target (see Section 2.2.4). We again found that mutual information was more correlated with out-degree ($R^2 = 0.21$, $p < 0.01$) during neuro-prosthetic control than during manual control ($R^2 = 0.15$, $p < 0.01$). In-degree was uncorrelated with mutual information during both control modes ($R^2 = 0.0006$, $p > 0.05$ and $R^2 = 0.002$, $p > 0.05$ for neuro-prosthetic and manual control, respectively).

We repeated the analysis using correlation-based approaches to compute the connectivity. Similar to before, we used the correlation method with the highest overall performance in simulations (Conv/Corr with 5ms sliding window and $p < 0.001$ threshold). Since correlation is a symmetric measure (i.e. correlation between X and Y is the same as between Y and X), we could only evaluate the number of connections

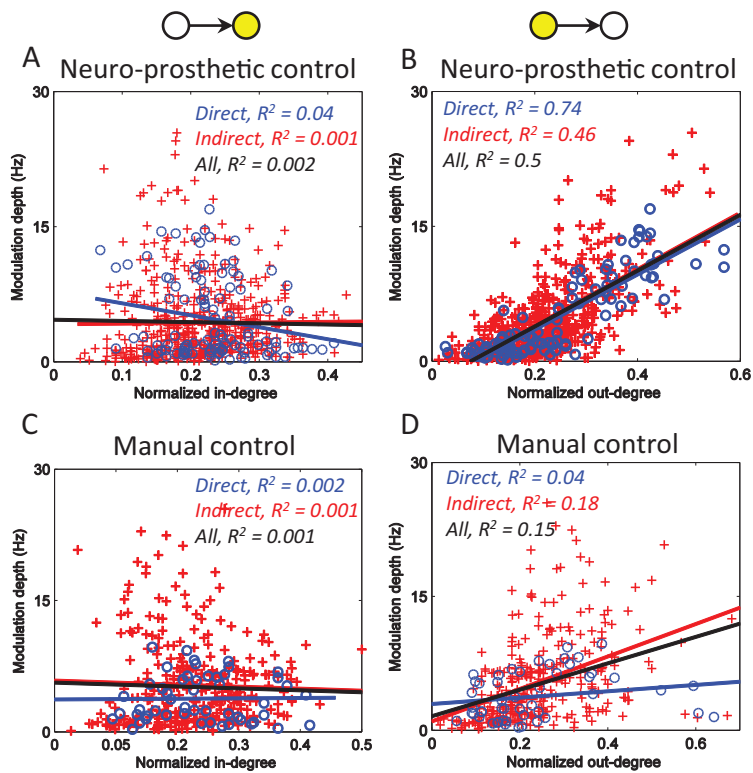


Figure 3.5: Modulation depth plotted against normalized in-degree and out-degree for neuro-prosthetic control (A, B) and manual control (C, D). Circles refer to direct neurons, crosses refer to indirect neurons.

for each unit. Overall, we found modulation depth to be significantly correlated with the (normalized) number of connections during both modes of operation ($R^2 = 0.58$ and $R^2 = 0.54$ for neuro-prosthetic and manual control, respectively). However, this analysis could not provide further insight on whether behaviourally relevant neurons correspond to those that influence many other neurons, or being influenced by many others.

3.5.3 Stability of connections

In addition, we examined whether the connection strengths changed over the 16 days of neuro-prosthetic control. If the changes in the network connectivity reflected changes due to learning, the connection strengths between the same neurons should show gradual consolidation as the task was learned. However, due to the impossibility of recording from the same neurons across the whole array over days (due to head movement, scar tissue buildup, etc.), we confined the analysis to the direct neurons, which were closely inspected to ensure that the action potential and interspike interval were stable over time. The connection strengths computed using CrossVal-DI between the direct units for days 5, 13 and 14 are shown in figure 3.6A. To evaluate the stability of the connections across days, we performed pairwise correlations of the connection strengths for each unit across days. We found an initial increase in the stability of the connections, followed by saturation (figure 3.6B). This result reveals striking similarity with the stability of tuning properties shown by Ganguly and Carmena[22].

We replicated the analysis using Conv/Corr (5ms sliding window) to measure connection strength. Although both CrossVal-DI and Conv/Corr showed increased stability in the connection strength over days, the correlation in connection strength for Conv/Corr was much weaker, and the consolidation occurred at a later stage. This may have resulted from a lower sensitivity of Conv/Corr to detect weaker connectivity at the earlier stages of learning.

3.6 Discussion

Directed information has been shown to be a theoretically well-grounded framework for inferring functional connectivity [4]. Here, we provide a practical assessment on the usefulness of directed information, compared to standard correlation methods, when applied to both simulated and experimental data. In addition, we present a more computationally feasible approach, CrossVal-DI, for estimating the directed information that shows higher performance than previously proposed estimators. We tested CrossVal-DI against a number of correlation procedures (Bin/Corr, Conv/Corr




Figure 3.6: (A) Connection strengths between the direct neurons for 3 different days as measured using CrossVal-DI. R^2 value below refer to the correlation between the connection strengths on different days, averaged across the direct neurons. (B,C) Pairwise correlations of the connection strengths across all 16 days, with connection strength measured using CrossVal-DI and Conv/Corr respectively. Both metrics show increased stability of connections over time.

and XCorr with various binning/window sizes) across 6 different parameter sets at 4 levels of increasing connection strengths. At all connection strengths, CrossVal-DI shows high accuracy in estimating the connectivity in both directions on average. Furthermore, CrossVal-DI is able to correctly identify the direction of influence, which cannot be determined with correlations. Most importantly, the method is able to robustly detect connections even when the underlying simulation model is not a perfect match, but contains different non-linearities and terms that are not accounted for in the a priori model imposed by CrossVal-DI. This robustness is also reflected in the proxy topology where the connection between the two neurons is further obscured by an intermediate neuron. Again, CrossVal-DI can still accurately detect the direction of influence.

We further demonstrated the robustness of CrossVal-DI by applying the method to both rodent and primate data. All implementation parameters were kept the same in all applications. Applying our method to rodent recordings, we show that directed information can correctly infer the degree of influence of intracortical stimulation in M1 on neural spiking activity in DLS. The conduction delay between M1 and DLS can also be approximated from the estimated parameters in the CrossVal-DI model. We replicated the analysis using correlations for comparison. We first applied the method with the highest accuracy in simulations (Conv/Corr with 5ms sliding window), but could not find any significant connections. Instead, XCorr was found to be a more suitable approach in elucidating these connections.




Figure 3.7: (A) Illustration of high out-degree and in-degree neurons in relation to network. (B) Schematic of the task-related network under manual (*top*) and neuro-prosthetic control (*bottom*). Under manual control, the task-related network is likely wider spread, and so neurons with high out-degree may not be task related. For neuro-prosthetic control, there is a high correlation between neurons with high out-degree and modulation depth, which suggests that the observed network is more relevant to the task.

Next, we analyzed neural recordings from a macaque monkey performing a center-out reaching task under manual and neuro-prosthetic modes of operation. We examined the degree centrality of each neuron, as such measures (i.e. in- and out-degree) are often useful indicators of importance in a network. We found that a neuron’s out-degree was more strongly correlated with its modulation depth than in-degree. Neurons with high out-degree were neurons that influenced the activity of many others and could be considered as important drivers of the observed network activity. Therefore, if these neurons were also highly modulated by the target direction, then this suggested that the observed network was highly relevant to the motor task. Conversely, neurons that were influenced by many others (i.e. high in-degree) may have received input from non-task related neurons that acted as noise with respect to the motor task (see figure 3.7A). Hence, such neurons would not necessarily have high modulation depth, as evident from the lack of correlation with modulation depth. Comparing between manual and neuro-prosthetic control, the out-degree/modulation depth correlation was weaker in manual than neuro-prosthetic control. Correspondingly, this weaker correlation in manual control suggested that the observed network was less related to the task. Even though there were neurons with high modulation depth, they did not influence the activity of many others, while there were highly influential neurons that were unrelated to the task. This was likely a result of the higher complexity and wider spread of the network responsible for controlling the joints and muscles associated with natural arm movement. Our electrode array may only have been recording from a portion of the relevant network mixed with neurons from unrelated networks, resulting in a weak relation between observed network activity and modulation depth (see figure 3.7B). In neuro-prosthetic control, however, movement control was solely determined by the network of direct neurons. Hence, the strong positive correlation between out-degree and modulation depth for the direct neurons in neuro-prosthetic control reflects a change in neural activity that focuses the task-relevant control around the direct neurons. A recent study [62] found widespread changes in modulation depth between neuro-prosthetic and natural arm control that occur differently in direct and indirect neurons. Here, our findings suggest that these widespread changes in modulation depth are coupled with changes in functional connectivity in a coordinated fashion that results in a more focused network of control.

Lastly, we analyzed the stability of the connections over time. Our method revealed a consolidation of the connectivity over time that closely follows the consolidation of directional tuning of the neurons shown by Ganguly and Carmena [22]. In that study, the authors showed that the directional tuning properties of the direct neurons stabilized as neuro-prosthetic control became proficient. Our analysis of the functional connectivity of these neurons indicates that not only is the directional tuning of each neuron *individually* stabilizing, but the connections *among* the neurons are also consolidating, which provides strong evidence that long-term neuro-prosthetic

control can lead to the formation of a stable cortical map [22].

In conclusion, these results provides further evidence that the neural circuitry changes during BMI operation. In addition, our method offers a robust, applicable approach based on directed information for identifying functional connectivity in intracortical recordings.

Chapter 4

Tackling the closed-loop BMI problem

In the previous chapters, we have discussed the various changes that go on in the brain as a subject learns to control a BMI in closed-loop. In this chapter, we shift our focus towards designing and implementing algorithms that address the closed-loop nature of BMI, as well as take advantage of different neural inputs (such as local field potentials) to maximize robustness and long-term stability.

4.1 Brain-machine interface control using local field potentials

Intracortical BMI studies have traditionally used single- and multi-unit activity as control signals, due to the fact that the discharge of motor cortical neurons correlates with different movement parameters, and that these cells can be volitionally modulated irrespective of physical movement using biofeedback [63]. On the other hand, certain features of local field potentials (LFPs), which are the extracellular potentials resulting from the synaptic activity in a neuronal population, have also been shown to exhibit similar tuning properties [11]. For instance, initiation of arm movement has been shown to reliably modulate power in the LFP beta band ($\sim 15\text{--}30$ Hz) [64]. In addition, other studies have shown that muscle activity, eye movements, and reaching and grasping kinematics can all be decoded from LFP activity to varying degrees of accuracy [9, 10, 65]. This makes LFPs an attractive alternative to single- and multi-unit activity in BMIs because they are more robust to signal degradation over time [66] and contain more information than other less invasive field potentials such as electroencephalography (EEG) and electrocorticography (ECoG) signals.

Despite the promising properties of LFPs for BMI control, only a few studies to date have used them explicitly as an input signal in BMIs [67–69], with one study demonstrating 2-D continuous BMI control [69]. In that study, Flint et al. trained decoders offline biomimetically (using neural activity recorded during overt arm movements) and held them fixed thereafter. While proficient BMI control was achieved in this way, biomimetic decoders may not be viable in certain clinical scenarios where the intended user cannot perform arm movements to fit the BMI decoder. Furthermore, the biomimetic approach may constrain the subject to modulate neural activity during a closed-loop BMI task in a way similar to activity evoked during natural arm movements. For instance, a low frequency feature known as the local motor potential (LMP) has been found to be informative in decoding kinematic parameters of arm-related movements [70], and Flint et al. indeed reported that the LMP and the 0–4 Hz band were the most informative features for LFP-based BMI control [71]. However, little is known about the degree of flexibility of the brain to modulate LFP oscillations of higher frequencies, or whether a broader range of frequencies can be used for BMI control with LFPs.

Here, we apply closed-loop decoder adaptation (CLDA) [72] during LFP-based BMI control to adapt the decoder to subject-specific modulations of different LFP frequency bands. Using CLDA in combination with an assistive control paradigm, we trained two non-human primates to perform a center-out task using LFP activity in the 0–150 Hz range. We then analyzed the relative importance of the different frequency bands towards each monkey’s cursor control. While both monkeys obtained proficient control of the cursor under identical task settings, the frequency bands most important to control were characteristically different between both monkeys. Next, we constrained the BMI control input to various frequency sub-ranges (0–40 Hz, 40–80 Hz, and 80–150 Hz) and found that our offline analyses accurately predicted the bands with which the monkeys achieved the best online performance. Interestingly, while each monkey performed better using certain frequency ranges, both monkeys were able to achieve BMI control well above chance level with all sub-ranges after decoder adaptation, suggesting that broad ranges of LFP frequencies can potentially be used for closed-loop BMI control. Finally, we found that the particular frequency bands most important for each subject’s control influenced the number of channels they needed, due to differences in the correlations across channels at different frequencies. Overall, our results demonstrate proficient, continuous BMI control using LFPs and shed light on the range of frequencies that can potentially be used, with implications for channel and feature selection.

4.2 Materials & methods

4.2.1 Electrophysiology and behavioral task

Two adult male rhesus macaques (*macaca mulatta*) were chronically implanted in the brain with bilateral microwire arrays of 128 teflon-coated tungsten electrodes (35 μm diameter, 500 μm wire spacing, 8×16 array configuration; Innovative Neurophysiology, Durham, NC), targeting the arm areas of primary motor cortex (M1) and dorsal premotor cortex (PMd). All procedures were conducted in compliance with the NIH Guide for Care and Use of Laboratory Animals and were approved by the University of California–Berkeley Institutional Animal Care and Use Committee.

LFP signals were sampled at 1 kHz using a 128-channel MAP recording system (Plexon Inc., Dallas, TX) and streamed to a dedicated computer running MATLAB (The Mathworks, Natick, MA) to implement feature extraction and closed-loop BMI control. Channels were referenced to ground and signal quality was visually inspected each day (channels with clear artifacts were removed). For each channel, we estimated the spectral power in 15 consecutive 10-Hz bands from 0–150 Hz using the multi-taper method [73]. Spectral estimation was performed every 100 ms using a sliding window containing the most recent 200 ms of raw LFP activity. The estimates of log spectral power in different frequency bands, across multiple LFP channels, were used as neural features for closed-loop BMI control (Figure 4.1A).

The monkeys were head-restrained in a primate chair and had their arms confined within the chair while performing a self-paced 2-D center-out task. Monkeys were previously trained to perform this task using their right arm. Trials were initiated by moving the cursor under neural control to the center target and holding for 300 ms. after which the monkeys had to reach to one of eight peripheral targets uniformly spaced about a 13 cm diameter circle and hold for 400 ms to receive a liquid reward (target radii = 1.7 cm). The monkeys were then required to move the cursor back to the center target to initiate the next trial. If the monkeys failed to hold or reach the target within 10 s, the trial was restarted without reward (Figure 4.1B).

4.2.2 Decoder design

A Kalman filter (KF) decoding algorithm (“decoder”) was used to implement closed-loop BMI control. The KF assumes the following models for how the kinematic state of the cursor (x_t) evolves over time and how the observed vector of extracted LFP features (y_t) relates to the state:

$$x_t = Ax_{t-1} + w_t \quad (4.1)$$

$$y_t = Cx_t + q_t \quad (4.2)$$




Figure 4.1: **Experimental setup.** (A) Illustration of the BMI task set-up. (B) Trial timeline for the BMI center-out task.

where $w_t \sim \mathcal{N}(0, W)$ and $q_t \sim \mathcal{N}(0, Q)$ are additive Gaussian noise terms. At each iteration, the KF estimates x_t using y_t and its previous estimate of x_{t-1} (see [74] for the actual filter equations that perform the state estimation at each iteration). While there is no explicit delay between the kinematic state and neural activity in the KF model, each measurement y_t contains information from neural activity up to 200 ms in the past.

We used a position-velocity KF, in which the KF state vector x_t is defined to include the position (p) and velocity (v) of the cursor at time t , along both the horizontal and vertical directions of the screen:

$$x_t = \begin{bmatrix} p_{\text{horizontal},t} & p_{\text{vertical},t} & v_{\text{horizontal},t} & v_{\text{vertical},t} & 1 \end{bmatrix}^T.$$

The constant term 1 and a corresponding extra column of C are added to account for neural observations y_t with non-zero means. The structures of A and W were constrained as in [5] such that integrating the velocity from one KF iteration perfectly explains position at the next iteration. C and Q were initially constructed by shuffling the rows and columns of a previous set of C and Q matrices that were trained with native arm reaches from the subject. These initial parameters were then updated during closed-loop BMI operation using the SmoothBatch CLDA algorithm (see 4.2.3). We typically adapted these parameters only during initial closed-loop control, and did not perform additional CLDA on the same decoder throughout the rest of the experiment.

In addition to producing estimates of the state vector x_t , the KF also generates an error covariance matrix P_t at each time-step that represents the filter’s confidence in its own state estimates. Specifically, the diagonal terms of P_t represent the KF’s error variances for its estimates of the state vector variables (i.e., larger values reflect lower

confidence). Note that during standard operation of the filter, P_t quickly converges to a steady-state matrix P that depends only the model parameters (A, W, C, Q) .

During the first training session of each experiment, we utilized an assistive control paradigm [75] simultaneously with CLDA. After decoder initialization, the cursor was temporarily assisted towards the target. Specifically, the cursor trajectory was determined by:

$$\overrightarrow{v_{cursor}} = \alpha \cdot \overrightarrow{v_{assist}} + (1 - \alpha) \cdot \overrightarrow{v_{user}}$$

where $\overrightarrow{v_{user}}$ is the decoded output from the Kalman filter, $\overrightarrow{v_{assist}}$ is a vector that points directly towards the current target, and $\overrightarrow{v_{cursor}}$ is a weighted average of the two that determines the final cursor output shown to the subject. The weighted average was set by the assist level $\alpha \in [0, 1]$, which was manually adjusted in real-time depending on the subject’s performance (at $\alpha = 0$, the cursor was fully controlled by the subject). The assist speed was set to 0.8 cm/s to match the natural speed of the cursor when the subject performed the center-out task under manual control. All analyses were performed on data with the cursor under full volitional control (0% assist).

4.2.3 Closed-loop decoder adaptation

Closed-loop decoder adaptation (CLDA) is an emerging paradigm for improving or maintaining the online performance of brain-machine interfaces. By adapting the decoder’s parameters during closed-loop BMI operation (i.e., while the subject is using the BMI), CLDA algorithms aim to match the decoder’s output to the subject’s particular pattern of neural activity [5–8, 76–81]. In our experiments, we used the SmoothBatch CLDA algorithm to adapt the KF decoder’s parameters during initial closed-loop control. The SmoothBatch CLDA algorithm has been previously used in spike-based BMI experiments, where it was demonstrated to rapidly improve BMI performance independent of the decoder’s initialization method [6]. Furthermore, SmoothBatch has also been shown to possess a variety of favorable algorithmic convergence properties [72].

To illustrate how the SmoothBatch CLDA algorithm operates, let $C^{(i-1)}$ and $Q^{(i-1)}$ denote the current observation model parameters of the KF decoder. SmoothBatch first collects a batch of neural activity and cursor kinematics as the subject operates the BMI cursor in closed-loop control using a decoder with these parameters. Next, the algorithm generates an estimate of the user’s intended cursor movements. For example, in our experiments, we used Gilja et al.’s method for inferring intended cursor kinematics (“innovation 1” of the ReFIT-KF algorithm), which assumes that

the subject always intends to move to directly towards the current target [5]. Other methods for estimating intended movements could also be used in conjunction with SmoothBatch, such as Shpigelman et al.’s supervised method [7] or Li et al.’s unsupervised method [79]. Using the estimate of intended cursor kinematics and the recorded neural activity, SmoothBatch then constructs batch maximum likelihood estimates, \hat{C} and \hat{Q} , of the C and Q matrices using the following equations:

$$\hat{C} = YX^T (XX^T)^{-1} \quad (4.3)$$

$$\hat{Q} = \frac{1}{N} (Y - \hat{C}X) (Y - \hat{C}X)^T \quad (4.4)$$

Finally, the algorithm updates the observation model parameters using a weighted average:

$$C^{(i)} = (1 - \rho) \hat{C}^{(i)} + \rho C^{(i-1)} \quad (4.5)$$

$$Q^{(i)} = (1 - \rho) \hat{Q}^{(i)} + \rho Q^{(i-1)} \quad (4.6)$$

where i indexes discrete batch periods and the weighting parameter $\rho \in [0, 1]$ controls the influence of \hat{C} and \hat{Q} on the new parameter settings. Note that SmoothBatch does not update the KF transition model parameters A and W . We typically reparametrize ρ in terms of a half-life h (i.e., the time it takes for a previous maximum likelihood estimate’s weight in the decoder to be reduced by a factor of 2):

$$\rho^{h/T_b} = \frac{1}{2}.$$

where T_b is the size of the batch (“batch period”) measured in units of time. In our experiments, batch periods and half-lives were typically chosen in the ranges $T_b \in [1, 6]$ mins and $h \in [0.5, 15]$ mins.

4.2.4 Performance evaluation

We quantified performance using target acquisition rate and percent of initiated trials that succeeded, which was computed using a 5 minute sliding window with 30 s steps. For comparison, we estimated the chance rate (measured to be 0.3 trials/min at a 0% assist level) by simulating 30 minutes of the task performed using a random walk trajectory (generated with the same transition model parameters as the Kalman filters used for BMI control). We also measured trajectory quality using reach time (RT), normalized path length (NPL), movement error (ME), and movement variability (MV); see [6] for details on each metric.

To perform approximate comparisons of BMI cursor control performance across studies (see Discussion), we calculated a summary statistic using the Fitts Law derived

index of difficulty, which has been proposed as a standardized assessment method for neural prostheses [82]. The index of difficulty, measured in bits, is calculated based on the distance to target (“Distance”) and target diameter (“Window”) as:

$$\text{Index of difficulty} = \log_2 \left(\frac{\text{Distance} + \text{Window}}{\text{Window}} \right)$$

From this value and the mean reach time, we calculate throughput in Fitts bits/sec as:

$$\text{Throughput} = \frac{\text{Index of difficulty}}{\text{Mean reach time}}$$

We refer the reader to [5] for further details on throughput calculation.

4.3 Results

During an initial set of experiments, we randomly chose 20 channels from the right M1/PMd implants and extracted the spectral power from fifteen 10-Hz bands, spanning 0–150 Hz, as input features to the decoder. Both animals attained full control of the 2-D cursor after an initial session using CLDA with the assistive control paradigm. The assistive control was reduced to zero after day 1 and CLDA was stopped after day 2. Both animals performed the task under full volitional control for two additional days. The target acquisition rate without assist improved from day 2 to 4 for both monkeys (Pearson’s correlation coefficient, Monkey S: $R = 0.34$, $p < 0.001$; Monkey J: $R = 0.30$, $p < 0.001$). The average target acquisition rates on the 4th day for Monkey S and Monkey J were 10.6 ± 2.4 trials/min and 5.8 ± 2.9 trials/min, respectively. The average success percentage was $78\% \pm 6\%$ for Monkey S and $72\% \pm 8\%$ for Monkey J. Figure 4.2A shows typical reach trajectories, rate of target acquisition, and chance level for both monkeys on the 4th day of the series.

Trial-averaged spectrograms for each of the 8 target directions, along with tuning curves for each frequency band, are shown in Figure 4.2B–C for the two monkeys. The modulation depth, defined as the difference between the highest and lowest points of the tuning curve, is shown for each frequency in Figure 4.2D. Monkey S showed highest modulation in the 0–20 Hz frequency range, with another peak in the 50–80 Hz range, but decreasing modulation at higher frequencies. In contrast, Monkey J showed highest modulation depth in both 0–20 Hz and 80–150 Hz range. In addition, the LFP modulations during BMI control were also considerably different

than those during natural arm control. Figure 4.3 shows the trial-averaged spectrograms, tuning curves and modulation depth plots in prior training sessions where Monkey S performed the same center-out task using natural arm reaches (no brain control). The characteristic decrease in beta band (15–40 Hz) power during movement was clearly evident during natural arm movement, but did not occur during BMI control.

Although tuning curves are one representation of the relationship between cursor movement and neural activity, they do not fully capture the exact mechanism by which the cursor position is generated from the neural activity at each iteration, since this is specified by the BMI decoder. To provide further insight into the LFP activity modulated during control, we examined the KF decoder’s observation model parameters (C and Q), which capture the relationship between cursor kinematics and observed neural activity. Since our KF state vector models cursor position and velocity, the rows of C reflect the preferred position and velocity directions of different LFP features. For instance, Figure 4.4A illustrates the preferred velocity directions for each 10-Hz frequency band (averaged across channels) as assigned in C . In combination with the covariance structure modeled in Q , these preferred velocities ultimately determine how feature modulations of different features contribute to cursor movement at each iteration. For example, Figure 4.4B (top) shows the contributions to the cursor displacement from each LFP frequency band (averaged across channels) during an example trial. Each individual vector represents the net cursor displacement during the reach due to the modulations of a particular frequency band (together, the vectors add to produce the overall cursor displacement). The underlying modulations of neural activity (averaged over the course of the reach) that were generated to create this trajectory are shown in Figure 4.4B (bottom). Hence, to generate the trajectory shown, Monkey S increased the 0–20 Hz activity while reducing activity from 30–80 Hz.

In order to identify the LFP frequencies within the 0–150 Hz range that were most important for each monkey’s overall BMI control, we calculated a movement contribution profile for each monkey by similarly decomposing the cursor movement across all successful trials. We considered three broad groups of frequencies — 0–40 (approximately the delta, theta, mu, and beta bands), 40–80 (low gamma), and 80–150 Hz (high gamma) — and determined the contribution of each frequency group, measured as a percentage of the total displacement. The resulting percentages directly measure which frequencies were responsible for moving the cursor on the screen (Figures 4.4C–D, pie charts). For Monkey S, the 0–40 Hz and 40–80 Hz bands together accounted for an average of 84% of the cursor movement per reach, while the higher frequencies (80–150 Hz) only accounted for 16% of the movement. Interestingly, Monkey J’s movement contribution profile was markedly different than that of Monkey S. For Monkey J, the 80–150 Hz and 0–40 Hz bands were most responsible for

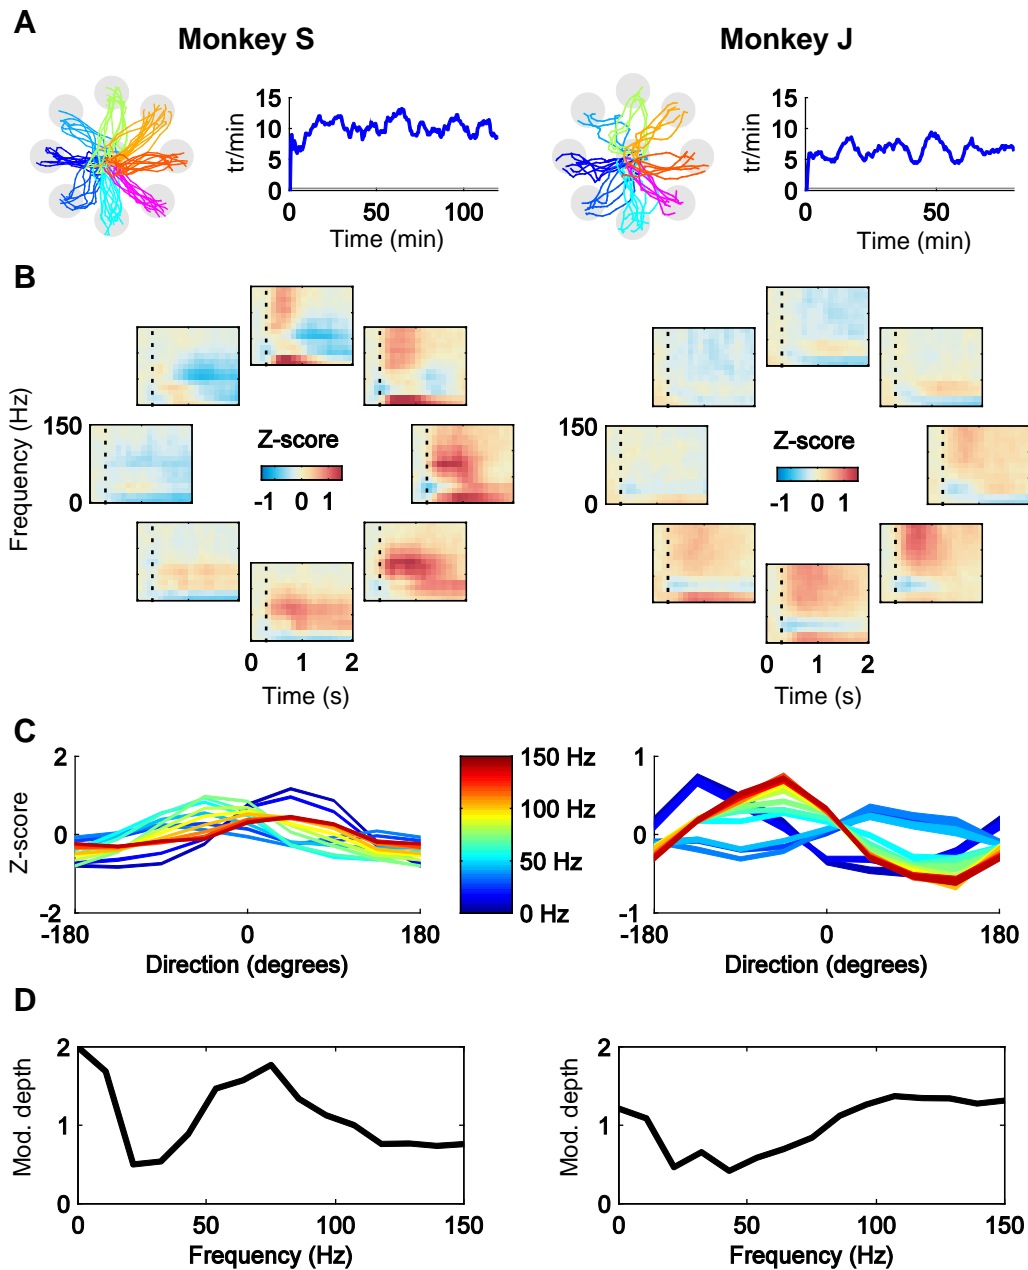


Figure 4.2: Neural activity during closed-loop BMI control. (A) *Left*: Typical reach trajectories to the eight center-out targets on the fourth day of the series for both monkeys. *Right*: Target acquisition rate on 4th day. Gray line indicates chance rate. (B) Trial-averaged spectrograms illustrating the z-scored log power on one channel during successful reaches for both subjects. (C) Tuning curves for each of the 15 spectral power features (on one particular channel) used for closed-loop control. Line thickness reflects standard error of the mean. (D) Modulation depth as a function of LFP frequency.

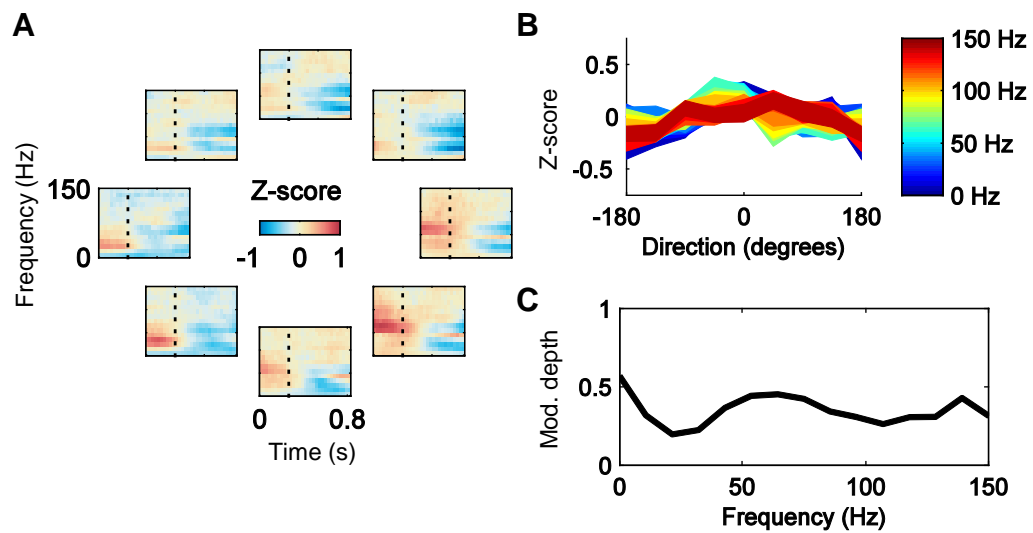


Figure 4.3: **Neural activity during natural arm movement.** (A) Trial-averaged spectrograms illustrating the z-scored log power on one channel in the contralateral hemisphere during natural right arm reaches in a center-out task for Monkey S. Dotted line indicates the go cue. (B) Tuning curves for each of the 15 spectral power features (on one particular channel). Line thickness reflects standard error of the mean. (C) Modulation depth as a function of LFP frequency.

cursor movement (46% and 34%, respectively).

As another way to measure the importance of different LFP frequencies for closed-loop control, we performed a leave-one-out offline sensitivity analysis by removing frequency groups from the Kalman filter decoder model and measuring the resulting increases in the steady-state Kalman error variances. These increases measure the impact on the decoder’s confidence in its state estimates if the corresponding features were not used in closed-loop decoding (Figures 4.4C–D, bar graphs). For Monkey S, removing the 0–40 Hz and 40–80 Hz bands resulted in larger increases (34% and 17%) in the Kalman error variances than removing the 80–150 Hz frequencies (5%), suggesting that the 0–80 Hz frequencies were relatively more important. On the contrary, for Monkey J, removing the 80–150 Hz frequencies resulted in the largest increase (25%) in the error variance compared to removing the 0–40 Hz or 40–80 Hz bands (12% and 6% increase, respectively), suggesting that 80–150 Hz was more important than the other two frequency groups for Monkey J.

4.3.1 Testing different frequency bands

To test the predictions derived from our analyses, we conducted multiple experimental series in which both monkeys were restricted to only use features extracted from various sub-ranges (0–40 Hz, 40–80 Hz, and 80–150 Hz) of the full 0–150 Hz (from the same 20 channels) for closed-loop BMI control. Each series was composed of 3–4 days, with CLDA and assistive control performed on the first day starting from shuffled decoder parameters as before. On subsequent days, CLDA was used on rare occasions only if performance decreased drastically. Figure 4.5 depicts both monkeys’ performance, measured using average success percentage, mean reach time and normalized path length, across these multiple frequency series.

Among the three frequency groups, Monkey S performed best in all metrics (success percentage, RT, NPL, MV, and ME) when using 40–80 Hz for closed-loop control. This result is consistent with his movement contribution profile (Figure 4.4C), which showed that the 40–80 Hz band accounted for an average of 41% of cursor movement during trials. However, the contribution profile showed that the 0–40 Hz band also accounted for a large percentage of movement (43%), and furthermore, our Kalman sensitivity analysis predicted 0–40 Hz to be the most important band for Monkey S. These analyses suggested that the combination of the two frequency bands (0–40 and 40–80 Hz) might allow Monkey S to achieve a higher level of performance. As such, we conducted an additional frequency series in which Monkey S used 0–80 Hz for control. As expected, with respect to all metrics, Monkey S performed best across all frequency series when using 0–80 Hz. For instance, Monkey S’s reach times were significantly lower when using 0–80 Hz than any other band ($p < 0.001$ for comparisons to each of the other four frequency series; Kruskal-Wallis ANOVA). Differences




Figure 4.4: **Importance of different LFP features to closed-loop BMI control.** (A) Preferred velocity directions for each 10-Hz frequency band as assigned in the KF observation model matrix C for Monkey S. (B) *Top*: Cursor trajectory (dashed gray line) during an example reach to the top target. Each colored vector indicates the net cursor displacement during the reach resulting from a particular frequency band. The vectors from all frequency bands add to produce the overall cursor displacement. *Bottom*: The underlying modulations of neural activity generated during the sample reach. (C–D) Movement contribution profiles (pie charts, *left*) and sensitivity analyses (bar graphs, *right*) across three groups of frequencies — 0–40, 40–80, and 80–150 Hz, for both monkeys. Movement contribution profiles were calculated by decomposing the subjects’ movements on successful trials into contributions from the three groups of frequencies, as in (B, *top*). Sensitivity analyses reflect increases in the Kalman error variances when certain groups of frequencies are removed from the KF decoder model.

in mean normalized path length, movement error, and movement variability when using 0–80 Hz versus any other band were also all significant ($p < 0.001$ for ME, MV, and NPL; Kruskal-Wallis ANOVA).

Conversely, among all bands, Monkey J performed best using 80–150 Hz in all metrics ($p < 0.001$ for differences in RT, NPL, ME, and MV; Kruskal-Wallis ANOVA). Indeed, these findings confirm the results of Monkey J’s movement contribution profile and Kalman sensitivity analysis, both of which predicted 80–150 Hz to be most important for his cursor control (Figure 4.4D). Moreover, just as those analyses had predicted 0–40 Hz to be second most important, Monkey J indeed performed second-best across all frequency series when using 0–40 Hz. Finally, Monkey J performed worst across all metrics when using 40–80 Hz, a result that is also consistent with predictions from the analysis.

Although we found that both monkeys performed better with some frequency sub-ranges than others, it is interesting to note that they were still able to achieve BMI control with all sub-ranges after initial decoder adaptation. Overall, these results suggest that there may be broad flexibility in the frequencies that could potentially be used for closed-loop LFP-based BMI control.

4.3.2 Testing different numbers of LFP channels

We also investigated the impact of reducing the number of channels on BMI control. Due to the close proximity of adjacent channels on the electrode array, the correlation between the raw LFPs from channels on the same array often exceeds 0.95. Hence, it is possible that multiple channels on the same array may not contribute additional information or degrees of freedom.

We varied the number of channels used for control in Monkey S from 50 channels to 1 channel (all from the right hemisphere, split between M1 and PMd) and tested each configuration for 3–4 days. We found that reach times decreased over time irrespective of the number of channels used (Figure 4.6A). We repeated the 50 channel configuration after the last series (1 channel series) and the reach times rapidly approached that of the previous series and kept decreasing. As mentioned above, one reason for the apparent low impact of the number of channels may be due to the high correlation between all LFP channels on a single electrode array. Comparing the power of each frequency band across 20 channels, we found high correlations between channels at frequencies below 80 Hz, but lower correlations at higher frequencies (Figure 4.6C). Since Monkey S relied primarily on modulating frequencies below 80 Hz for control, the redundancy across channels in the input was especially prominent. For Monkey S, our results suggested that the volitional modulation of different frequency bands, from even a single channel, was sufficient for 2-D cursor control.

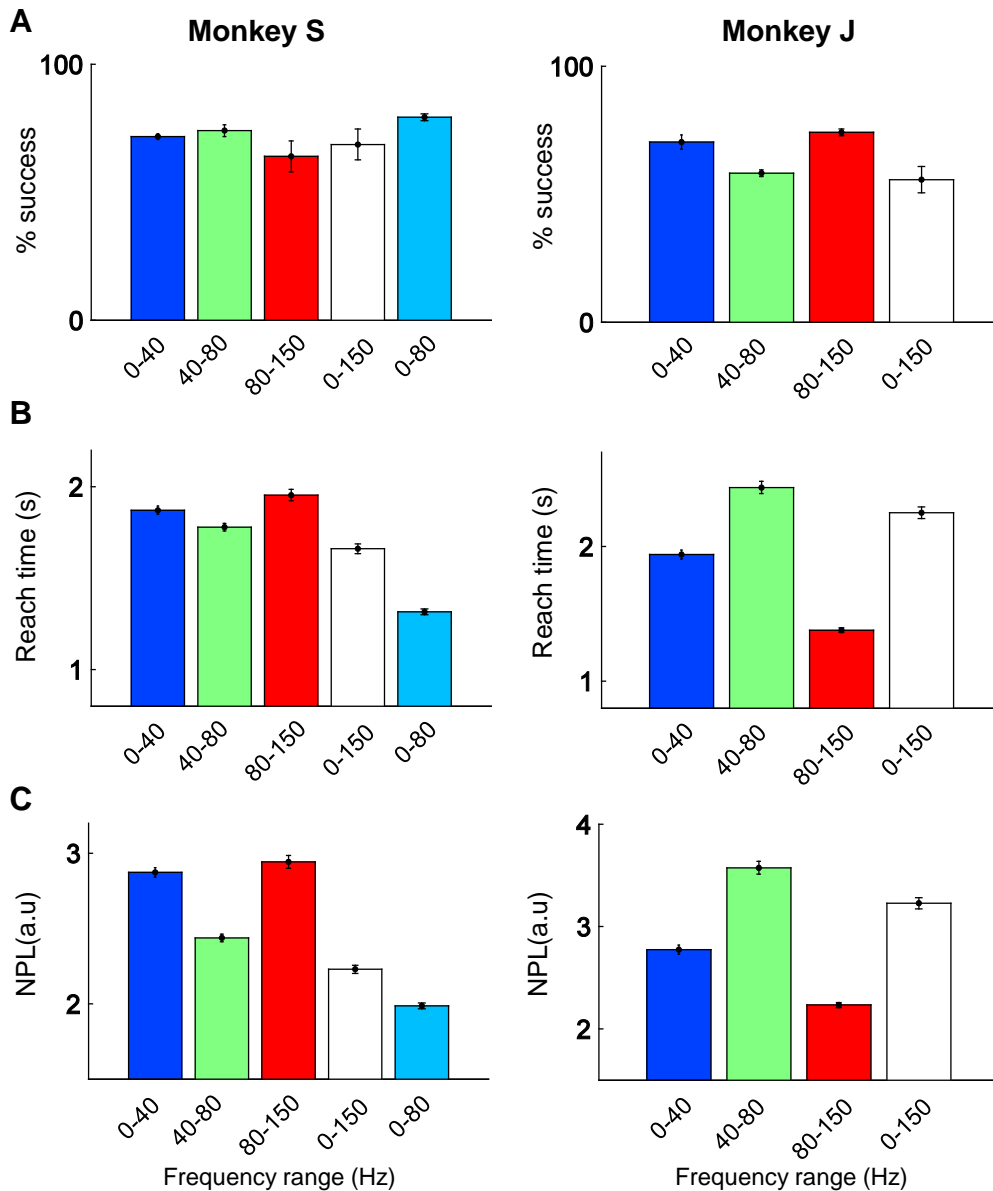


Figure 4.5: **Closed-loop BMI control with different sub-ranges of LFP frequencies.** Experimental performance, measured using (A) success percentage, (B) reach time (RT, seconds) and (C) normalized path length (NPL, arbitrary units), of both subjects across multiple experimental series in which various sub-ranges of the full 0–150 Hz range were used as neural features for closed-loop control. Bars indicate the mean across the entire series, and error bars indicate standard error of the mean.




Figure 4.6: **Closed-loop BMI control with different numbers of LFP channels.** Experimental performance of (A) Monkey S and (B) Monkey J, measured using reach time, across multiple experimental series in which varying numbers of LFP channels were used in the decoder for closed-loop control. (C) Correlation matrices for Monkey S depicting the correlations across channels for the 0–10 Hz and 110–120 Hz frequency features (i.e., log power in different frequency bands). (D) Average correlation between features values on different channels for both subjects.

We also tested BMI control in Monkey J with a 20 channel and 1 channel configuration (Figure 4.6B). In contrast to Monkey S, Monkey J’s initial reach times with 1 channel input were dramatically higher than the 20 channel input series that was used earlier. While the reach times decreased with training, the high initial increase in reach times when switching from 20 channels to a single channel for Monkey J suggested that the removal of channels from the same hemisphere had a greater detrimental effect for Monkey J than for Monkey S. The disparity may result from the difference in frequency bands most important for cursor control between Monkey S and J, and the correlations across channels for those frequency bands. For Monkey J, the 80–150 Hz band accounted for most of the cursor movement, but was also less correlated among channels (Figure 4.6D). Hence, there was less redundancy across channels compared to Monkey S, resulting in a greater effect on performance when channels were removed. Overall, we found that the effect of the number of LFP channels used on BMI control can vary depending on the strategy of frequency modulations used by a particular subject.

4.4 Discussion

By using an adaptive approach (CLDA) to fit decoder parameters, we match the decoder output to the modulations of LFP evoked by the monkeys as they attempted to move the cursor towards the target. This paradigm allows the cursor to be controlled by LFP modulations specific to the subject. As a result, while both monkeys were trained on identical task settings, our movement contribution and sensitivity analyses revealed differences between monkeys in the contribution and relative importance of each LFP frequency band. We tested the analysis results empirically by requiring the monkeys to perform the same task using only sub-ranges of the entire 0–150 Hz band. For Monkey S, the 0–40 Hz (delta, theta, mu, and beta bands) and 40–80 Hz (low gamma), contributed significantly to the cursor movement (43% and 41%, respectively) based on the movement contribution profile. Subsequently, Monkey S performed best across all performance metrics when using frequencies 0–80 Hz and worst when these frequencies were removed. In contrast, for Monkey J, the 80–150 Hz (high gamma) range contributed the most out of the three groups to cursor movement in our analysis (46%), and indeed Monkey J performed significantly better across all metrics when using 80–150 Hz for closed-loop control. However, it is interesting to note that both monkeys were able to achieve above chance level performance using *all* of the frequency sub-ranges, even when they didn’t include the most important frequency bands for each monkey.

A prior LFP-based BMI study by Flint et al. also extracted LFP spectral power as neural features, but sub-divided frequencies in a different way: 0–4, 7–20, 70–115, 130–200, and 200–300 Hz [69]. In addition, Flint et al. utilized the local motor potential (LMP) — a very low-pass filtered version of the time-domain LFP signal — as a feature for closed-loop control. In a related initial study [71], Flint et al. reported that the LMP and the power in the 0–4 Hz band appeared to be among the most informative features for control for both subjects (together, they accounted for 80% and 43% of the features used in their two monkeys). A key difference compared to the current study was that Flint et al. used a fixed decoder trained from neural activity collected as the monkey performed natural reaches. Since the decoder parameters did not adapt, the subjects may have been constrained to modulate neural activity in ways similar to the activity used to fit the decoder; namely, those generated during natural reaches. Given that the LMP has been shown to be informative in decoding movement trajectories and reach direction [70], it is not surprising that the LMP was found to be the one of the most informative features for control when using a fixed decoder trained from natural reaches. In the current study, the decoder was initialized using a random shuffling method and fit with an adaptive procedure, using neural activity evoked during closed-loop BMI control. Hence, subjects were not constrained to modulate neural activity in a predefined way.

The range of informative frequencies in our results indicate that a broader band of LFP frequencies can potentially be effective for closed-loop BMI control. This result is consistent with recent findings that LFP power in the 30–50 Hz range and ECoG power in the 75–105 Hz range can be volitionally modulated during closed-loop BMI in a operant conditioning paradigm [68, 83, 84].

Are there frequencies that are inherently more readily modulated than others in a closed-loop BMI setting? Hwang et al. [67] reported that LFP power in the 0–10 Hz and 20–40 Hz bands in the parietal reach region could be volitionally modulated to initiate trials. Engelhard et al. [68] found that the 30–50 Hz band could be robustly modulated when a constraint to use this range is directly imposed. When such constraints are relaxed, our results indicate that the range of frequencies that subjects modulate can vary from subject to subject. One potential reason for the variation may be due to the recording quality of high frequencies from the electrode array. From the time of initial array implantation to the time of the present study, the LFP power in frequencies above 80 Hz decreased by approximately 10 dB for Monkey S (~3 years post implant) but only 2 dB for Monkey J (~1.5 years post implant). Hence, the low contribution of the 80–150 Hz band to BMI control for Monkey S may be a result of the low signal-to-noise ratio in that frequency band. Nonetheless, Monkey S was able to obtain control well above the chance level using both these and other frequency bands (e.g., 0–80 Hz), indicating that broad ranges of LFP frequencies can potentially be used for a closed-loop BMI.

Another difference between the two subjects is the randomized decoder initialization. While this difference may have some effect on the subject’s control strategy, we believe the effect is almost negligible. The trained decoder that is achieved using the SmoothBatch CLDA algorithm ultimately depends on the patterns of spectral power modulations evoked by the subject, and the co-occurrence/correlation of those patterns with different directions of cursor movement. Therefore, if the subject uses roughly the same control strategy each time a new decoder is trained using CLDA, then the final trained decoder parameters will end up very similar, regardless of the initialization. Secondly, due to the high level of assistive control during initial CLDA, the cursor moves accurately regardless of the decoder’s initialization, thereby significantly reducing any effect of the decoder’s initialization on the subject’s visual feedback and resulting control strategy.

An important implication that follows from the monkeys’ choice of modulated frequencies is the effect on BMI control when the number of channels used are reduced, which relates to broader questions about the range of LFP signals. Past studies have suggested that LFPs can reflect neuronal processes 0.5–2 mm away [85]. Modulations of low frequencies (8–30 Hz) have been observed across wide areas of the motor cortex [86], while ECoG studies have reported spatially specific activity at higher frequencies (>40 Hz) [70]. The correlation between channels for the different frequency

bands in the current study are consistent with these previous studies —namely, low frequencies were more correlated than high frequencies across the array. Moreover, the monkeys’ task performance when using varying numbers of channels for control also confirms these results, as Monkey J performed worse than Monkey S when using fewer channels due to his greater dependence on higher LFP frequencies.

4.4.1 Comparison of BMI performance

Development of brain-machine interfaces to date have largely focused on spike activity as the only source of control input. The present study adds to a growing body of evidence that extracellular field potentials, measured through either intracortical arrays (i.e. LFP) or electrodes on the cortical surface (i.e. ECoG), are viable alternatives for BMIs. To compare the level of BMI performance reported here with that of other studies in the literature, we calculated the throughput in bits per second for various studies (see Methods). Throughput is a summary statistic that uses the mean reach time and Fitts Law derived index of difficulty [87, 88], which has been proposed as a standardized assessment method for neural prostheses [82]. Table 4.1 compares the BMI performance achieved in this study with a sampling of other notable closed-loop BMI cursor control studies using spikes, LFPs, ECoG, and EEG. For consistency, all studies listed in the table used a 4- or 8-target center-out task, except for Flint et al., 2013 [69] (we chose to include this study as well because it represents essentially the first demonstration of closed-loop continuous 2-D control with LFPs). To calculate throughput for Monkey S and J in the present study, we used the mean reach times from their 0–80 Hz and 80–150 Hz experimental series, respectively. Due to a wide range of varying task parameters in the literature (e.g., center-out vs. point-to-point reaching), it should be noted that the throughput metric is imperfect and does not allow for an exact comparison of cursor control performance across BMI studies. For instance, the calculation of the index of difficulty does not incorporate target hold time, an important aspect which undoubtedly affects the difficulty of the task (we have therefore listed the target hold time for each study as an additional indicator of task difficulty). However, despite its limitations, the throughput metric captures important aspects of target acquisition tasks such as target size, target distance, and reach time, and it therefore allows us to perform approximate comparisons between studies. Overall, with respect to this metric, we found that our LFP-based BMI performance compares favorably with that of spike-based studies, and exceeds the performance of other studies using LFPs, ECoG, and EEG.

Table 4.1: Performance comparison to other BMI cursor reaching studies. Comparison of the BMI performance achieved in this study with a small sampling of recent closed-loop BMI cursor control studies using spikes, LFPs, ECoG, and EEG. Throughput in bits per second is computed by calculating an index of difficulty (measured in bits) based on task settings and dividing by the mean reach time (measured in seconds). The index of difficulty of a task is calculated based upon the distance to targets and the target size (see Methods).

Neural signal	Study	Target hold time (ms)	Index of difficulty (bits)	Reach time (s)	Throughput (bits/s)
Spikes	Taylor et al., 2002 [8]	N/A	0.80	1.50	0.53
	Kim et al., 2008 [89]	500	2.89	5.51	0.52
	Ganguly & Carmena, 2009 [22]	100	2.37	2.30	1.03
	Orsborn et al., 2012 [6]	400	1.36	1.23	1.10
	Gilja et al., 2012 [5]	500	1.07	0.59	1.81
LFP	Flint et al., 2013 [69]	100	N/A	N/A	0.73
	Present study (Monkey S)	400	1.27	1.31	0.97
	Present study (Monkey J)	400	1.27	1.35	0.94
ECoG	Schalk et al., 2008 [90]	0	1.71	2.13	0.80
EEG	Wolpaw et al., 2004 [91]	0	1.18	1.90	0.62

4.4.2 Practical considerations for neuroprosthetics

LFPs have potential advantages over single-unit activity as control signals for neuroprosthetics because of their robustness and longevity. Our paradigm for achieving LFP-based BMI control is an important step forward towards clinically viable neuroprosthetics. We show that BMI control with LFPs can be achieved from initial arbitrary decoder parameters by using CLDA. In a similar study involving closed-loop LFP-based BMI control, Flint and colleagues recorded neural activity while the subject performed the cursor control task using overt arm movements to train a biomimetic decoder. However, this procedure may not be possible for paralyzed individuals who cannot move their limbs.

Our BMI paradigm allows the decoder’s parameters to be uniquely optimized for the specific set of neural features preferred by the subject. Using CLDA and an assistive control paradigm, both monkeys in this study achieved proficient performance with different ranges of LFP frequencies. Such flexibility in the neural features used for closed-loop BMI control could be important for patients with deficits in their ability to volitionally modulate certain LFP frequencies.

4.5 Multi-tasking using LFP BMI

As the subject operated the BMI system using LFP, we observed stereotypical arm and leg movements as the subject attempted to move the cursor towards the target. Hence, we were interested in determining whether such overt movements were necessary for the modulation of LFPs, or simply a “good enough” solution that the subject happened to settle on. To test this hypothesis, we used a multi-tasking experimental paradigm that required the subject to simultaneously perform the center-out task with the LFP BMI system while maintaining a steady isometric force with their hand on a force sensor. Under this paradigm, the force constraint will require the subject to modulate their LFP without the previously observed overt hand movements. Hence, if multi-tasking can be achieved, this would suggest that overt movements are not necessary for modulation of LFPs.

4.5.1 Task setup and training

The experimental setup was largely the same as before, with the additional requirement of maintaining a steady force in the range of $0.75\text{N}\pm 0.5\text{N}$ on the force sensor (see Figure 4.7). The force sensor was placed within the chair, positioned approximately where the heel of the animal’s palm would normally lie. A circle with a radius proportional to the force measured was used to provide visual feedback. A




Figure 4.7: Experimental setup for multi-tasking paradigm

teal donut was used to demarcate the accepted force limits. The cursor only appeared on the screen when the force applied was within the limits, and the trial was restarted without reward if the force applied was outside of the limits at any time during the course of the trial. To instruct the subject on the multi-tasking paradigm, we trained the subject on three intermediary stages: force-only, BMI/force block switching, BMI/force task switching.

1. Force-only stage: The subject was required to press and hold the force at one of three levels of force: 0.5, 0.75 and 1N (± 0.25 N at each level) for 400ms in order to receive reward for the trial. The subject was then required to release the force level back to baseline to start the next trial. The force target was randomly chosen each trial.
2. BMI/force block switching: The subject performed interleaving 10 minute blocks of force-only and BMI center-out.
3. BMI/force task switching: The subject performed interleaving trials of force-only and BMI center-out.

CLDA for multi-tasking: posture-adjusted vs. original

Since the placement of the force sensor affects the positioning of the subject's hand, it may be possible that the change in posture can render the previously trained decoder to no longer be usable. Hence, we also trained a separate "posture-adjusted" decoder, whereby CLDA was used to update the decoder during a pure BMI session in which the subject's hand was placed over the force sensor (even though force was not required).

4.5.2 Results

The intermediary training stages took 15 days in total, and multi-tasking sessions were performed for another 23 days. Figure 4.8 shows the average target acquisition rate and average reach times for each day for the various conditions, including the pure BMI sessions done before multi-tasking training began, as a comparison. Multi-tasking was initially performed with no target hold time requirement, and was increased to 200ms once performance improved. The overall performance using the posture-adjusted decoder steadily improved over a period of 8 days. While the average reach time appeared to increase over time, the target acquisition rate was actually improving since initially, the subject had a hard time maintaining the force, and so the only success trials were extremely quick reaches, leading to a low average reach time, but also low target acquisition rate.

After switching to the original, non-posture-adjusted decoder, multi-tasking performance initially dropped. Interestingly, performance readily increased back up prior levels over the next 4 days. During these sessions, CLDA was still performed at times when performance was consistently poor, but the force sensor was removed, to allow the subject to maintain their prior posture.

We subsequently increased the force requirement to $1\text{N}\pm 0.5\text{N}$, effectively doubling the minimum force required. As expected, target acquisition rate dropped substantially as the subject was not used to sustaining the higher force minimum. However, the subject was eventually able to apply the higher force level while modulating LFP activity, leading to higher target acquisition rate (although reach times remained roughly constant).

Neural activity

Figure 4.9 shows trial-averaged spectrograms for pure BMI, BMI trials during task-switching, multi-tasking with posture-adjusted decoder and multi-tasking with original decoder. Although a slight drop in performance was seen when the subject switched from the posture-adjusted decoder back to the original, the spectrograms were nearly identical, suggesting that both decoders were in fact quite similar. Interestingly, comparing pure BMI with multi-tasking, we see largely similar neural activity, except for a decrease in power in the 30–50 Hz range approximately 200ms after the Go cue. In fact, this decrease in 30–50 Hz range was also visible during the BMI trials in the task-switching paradigm, where the subject was required to apply force on the sensor in order to initiate the trial.

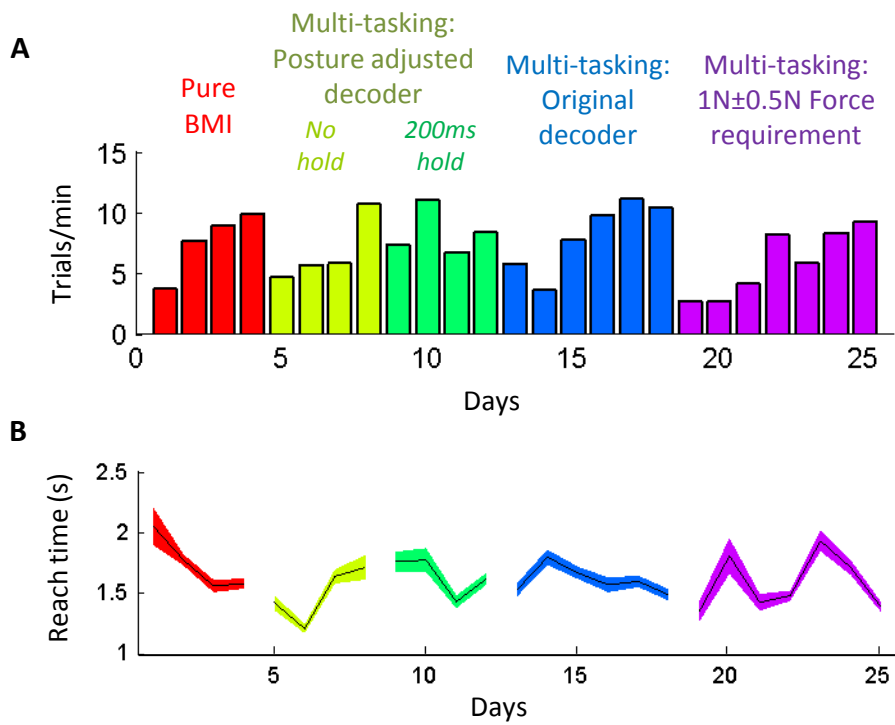


Figure 4.8: Behavioral performance for multi-tasking, measured in (A) target acquisition rate and (B) average reach time.

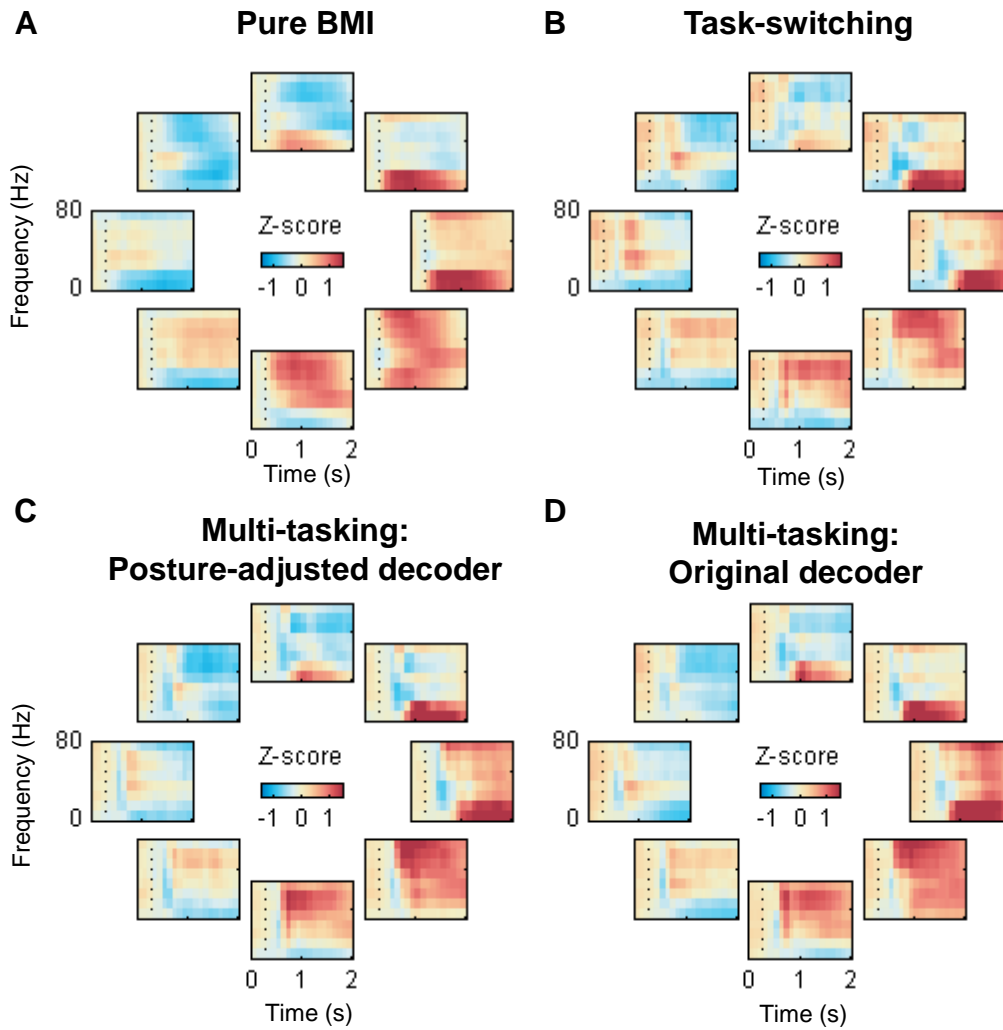


Figure 4.9: Trial-averaged spectrograms for (A) pure BMI, (B) BMI trials during task-switching, (C) multi-tasking with posture-adjusted decoder and (D) multi-tasking with original decoder.

4.5.3 Discussion

In summary, we show that LFP BMI can still be accomplished while simultaneously applying a sizable isometric force with their right hand. With practice, the target acquisition rate approached that of pure BMI. We also tested whether CLDA was needed with his hand on the sensor to account for posture differences (posture-adjusted decoder) and found that the original decoder largely worked just as well, suggesting minimal effect on the decoder due to posture. Comparing the trial-averaged spectrograms between pure BMI and multi-tasking, the primary difference was a decrease in 30–50 Hz activity shortly after the Go cue. This may be a result of the need to maintain the isometric force over the course of the trial. Overall, our results suggest that modulation of LFP is not necessarily tied to overt arm movements. Furthermore, our results indicate that BMI using LFP can still be used while the subject’s arm is performing a different (albeit, simple) task, which can be useful from a practical standpoint.

Chapter 5

Conclusion

Realizing robust closed-loop brain-machine interface control presents challenges in both neuroscience and engineering. The progress described here represents important steps towards understanding how the brain interacts with a BMI, and how to leverage that interaction to enable stable, long-term control.

5.1 Thesis contributions

Specifically, this thesis contributes to a number of areas regarding BMI control:

- Applying a metric from information theory, mutual information, we quantified the amount of information encoded in a neural ensemble about the intended direction of reach. We report two main findings. First, information encoding was largely redundant, consistent with past studies. Secondly, information encoded in the output neurons during BMI control was highest, even when they were located in the ipsilateral hemisphere where information encoding was lower during natural reaching [92, 93].
- We examined the connectivity of the neural network using a recently proposed information-theoretic metric known as directed information. We applied the metric extensively on simulated and rodent neural data and showed that directed information provides higher accuracy and more insight than conventional correlation methods [94]. Applying the method to neural data collected during BMI control, we found a gradual consolidation of the neural network connections that parallel the changes in behavioral performance and in directional tuning, providing further evidence suggesting that BMI control can lead to formation of a stable cortical map [95].

- We implemented 2-D closed-loop BMI control using local field potentials as input, with a CLDA algorithm (SmoothBatch) to adapt to the neural changes that occur during BMI operation. This study adds to the growing body of evidence that field potentials (LFP or ECoG) can be volitionally modulated well enough to provide 2-D cursor control. Moreover, we found substantial differences in the pattern of LFP modulations between the two subjects. Hence, adaptive methods such as CLDA may be useful to account for these inter-subject differences [96, 97].

5.2 Future work

Our preliminary result demonstrating volitional modulation of LFP spectral power through a BMI paradigm serves as a springboard for a number of other research directions. The relationship between spiking activity and LFP are still unclear; how independent are these two signal sources? Do they carry redundant or independent information? If spiking activity is coupled to the phase of particular frequencies of the LFP, how strong is this connection and can it be changed? Using a BMI paradigm, we can begin to explore these directions by monitoring the spiking patterns as a subject learns to volitionally modulate LFP oscillations and vice versa. These insights can in turn inform BMI decoder design that can leverage both spiking activity and LFP, either treating both as relatively independent sources of information to contribute to higher degrees of freedom, or using them as observations of the same process in order to improve signal-to-noise ratio and robustness.

Outside of BMI development, using the tools developed for inferring functional connectivity, we can move beyond analyzing the local network of neurons in M1 to understanding how connections may be formed with other brain structures that govern learning and higher cognitive processes. Combined with advances in multi-site long-term intracortical recordings, we are in a position to gain fundamental insight on how different parts of the brain communicate.

Bibliography

- [1] Green A M and Kalaska J F 2011 Learning to move machines with the mind. *Trends in Neurosciences* **34** 61–75
- [2] Shadlen M N and Newsome W T 1998 The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience* **18** 3870–3896
- [3] Averbeck B B and Lee D 2006 Effects of noise correlations on information encoding and decoding. *Journal of Neurophysiology* **95** 3633
- [4] Quinn C, Coleman T, Kiyavash N and Hatsopoulos N 2010 Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of Computational Neuroscience* **30** 17–44
- [5] Gilja V, Nuyujukian P, Chestek C A, Cunningham J P, Yu B M, Fan J M, Churchland M M, Kaufman M T, Kao J C, Ryu S I and Shenoy K V 2012 A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience* **15** 1752–1757
- [6] Orsborn A L, Dangi S, Moorman H G and Carmena J M 2012 Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE transactions on neural systems and rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society* **20** 468–477
- [7] Shpigelman L, Lalazar H and Vaadia E 2008 Kernel-ARMA for hand tracking and brain-machine interfacing during 3D motor control *Proc. Neural Inf. Process. Syst* pp 1489–1496
- [8] Taylor D M, Tillery S I H and Schwartz A B 2002 Direct cortical control of 3D neuroprosthetic devices | Science/AAAS. *Science* **296** 1829–1832
- [9] Bansal A K, Vargas-Irwin C E, Truccolo W and Donoghue J P 2011 Relationships among low-frequency local field potentials, spiking activity, and three-dimensional reach and grasp kinematics in primary motor and ventral premotor cortices. *Journal of Neurophysiology* **105** 1603–1619
- [10] Markowitz D A, Wong Y T, Gray C M and Pesaran B 2011 Optimizing the decoding of movement goals from local field potentials in macaque cortex. *Journal of Neuroscience* **31** 18412–18422
- [11] Rickert J, de Oliveira S C, Vaadia E, Aertsen A, Rotter S and Mehring C 2005 Encoding of movement direction in different frequency ranges of motor cortical local field potentials. *The Journal of Neuroscience* **25** 8815–8824

- [12] Gawne T J and Richmond B J 1993 How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience* **13** 2758–2771
- [13] Zohary E, Shadlen M N and Newsome W T 1994 Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370** 140–143
- [14] Maynard E M, Hatsopoulos N G, Ojakangas C L, Acuna B D, Sanes J N, Normann R A and Donoghue J P 1999 Neuronal interactions improve cortical population coding of movement direction. *J. Neurosci.* **19** 8083–8093
- [15] Oram M W, Hatsopoulos N G, Richmond B J and Donoghue J P 2001 Excess synchrony in motor cortical neurons provides redundant direction information with that from coarse temporal measures. *Journal of Neurophysiology* **86** 1700–1716
- [16] Hatsopoulos N G, Ojakangas C L, Paninski L and Donoghue J P 1998 Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proceedings of the National Academy of Sciences of the United States of America* **95** 15706–15711
- [17] Abbott L F and Dayan P 1999 The effect of correlated variability on the accuracy of a population code. *Neural Computation* **11** 91–101
- [18] Latham P E and Nirenberg S 2005 Synergy, redundancy, and independence in population codes, revisited. *Journal of Neuroscience* **25** 5195
- [19] Shamir M and Sompolinsky H 2011 Nonlinear population codes. *Neural Computation* **16** 1105–1136
- [20] Narayanan N S, Kimchi E Y and Laubach M 2005 Redundancy and synergy of neuronal ensembles in motor cortex. *Journal of Neuroscience* **25** 4207–4216
- [21] Fetz E E and Cheney P D 1980 Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *Journal of Neurophysiology* **44** 751–772
- [22] Ganguly K and Carmena J M 2009 Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology* **7** e1000153
- [23] Quiñero Quiroga R and Panzeri S 2009 Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience* **10** 173–185
- [24] Panzeri S, Senatore R, Montemurro M A and Petersen R S 2007 Correcting for the sampling bias problem in spike train information measures. *Journal of Neurophysiology* **98** 1064–1072
- [25] Rolls E T, Franco L, Aggelopoulos N C and Reece S 2003 An information theoretic approach to the contributions of the firing rates and the correlations between the firing of neurons. *Journal of Neurophysiology* **89** 2810–2822
- [26] Schneidman E, Bialek W and Berry M J 2003 Synergy, redundancy, and independence in population codes. *Journal of Neuroscience* **23** 11539–11553
- [27] Ashe J and Georgopoulos A P 1994 Movement parameters and neural activity in motor cortex and area 5. *Cerebral Cortex* **4** 590–600

- [28] Thach W T 1978 Correlation of neural discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. *Journal of Neurophysiology* **41** 654–676
- [29] Georgopoulos A P, Kalaska J F, Caminiti R and Massey J T 1982 On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience* **2** 1527–1537
- [30] Fu Q G, Flament D, Coltz J D and Ebner T J 1995 Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of Neurophysiology* **73** 836–854
- [31] Carmena J M, Lebedev M A, Crist R E, O’Doherty J E, Santucci D M, Dimitrov D F, Patil P G, Henriquez C S and Nicolelis M A L 2003 Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol* **1** 193–208
- [32] Barlow H 2001 Redundancy reduction revisited. *Network (Bristol, England)* **12** 241–253
- [33] Nudo R J, Milliken G W, Jenkins W M and Merzenich M M 1996 Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *Journal of Neuroscience* **16** 785–807
- [34] Kleim J A, Barbay S and Nudo R J 1998 Functional reorganization of the rat motor cortex following motor skill learning. *Journal of Neurophysiology* **80** 3321
- [35] Johnson M T, Mason C R and Ebner T J 2001 Central processes for the multiparametric control of arm movements in primates. *Current Opinion in Neurobiology* **11** 684–688
- [36] Jarosiewicz B, Chase S M, Fraser G W, Velliste M, Kass R E and Schwartz A B 2008 Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences* **105** 19486
- [37] Serruya M D, Hatsopoulos N G, Paninski L, Fellows M R and Donoghue J P 2002 Brain-machine interface: Instant neural control of a movement signal. *Nature* **416** 141–142
- [38] Pascual-Leone A, Amedi A, Fregni F and Merabet L B 2005 The plastic human brain cortex. *Annual review of neuroscience* **28** 377–401
- [39] Lee L, Harrison L M and Mechelli A 2003 A report of the functional connectivity workshop, dusseldorf 2002. *NeuroImage* **19** 457–465
- [40] Horwitz B 2003 The elusive concept of brain connectivity. *NeuroImage* **19** 466–470
- [41] Eguíluz V M, Chialvo D R, Cecchi G A, Baliki M and Apkarian A V 2005 Scale-free brain functional networks. *Physical Review Letters* **94** 018102
- [42] Hampson M, Peterson B S, Skudlarski P, Gatenby J C and Gore J C 2002 Detection of functional connectivity using temporal correlations in MR images. *Human Brain Mapping* **15** 247–262
- [43] Horwitz B, Rumsey J M and Donohue B C 1998 Functional connectivity of the angular gyrus in normal reading and dyslexia. *Proceedings of the National Academy of Sciences of the United States of America* **95** 8939–8944

- [44] Pfurtscheller G and Andrew C 1999 Event-related changes of band power and coherence: methodology and interpretation. *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society* **16** 512–519
- [45] Salvador R, Suckling J, Schwarzbauer C and Bullmore E 2005 Undirected graphs of frequency-dependent functional connectivity in whole brain networks. *Philosophical Transactions of the Royal Society B: Biological Sciences* **360** 937–946
- [46] Paninski L 2003 Estimation of entropy and mutual information. *Neural Computation* **15** 1191–1253
- [47] David O, Kiebel S J, Harrison L M, Mattout J, Kilner J M and Friston K J 2006 Dynamic causal modeling of evoked responses in EEG and MEG. *NeuroImage* **30** 1255–1272
- [48] Friston K J, Harrison L and Penny W 2003 Dynamic causal modelling. *NeuroImage* **19** 1273–1302
- [49] Kamiński M, Ding M, Truccolo W and Bressler S 2001 Evaluating causal relations in neural systems: granger causality, directed transfer function and statistical assessment of significance. *Biological cybernetics* **85** 145–157
- [50] McIntosh A and Gonzalez-Lima F 1994 Structural equation modeling and its application to network analysis in functional brain imaging. *Human Brain Mapping* **2** 2–22
- [51] Chen Z, Putrino D, Ghosh S, Barbieri R and Brown E 2011 Statistical inference for assessing functional connectivity of neuronal ensembles with sparse spiking data. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* **19** 121–135
- [52] Kim S and Brown E 2010 A general statistical framework for assessing granger causality *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on* pp 2222–2225 ISBN 1520-6149
- [53] Massey J L 1990 Causality, feedback and directed information. *Proceedings of the International Symposium on Information Theory and its Applications* 303–305
- [54] Granger C W J 1969 Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37** 424–438
- [55] Truccolo W, Eden U T, Fellows M R, Donoghue J P and Brown E N 2005 A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology* **93** 1074
- [56] Truccolo W, Hochberg L R and Donoghue J P 2009 Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nature Neuroscience* 105–111
- [57] Pillow J W, Shlens J, Paninski L, Sher A, Litke A M, Chichilnisky E J and Simoncelli E P 2008 Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454** 995–999
- [58] Pawitan Y 2001 *In All Likelihood: Statistical Modelling and Inference Using Likelihood* 1st ed (Oxford University Press, USA) ISBN 9780198507659

- [59] Song D, Chan R H, Marmarelis V Z, Hampson R E, Deadwyler S A and Berger T W 2009 Sparse generalized laguerre-volterra model of neural population dynamics *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE* 4555–4558
- [60] McGeorge A J and Faull R L 1989 The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* **29** 503–537
- [61] Charles J W 1986 Postsynaptic potentials evoked in spiny neostriatal projection neurons by stimulation of ipsilateral and contralateral neocortex. *Brain Research* **367** 201–213
- [62] Ganguly K, Dimitrov D F, Wallis J D and Carmena J M 2011 Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat Neurosci* **14** 662–667
- [63] Fetz E E 2007 Volitional control of neural activity: implications for brain-computer interfaces. *Journal of Physiology* **579** 571–579
- [64] Sanes J N and Donoghue J P 1993 Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proceedings of the National Academy of Sciences of the United States of America* **90** 4470–4474
- [65] Flint R D, Ethier C, Oby E R, Miller L E and Slutzky M W 2012 Local field potentials allow accurate decoding of muscle activity. *Journal of Neurophysiology* **108** 18–24
- [66] Flint R D, Lindberg E W, Jordan L R, Miller L E and Slutzky M W 2012 Accurate decoding of reaching movements from field potentials in the absence of spikes. *Journal of Neural Engineering* **9** 046006
- [67] Hwang E J and Andersen R A 2009 Brain control of movement execution onset using local field potentials in posterior parietal cortex. *Journal of Neuroscience* **29** 14363–14370
- [68] Engelhard B, Ozeri N, Israel Z, Bergman H and Vaadia E 2013 Inducing gamma oscillations and precise spike synchrony by operant conditioning via brain-machine interface. *Neuron* **77** 361–375
- [69] Flint R D, Wright Z A, Scheid M R and Slutzky M W 2013 Long term, stable brain machine interface performance using local field potentials and multiunit spikes. *Journal of Neural Engineering* **10** 056005
- [70] Schalk G, Kubánek J, Miller K J, Anderson N R, Leuthardt E C, Ojemann J G, Limbrick D, Moran D, Gerhardt L A and Wolpaw J R 2007 Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *Journal of Neural Engineering* **4** 264
- [71] Flint R, Wright Z and Slutzky M 2012 Control of a biomimetic brain machine interface with local field potentials: Performance and stability of a static decoder over 200 days *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* pp 6719–6722
- [72] Dangi S, Orsborn A L, Moorman H G and Carmena J M 2013 Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces. *Neural Computation* **25** 1693–1731

- [73] Haykin S S 2002 *Adaptive filter theory* 4th ed (Prentice Hall) ISBN 978-0132671453
- [74] Wu W, Black M J, Gao Y, Bienenstock E, Serruya M, Shaikhouni A and Donoghue J P 2003 Neural decoding of cursor motion using a kalman filter *Advances in Neural Information Processing Systems 15* (MIT Press) 133–140
- [75] Wang W, Collinger J L, Degenhart A D, Tyler-Kabara E C, Schwartz A B, Moran D W, Weber D J, Wodlinger B, Vinjamuri R K, Ashmore R C, Kelly J W and Boninger M L 2013 An electrocorticographic brain interface in an individual with tetraplegia. *PloS one* **8** e55344
- [76] Gage G J, Ludwig K A, Otto K J, Ionides E L and Kipke D R 2005 Naive coadaptive cortical control. *Journal of Neural Engineering* **2** 52–63
- [77] Heliot R, Venkatraman S and Carmena J M 2010 Decoder remapping to counteract neuron loss in brain-machine interfaces. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* 1670–1673
- [78] Mahmoudi B and Sanchez J C 2011 A symbiotic brain-machine interface through value-based decision making. *PloS one* **6** e14760
- [79] Li Z, O’Doherty J E, Lebedev M A and Nicolelis M A L 2011 Adaptive decoding for brain-machine interfaces through bayesian parameter updates. *Neural computation* **23** 3162–3204
- [80] Dangi S, Gowda S, Heliot R and Carmena J 2011 Adaptive kalman filtering for closed-loop brain-machine interface systems *IEEE/EMBS Conference on Neural Engineering (NER)* pp 609–612
- [81] Hochberg L R, Bacher D, Jarosiewicz B, Masse N Y, Simeral J D, Vogel J, Haddadin S, Liu J, Cash S S, Smagt P v d and Donoghue J P 2012 Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* **485** 372–375
- [82] Donoghue J P, Simeral J D, Kim S P, Friehs G M, Hochberg L R and Black M J 2007 Toward standardized assessment of pointing devices for brain-computer interfaces (Society for Neuroscience (SfN 2007))
- [83] Rouse A G, Williams J J, Wheeler J J and Moran D W 2013 Cortical adaptation to a chronic micro-electrocorticographic brain computer interface. *Journal of Neuroscience* **33** 1326–1330
- [84] Wander J D, Blakely T, Miller K J, Weaver K E, Johnson L A, Olson J D, Fetz E E, Rao R P N and Ojemann J G 2013 Distributed cortical adaptation during learning of a brain-computer interface task. *Proceedings of the National Academy of Sciences of the United States of America* **110** 10818–10823
- [85] Kajikawa Y and Schroeder C E 2011 How local is the local field potential? *Neuron* **72** 847–858
- [86] Crone N E, Miglioretti D L, Gordon B, Sieracki J M, Wilson M T, Uematsu S and Lesser R P 1998 Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. i. alpha and beta event-related desynchronization. *Brain* **121** 2271–2299
- [87] Fitts P M 1954 The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* **47** 381–391

- [88] MacKenzie I S 1989 A note on the information-theoretic basis of fitts' law. *Journal of Motor Behavior* **21** 323–330
- [89] Kim S P, Simeral J D, Hochberg L R, Donoghue J P and Black M J 2008 Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of Neural Engineering* **5** 455
- [90] Schalk G, Miller K J, Anderson N R, Wilson J A, Smyth M D, Ojemann J G, Moran D W, Wolpaw J R and Leuthardt E C 2008 Two-dimensional movement control using electrocorticographic signals in humans. *Journal of Neural Engineering* **5** 75–84
- [91] Wolpaw J R and McFarland D J 2004 Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proceedings of the National Academy of Sciences of the United States of America* **101** 17849–17854
- [92] So K, Ganguly K, Gastpar M and Carmena J 2011 Redundant information encoding in primary motor cortex during motor tasks *2011 5th International IEEE/EMBS Conference on Neural Engineering (NER)* pp 212–215
- [93] So K, Ganguly K, Jimenez J, Gastpar M C and Carmena J M 2012 Redundant information encoding in primary motor cortex during natural and prosthetic motor control. *Journal of computational neuroscience* **32** 555–561
- [94] So K, Gastpar M and Carmena J M 2011 Assessing directed information as a method for inferring functional connectivity in neural ensembles. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference* **2011** 7324–7327
- [95] So K, Koralek A C, Ganguly K, Gastpar M C and Carmena J M 2012 Assessing functional connectivity of neural ensembles using directed information. *Journal of neural engineering* **9** 026004
- [96] Dangi S, So K, Orsborn A L, Gastpar M C and Carmena J M 2013 Brain-machine interface control using broadband spectral power from local field potentials. *Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference* **2013** 285–288
- [97] So K, Dangi S, Orsborn A L, Gastpar M C and Carmena J M 2013 Subject-specific modulation of local field potential spectral power during brain-machine interface control in primates. *Journal of Neural Engineering* (in press)
- [98] Song D, Chan R H, Marmarelis V Z, Hampson R E, Deadwyler S A and Berger T W 2007 Non-linear dynamic modeling of spike train transformations for hippocampal-cortical prostheses. *Biomedical Engineering, IEEE Transactions on* **54** 1053–1066

Appendix

A.1 Raised cosine basis functions

$$b_j = \begin{cases} \frac{1}{2} \cos \left(a \log(t) - \frac{\pi}{2} (j - 1) \right) + \frac{1}{2} & a \log(t) \in \left[\frac{\pi}{2} (j - 1) - \pi, \frac{\pi}{2} (j - 1) + \pi \right] \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.1})$$

where $t \in \{1, \dots, 100\}$. For the intrinsic history effect, j ranges from 1 to 10, and $a = 3.5$. For the extrinsic history effect, j ranges from 1 to 4 and $a = 1.5$. Figure A.1 is a plot of the two sets.




Figure A.1: (left) Intrinsic history basis set and (right) extrinsic history basis set using raised cosine functions.

A.2 Sparse generalized Laguerre-Volterra model (SGLVM)

Under this model, we generate spike trains where the spiking probability μ_i for neuron Y is determined as follows:

$$\begin{aligned} \Phi^{-1}(\mu) = & \alpha_0 + \sum_{l=1}^4 \alpha_l \langle v_l \cdot Y_{i-100}^{i-1} \rangle + \sum_{m=1}^4 \beta_m \langle v_m \cdot X_{i-100}^{i-1} \rangle \\ & + \sum_{j=1}^4 \sum_{k=1}^j \beta_{j,k} \langle v_j \cdot X_{i-100}^{i-1} \rangle \langle v_k \cdot X_{i-100}^{i-1} \rangle \end{aligned} \quad (\text{A.2})$$

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp\left(-\frac{x^2}{2}\right) dx \quad (\text{A.3})$$

where $\Phi^{-1}(\cdot)$ is the inverse Gaussian CDF, X is the activity of another simulated neuron. The set $\alpha_0, \dots, \alpha_4$ are preset parameters that determine the effect of past spiking activity, while the sets β_1, \dots, β_4 and $\beta_{j,k}$ for $j, k \in \{1, \dots, 4\}$ determine the strength of the linear and nonlinear coupling between the two spike trains. v is a set of orthonormal Laguerre basis functions of the following form [98]:

$$v_j(t) = \begin{cases} (-1)^t \alpha^{\frac{2j-t}{2}} (1-\alpha)^{\frac{1}{2}} \\ \times \sum_{k=0}^t (-1)^k \binom{t}{k} \binom{2j}{k} \alpha^{t-k} (1-\alpha)^k & 0 \leq t < 2j \\ (-1)^{2j} \alpha^{\frac{t-2j}{2}} (1-\alpha)^{\frac{1}{2}} \\ \times \sum_{k=0}^{2j} (-1)^k \binom{t}{k} \binom{2j}{k} \alpha^{2j-k} (1-\alpha)^k & 2j \leq t \leq R \end{cases} \quad (\text{A.4})$$

where $\alpha = 0.6$, $R = 100$, and $j \in \{0, \dots, 3\}$, see figure A.2 for plot. To determine physiologically realistic values for the parameters ($\alpha_0, \dots, \alpha_4, \beta_1, \dots, \beta_4$ and $\beta_{j,k}$ for $j, k \in \{1, \dots, 4\}$), we first fit the SGLVM on experimental data. We use spike trains recorded from M1 of a rhesus monkey while performing a center-out reaching task using natural arm movement (see Section 3.5.1 for more details). The parameters are then determined by maximizing the likelihood function formed from (A.2) - (A.4), with the exception of α_0 , which is manually adjusted to provide an average firing rate of ~ 5 -10Hz. This procedure is repeated using six randomly selected pairs of neurons to obtain different parameter sets that reflect a variety of excitatory and inhibitory connections.

A.3 Calculating modulation depth

To compute modulation depth, we first calculate the mean firing rate for each target. A cosine tuning curve is fitted to the mean firing rates by estimating the set of coefficients in (A.5).




Figure A.2: Laguerre basis function set used in the simulation model.

$$f = B_1 + B_2 \sin \theta + B_3 \cos \theta \quad (\text{A.5})$$

$$MD = \sqrt{B_2^2 + B_3^2} \quad (\text{A.6})$$

where MD is the modulation depth. A standard F-test is used to assess the goodness of fit, and the modulation depth is subsequently computed if the fit is significant ($p < 0.05$).