# Optimizing electrical brain stimulation for seizure disorders

A Dissertation

SUBMITTED TO THE FACULTY OF

UNIVERSITY OF MINNESOTA

BY

Vivek Nagaraj

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Theoden I. Netoff

March 2017

# Acknowledgements

I would like to acknowledge the following people for training and support during my PhD:

*Dr. Theoden I Netoff* – Thanks so much for all of the training. I have enjoyed the opportunity to learn so many different subject areas in science and engineering under your mentorship.

*Netoff Lab* – Thank you for all the training and great conversations about science and life as a scientist! Dr. Oscar Miranda-Dominguez, Dr. Bryce Beverlin, Dr. Abbey Holt, Tyler Stigen, Jennifer Zick, Dr. David Darrow, Ken Louie, Logan Grado, Pantea Moghimi Nicholas Roberts, Benjamin Johnson, Tyler Nordmann, Pratik Parikh, Sarah West, and David Molinari.

*Family* – Thanks so much to my amazing wife Kseniya Nagaraj for providing incredible support during my PhD. I am also very thankful for all the support from my father Kesavamurthy, mother Indira, brother Aaditya, sister-in-law Apoorva, uncle Krishna, and cousin Nitin.

*Friends* – Thank you to my GPN class mates, John Paton, present and past Directors of Graduate Studies, and most importantly the GPN program for giving me the opportunity to be a student.

In addition, I would like to acknowledge my funding sources and my colleagues around the world for the many engaging conversations at conferences and workshops about Neuroscience, Epilepsy and Neuromodulation.

# Dedication

This dissertation is dedicated to my family.

# Abstract

Approximately 1% of the world population is afflicted with Epilepsy. For many patients, antiepileptic drugs do not fully control seizures. Electrical brain stimulation therapies have been effective in reducing seizure rates in some patients. While current neuromodulation devices provide a benefit to patients, efficacy can be improved by optimizing brain stimulation so that the therapy is tuned on a patient by patient basis. One optimization approach is to target deep brain regions that strongly modulate seizure prone regions. I will present data on the effects of stimulation of two different anatomical regions for seizure control, and establish my experimental platform for testing closed-loop algorithms.

There are two general methods to implementing closed-loop algorithms to modulate neural activity: 1) Model-free algorithms that require a learning period to establish an optimal mapping between neural states and best therapeutic parameters, and 2) Model-based algorithms that use forward predictions of the neural system to determine the appropriate stimulation therapy to be administered. In this thesis, I will propose and test two closed-loop control schemes to control the brain activity to prevent epileptogenic activity while reducing stimulation energy. I will also present techniques to remove stimulation artifacts so that neural biomarkers can be measured while simultaneously applying stimulation. The methods I will present could potentially be implemented in next generation electrical brain stimulation hardware for seizure disorders and other neurological diseases.

# Table of Contents

# List of Tables

# List of Figures

# Chapter 1 Introduction

*Excerpts, including text and figures from (Nagaraj et al., 2015) have been included in this chapter. This is a non-final version of an article published in final form in The Journal of Clinical Neurophysiology. The final peer-reviewed manuscript can be found at:

http://journals.lww.com/clinicalneurophys/Abstract/2015/06000/Future_of_Seizure_Prediction_and_Intervention__.3.aspx

## *1.1 Introduction*

Epilepsy affects nearly 3 million people in the United States, with nearly 500 new cases of epilepsy diagnosed every day (Hauser et al., 1993; CJL and AD, 1994). Current antiepileptic drugs can have major negative side effects, and approximately one third of patients with epilepsy are drug refractory (Kwan et al., 2010). Of those, only patients with a well-localized focus in an area outside of the eloquent cortex are good surgical candidates. For the remaining patients, there are few options. Closed-loop therapies may improve seizure control while reducing or eliminating side effects by limiting the therapy to times when the patient is in need. Closed-loop seizure therapies may also allow for stronger therapy doses that cannot be delivered chronically. The development of on-demand approaches will require effective seizure prediction or early detection algorithms and optimized intervention strategies while being reliable and safe for chronic implantation in humans. This thesis discusses both the progress and the future directions for each of these areas.

## *1.2 Need for new seizure detection and prediction devices*

Although many therapeutic devices have been developed for epilepsy, few devices have made it to clinical trials. There is a need for robust and accurate seizure detection and prediction devices. Self-reporting by patients of their seizures is often poor when compared with the detection of electrographic events, in part, because consciousness may be affected by the seizure (Heck et al., 2014). A monitoring device could therefore

dramatically improve the assessment of therapy both in treating patients and in clinical trials. Furthermore, a device that could detect changes in physiology before a seizure and deliver therapies to prevent the seizure would be transformative, enabling new approaches to treating epilepsy.

## *1.3 Defining seizure prediction*

Seizure prediction has had a long and storied history and has been well reviewed elsewhere (Litt and Lehnertz, 2002; Iasemidis, 2003; Mormann et al., 2007). A significant advance for the field came with simply defining "seizure prediction." Prediction is defined as identifying an event after which a seizure will occur within a fixed period of time. In contrast, if a heightened risk of seizure occurring is identified without a specific time window, it is considered "seizure forecasting." Four key criteria to measure the efficacy of a seizure prediction algorithm have been proposed: (1) developing algorithms on long-term recordings from patients, (2) assessing sensitivity and specificity about a range of prediction horizon times and the portion of time under false warning, (3) proving that the algorithm can perform above chance level by using statistical techniques, and (4) determine generalizability of the algorithm by testing it on out-of-sample data (Mormann et al., 2007).

## *1.4 Future of seizure therapy: closing the loop*

A monitoring device that can deliver closed-loop treatment enables new approaches to develop more personalized therapies. There are important considerations at all phases in

developing a closed-loop treatment: (1) the signals to be measured, (2) the features to be extracted and classification algorithms, (3) optimization of treatment, (4) the therapeutic actions to be taken, and (5) the devices to implement closed-loop therapies. These steps form a therapeutic loop, illustrated in Figure 1. Although seizure prediction and closed-loop therapies may be the ultimate goal, there are benefits to be gained in improving each phase of this loop. The development of closed-loop therapies may also help elucidate the effects of intervention on underlying biologic processes. This chapter outlines the current state and discusses potential future directions of seizure prediction and closed-loop therapies.

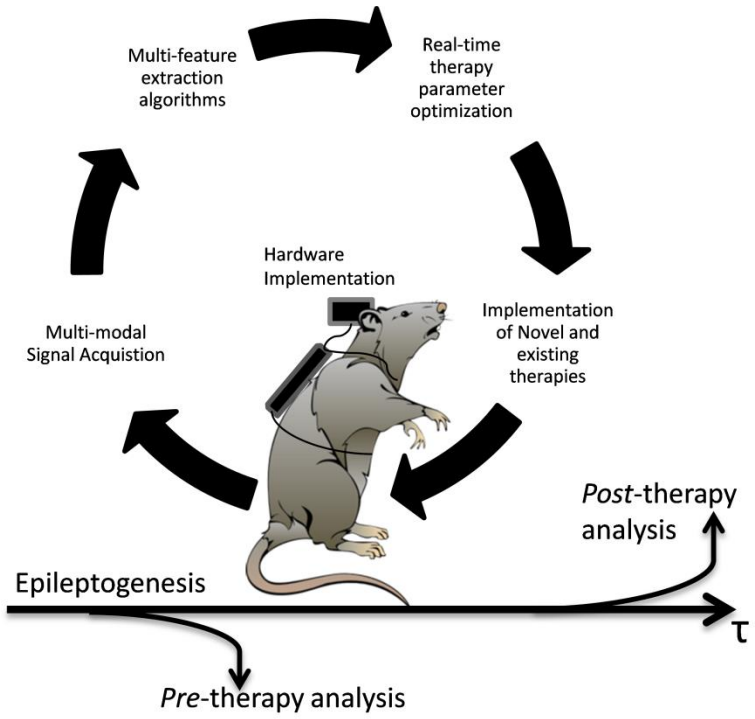


Figure 1. Hypothetical closed-loop experimental protocol for suppressing seizures using multimodal recordings of physiologic activity. From the recordings, features are extracted and a classifier is applied to detect seizure activity or a pre-seizure state for prediction. On detection of an event, the device triggers a therapy.

## *1.5 Detection of Pre-ictal and ictal states*

### 1.5.1 Signal Modalities

Scalp EEG and electrocorticography, measured at the surface of the brain, have been the mainstays for seizure prediction since the 1970s (Viglione and Walsh, 1975). The use of linear (Rogowski et al., 1981) and nonlinear methods (Iasemidis et al., 1990) helped conceptualize the "preictal" state as a period in which the brain exhibits an increasing prevalence of seizure-like behavior before seizure onset. However, it is questionable whether these signals contain sufficient information for accurate seizure prediction. Other recording methods, such as penetrating microelectrodes that provide higher spatial resolution and wider spectral content, may be necessary to better characterize the neural activity at the focus. Simultaneous macroelectrode and microelectrode recordings have revealed microbursts of activity that are not seen by the macroelectrodes alone (Schevon et al., 2009; Stead et al., 2010; Truccolo et al., 2011; Viventi et al., 2011). Applying these microbursts as an additional feature set could potentially increase predictive power. Furthermore, single-unit recordings may provide additional features for improved seizure prediction (Bower and Buckmaster, 2008; Keller et al., 2010; Bower et al., 2012; Einevoll et al., 2012).

Beyond passive recording techniques, important information may be gained through evoked potentials. Single electrical pulses may be used to detect changes in excitability before a seizure. Changes in excitability over space may also be used to localize a seizure focus. Excitatory synaptic activity can increase in epileptic tissue

(Avoli et al., 2005) through a variety of mechanisms, including potentiation of synaptic AMPA receptors (Abegg et al., 2004; Debanne et al., 2006; Lopantsev et al., 2009; Müller et al., 2013), potentiation of extrasynaptic NMDA receptors (Frasca et al., 2011; Müller et al., 2013), and an increase of intrinsic excitability of neurons (Tang et al., 2012). Importantly, a single pulse may evoke different epileptic responses depending on the level of excitability (Demont-Guignard et al., 2012). For example, an electrical brain-stimulation paradigm was successfully used for estimating both the seizure onset sites and the time to ictal transition in human temporal lobe epilepsy (Kalitzin et al., 2005). Evoked potentials may therefore be valuable for closed-loop stimulation protocols and optimization of stimulation parameters (McIntyre et al., 2004; Kent and Grill, 2012).

Additionally, improved seizure prediction may be achieved using other recording modalities. Several physiologic changes have been observed preceding electrographic seizure onset; in animal models and humans, changes in blood flow, blood oxygenation, and metabolism have all been shown to precede a seizure (Schwartz, 2007; Zhao et al., 2011; Patel et al., 2013). Additionally, in humans, electrochemical studies have shown distinct glutamate and adenosine dynamics associated with seizures (Van Gompel et al., 2014). Heart rate monitors and accelerometers have also been used to extract salient biomarkers of seizures (Zijlmans et al., 2002; Nijsen et al., 2005; Lockman et al., 2011). Incorporation of these other modalities with EEG may significantly improve seizure prediction and intervention.

### 1.5.2 Features

Although many features extracted from EEG signals have been used for seizure prediction (Mormann et al., 2007), no single or univariate measure has successfully characterized a pre-seizure state with high sensitivity and specificity. Individual features may have some small predictive power; however, as an aggregate, the feature set may achieve stronger predictive power. For example, changes in EEG at the focus may only be understood in the context of patterns seen in surrounding areas. Therefore, seizure prediction may benefit from a multivariate approach. A drawback to a multivariate approach, however, is that the relationship between the measured features and the predictions may be complex and unintuitive. It may therefore be difficult for clinicians to correlate changes in the raw signal with abstract features extracted through signal processing techniques. Similarly, classification of multivariate data for seizure prediction can be an engineering black box approach that provides limited insight to the mechanisms underlying the seizure onset.

An alternative approach is to fit physiologically realistic models to the data. The parameters of the model can then be used as the feature set. This may provide better insight into the relationship between mechanisms, such as changes in connection strengths between nodes, and state changes, such as preictal to ictal. Another challenge of a multivariate approach is that it is easy to produce an immense parameter space that cannot be fully explored. This is problematic for several reasons. First, a larger dimensional space can allow classifiers to overfit the data in training, resulting in poor

classification on new data. Second, as the number of features increases, the computational load increases as well. Therefore, when developing seizure prediction algorithms, the processing power of implantable devices should be considered.

There are two approaches that can be used to reduce feature space. One approach is to limit the number of measurements and features to those that provide some predictive power, for example, by using receiver operating characteristic analysis or Fisher discriminant analysis to select features. The other is to combine features in a way to reduce computational load, such as principal component analysis. Managing computational load is an important aspect of implementing seizure prediction and detection on implantable devices that needs to be developed further.

### 1.5.3 Linear vs Nonlinear features

Linear features, such as mean, standard deviation, and power spectral density, are well understood. However, there may be physiologic signals that cannot be well characterized by linear features. With the advent of chaos theory, many new measures were developed to analyze nonlinear systems, such as the Lyapunov exponent and the fractal dimension of the data. Presumably, these measures can detect changes undetectable with linear measures. However, in head-to-head comparisons, linear features often outperform the nonlinear ones in detecting a pre-ictal state (Jerger et al., 2003; Mormann et al., 2005). This is because linear features are extremely robust to noise (Netoff et al., 2004). Clearly, nonlinear measures may have great value for seizure prediction (Iasemidis, 2003) and should not be ignored; however, they must be compared in performance and

computational costs with linear measures. Furthermore, for implementation in an implantable device, linear features may be the only option because of current computational constraints.

## *1.6 Intervention*

In this section, we will review the need for new stimulation targets and interventional therapies such as optogenetics. We will also discuss methods for optimizing therapies using closed-loop approaches. Closed-loop therapies may not only provide hope for patients but also provide critical insight into the mechanisms underlying ictogenesis and epileptogenesis, which may lead to new approaches to therapies.

### 1.6.1 Stimulation Target Selection

Two different approaches have been tried for target selection in recent clinical trials. One approach (conducted by Medtronic) was to target an area with wide neuromodulatory effects (Fisher et al., 2010), the anterior nucleus of the thalamus (Medtronic, Minneapolis, MN). The other was to stimulate the focus directly (Heck et al., 2014) led by NeuroPace (NeuroPace, Mountain View, CA). Surprisingly, in both cases, there was about a 40% decrease in seizure frequency on average over a 3-month period. Even stimulation of the vagus nerve has demonstrated approximately 30% decrease in seizure frequency in patients after a 3-month blind period (Handforth et al., 1998). Interestingly, the efficacy of stimulation from each of these studies increased at the end of a 2-year follow-up (Fisher et al., 2010; Morris et al., 2013; Heck et al., 2014). These long-term

effects indicate that there may be direct and indirect mechanisms on seizure suppression caused by the stimulation.

Animal experiments have also tested many other promising stimulation locations that have not yet been tested in full clinical trials, such as the subthalamic nucleus (Loddenkemper et al., 2001; Chabardès et al., 2002), medial temporal structures (Vonck et al., 2005; Tellez-Zenteno et al., 2006; Velasco et al., 2007), centromedian thalamic nuclei (Mina et al., 2013; Pasnicu et al., 2013), and hippocampal commissural fiber (Koubeissi et al., 2013; Toprani and Durand, 2013) among others (Fisher, 2013). On-demand electrical stimulation has also shown to be effective in suppressing seizures in animal models (Bikson et al., 2001; Schiller and Bankirer, 2007; Good et al., 2009; Nelson et al., 2011).

Although many patients receive benefits, very few become seizure free. A single target may not be sufficient to control seizures. In some patients, seizures emerge from the interaction of multiple nodes within a network. Therefore, stimulation of more than one node in the network simultaneously may be required to improve seizure control. One approach to simultaneously modulate multiples nodes within a seizure network is to stimulate nuclei or fiber tracts that have widespread modulatory effects. Ideally, these targets would be far away from major blood vessels, to minimize risk during surgery, and have little effect on cognitive and motor functions.

### 1.6.2 Closed-loop intervention strategies

Different neurophysiologic mechanisms are engaged by DBS depending on the

stimulation localization, frequency, intensity, duration, and pattern. As noted above, there is large clinical discrepancy in the efficacy between patients, especially if they suffer from different subtypes of epilepsies. It is difficult to tune stimulus parameters, because the effect on seizure frequency cannot immediately be seen in the clinical setting. It may require statistical analysis over months to identify an effect of a stimulus parameter. Therefore, new methods must be developed for fine tuning of stimulation parameters in a patient-specific manner to maximize therapeutic effects.

How can electrical stimulation be optimized for a patient? There are 3 approaches to closed-loop intervention, as illustrated in Figure 2: (1) a controller that slowly adapts stimulation parameters over time to maximize therapy (e.g., a gradient descent controller) (Panuccio et al., 2013), (2) on-demand therapy (Armstrong et al., 2013; Krook-Magnuson et al., 2013; Heck et al., 2014), and (3) physiologically adapting closed-loop neuromodulation (Wilson et al., 2011; Little et al., 2013; Montaseri et al., 2013).

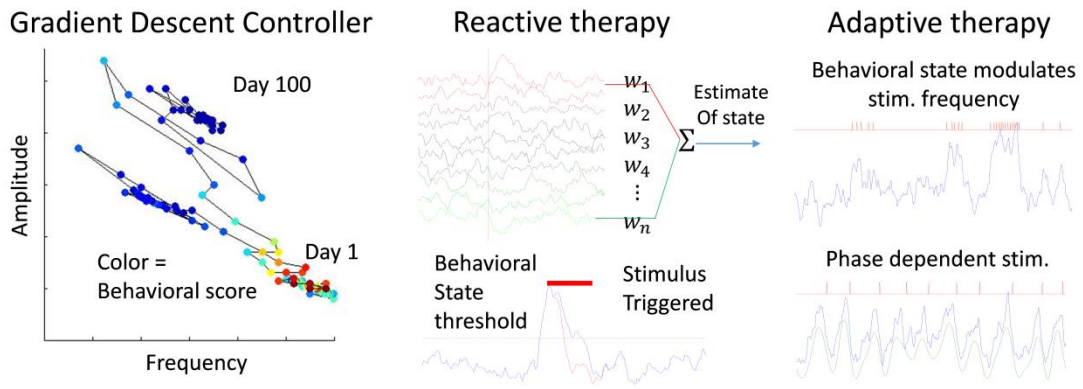


Figure 2. Types of closed-loop controllers. Left, a gradient descent controller adjusts the parameters and measures the effect, such as seizure frequency. Through small changes on a daily basis, the algorithm can traverse a parameter landscape to find an optimal solution. In a reactive therapy (middle), signals are processed when an event is detected in the estimated state of the patient. In this example, a stereotyped stimulus is applied; however, statistical tools can be applied to optimize stimulus parameters for reactive therapy. In an adaptive therapy (right), the stimulus is modulated by the state of the patient or can be used to trigger phasic stimulation with a millisecond precision.

### 1.6.3 Optimization of Stimulation Parameters

Optimization can be achieved by a gradient descent method (Figure 2, left) with offline analysis of data. This algorithm would establish a functional relationship between the parameters and the symptoms and iteratively adjust the parameters of the therapy to improve the symptoms. Adaptive reinforcement learning algorithms are well suited for problems with noisy, nonlinear, and nonstationary signals (Prokhorov and Wunsch, 1997; Panuccio et al., 2013). This approach requires the algorithm to "learn" from its mistakes to descend upon a local solution.

One concern with implementing a closed-loop algorithm is whether it is stable and safe. It is often more acceptable to implement a closed-loop algorithm that essentially does the same, but more formally, as a clinician or the patient might do given a user

interface. Optimization approaches like these have been tested in animal models with some success (Panuccio et al., 2013) but are yet to be used in a clinical trial. The first step toward obtaining Food and Drug Administration approval for a closed-loop device is in an advisory role to the patient or clinician, who makes the parameter changes. With proven safety of the device in an advisory role, it then may be possible to implement a fully automatic version.

Closed-loop Vagal Nerve Stimulation (VNS) optimization of stimulation parameters can be explored by controlling the activation of different vagus nerve fiber populations. An autonomous neural control algorithm has been developed to map the stimulation parameters to nerve activation profiles (Ward et al., 2014). Given this map, the algorithm dynamically adjusts the stimulus to maintain the nerve activation at a therapeutic level over time. This closed-loop algorithm adjusts stimulus parameters to maintain a stable stimulus response over time.

### 1.6.4 On-demand stimulation

The reactive nerve stimulation by NeuroPace stemmed out of the closed-loop algorithms used in cardiac pacemaking. NeuroPace's Responsive Neurostimulator delivers a stereotyped stimulus pattern when an abnormal event in the EEG is detected (Morrell, 2006; Morrell and Group, 2011). In clinical trials, the stimulus parameters were set by the clinician. The device is closed loop in which the stimulation is delivered on demand (i.e., is a responsive stimulator), but it is limited to fixed stimulation parameters (Figure 2, middle). The success of the responsive neurostimulator enables it to be a platform to test

and implement new therapies. The responsive neurostimulator device could provide a platform to test optimization algorithms (such as those discussed above) to tune stimulus parameters to maximize effects and could be coupled to other therapeutic modalities such as micropumps for drug delivery or optogenetics. It also could provide a platform to test optimization algorithms to tune stimulus parameters to maximize effects.

### 1.6.5 Physiologically adapting closed-loop neuromodulation

Closed-loop neuromodulation generally uses fixed stimulus patterns and is not yet adapted to the patient. Therapeutic efficacy may be improved by providing therapy modulated by physiologic activity of the patient. Closed-loop methods that adjust the timing of stimulation based on measured physiologic activity (Ward et al., 2014) have been proposed for treatment of tremor but have not yet been applied to epilepsy. This may be for several reasons including that seizures are changing so rapidly that it is difficult to have a clear signal upon which to determine the stimulation timing. To develop an optimal control algorithm, a model of the activity and response to stimulation must be made, for which a controller can be designed (e.g., a firing rate model of different neuronal populations involved in the seizure focus). Once a model is fit to the data and a control objective is defined, there are several approaches to designing optimal control algorithms. For linear and time invariant systems, there are well-developed engineering tools for designing closed-loop controllers (Ogata, 2010). Because EEG changes on a second-by-second basis, new adaptive nonlinear controllers need to be developed. Dynamic programming is a general approach to finding an optimal solution to

a control problem (Bellman and Rand Corporation, 1957). Dynamic programming has been used to design optimal stimulus waveforms to minimize energy while maximally perturbing neuronal spikes (Nabi et al., 2013b). However, for systems with more than a few parameters, this approach becomes intractably difficult to solve. Therefore, there is a great need to develop minimalist models that describe the pathologic activity that can be used in control design algorithms.

### 1.6.6 Novel therapies

Aside from direct electrical stimulation strategies discussed above, other closed-loop approaches should be explored. For example, noninvasive closed-loop neuromodulation technologies using EEG biofeedback have also shown promising results in controlling seizures (Sterman, 2000; Nagai et al., 2004; Tan et al., 2009). EEG biofeedback, or neurofeedback, modulates EEG activity at specific bandwidths through means of operant conditioning. With these promising results, and because it is noninvasive, neurofeedback should be investigated further.

Closed-loop approaches have been applied to other therapeutic modalities as well. For example, on-demand transcranial–electrical stimulation has been shown to terminate absence seizures in a rodent model (Berenyi et al., 2012). Focal cooling has long been used by surgeons to stop seizures during surgery (Rothman et al., 2005) and has been shown to be effective in a rodent seizure model (Hill et al., 2000; Yang and Rothman, 2001; Yang et al., 2003; Yang et al., 2006) and suppresses tumor-related epileptic discharges (Karkar et al., 2002). Design of new implantable cooling therapy devices in

humans would therefore be of a great interest for refractory focal epilepsies (Smyth and Rothman, 2011). Caged drugs that can be uncaged optically on demand (Rothman et al., 2007) and optogenetic approaches have also shown promising results in seizure control but have not yet been tested in humans.

# Chapter 2 Determining optimal deep brain targets for seizure control

## *2.1 Introduction*

Recent clinical trials have tested the efficacy of electrical stimulation of the superior anterior nucleus of the thalamus (SANTE) for seizure suppression. While stimulation in this area was effective for some patients, it was not for others (Fisher et al., 2010). New stimulation sites may provide alternative targets for these patients. Stimulation targets may be divided into grey matter targets and white matter tracts. Grey matter targets are functionally distinct units of cells that are often a node in a larger network. In this chapter, we will first discuss the Endopiriform Nucleus (ENc) as a potential deep brain grey matter target for seizure control. White matter tracts have regions have diffuse activation over areas they connect. Stimulation of axonal fibers drive strong depolarization events at target cell body regions. The second deep brain target we will use for seizure suppression is the Ventral Hippocampal Commissure (VHC). In this chapter, we will reproduce seizure suppression data from literature, and set up the animal model for all further *in vivo* experiments.

### 2.1.1 Endopiriform Nucleus

Electrical stimulation of the Piriform Cortex (PC) has been to shown to modulate seizure activity (Hoffman and Haberly, 1996; Zhu-Ge et al., 2007) presumably through its strong projections to the hippocampal formation, the origin of the seizures. The ENc is located immediately subjacent to the PC, on the inferior surface of the rodent telencephalon, and at the mesial surface of the temporal stem in the human brain (see Figure 3). In rodents, the ENc is the deep layer of the piriform cortex which is in the anterior ventro-lateral

aspect of the brain.

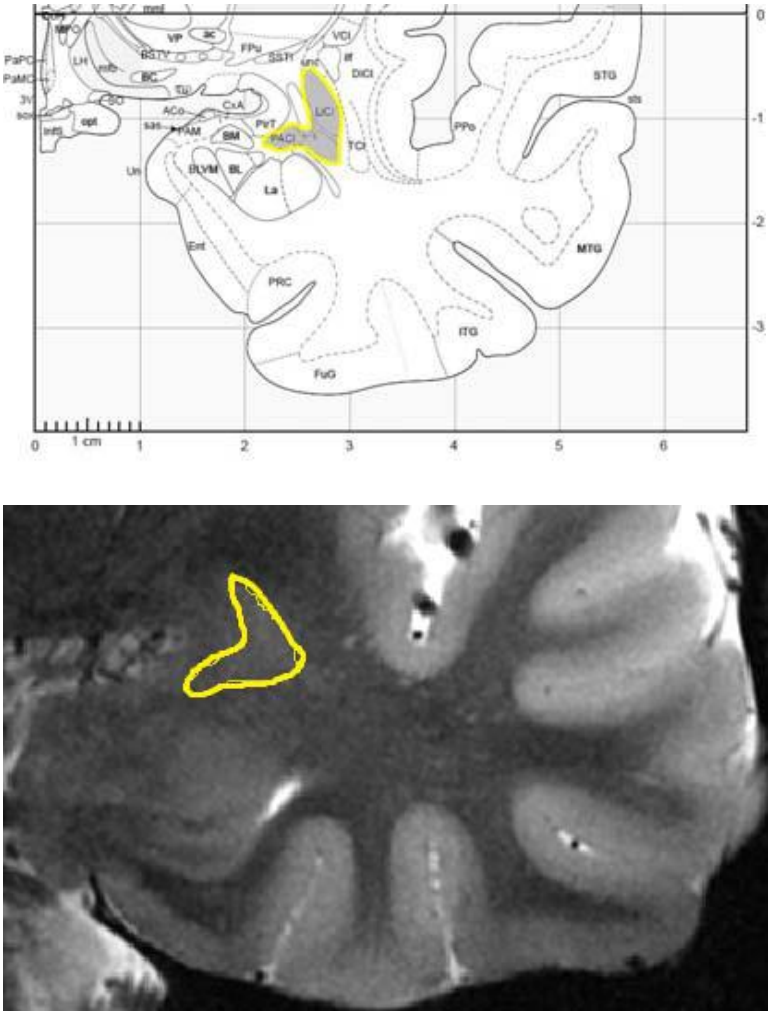


Figure 3. Location of the Endopiriform Nucleus in humans. This coronal T2-weighted image of a healthy control subject shows the endopiriform nucleus segmentation (in yellow). These images are the first to routinely detect ENc and distinguish it from amygdala in humans; they were acquired at the UMN CMRR's facility which provided the 7 Tesla MRI with resolution of 0.2 x 0.2 mm.

The ENc has dense reciprocal connections with the overlying PC, and it is likely that the ENc and PC act together in seizure antagonism (or promotion). For human DBS electrode implantation in the ENc is expected to be much safer than implantation in the

PC, given the proximity of the ENc to the amygdala which is routinely implanted with depth electrodes for seizure recoding via a lateral approach, which avoids sites of large blood vessels (Ross et al., 1996), and also the proximity of the PC to the middle cerebral artery.

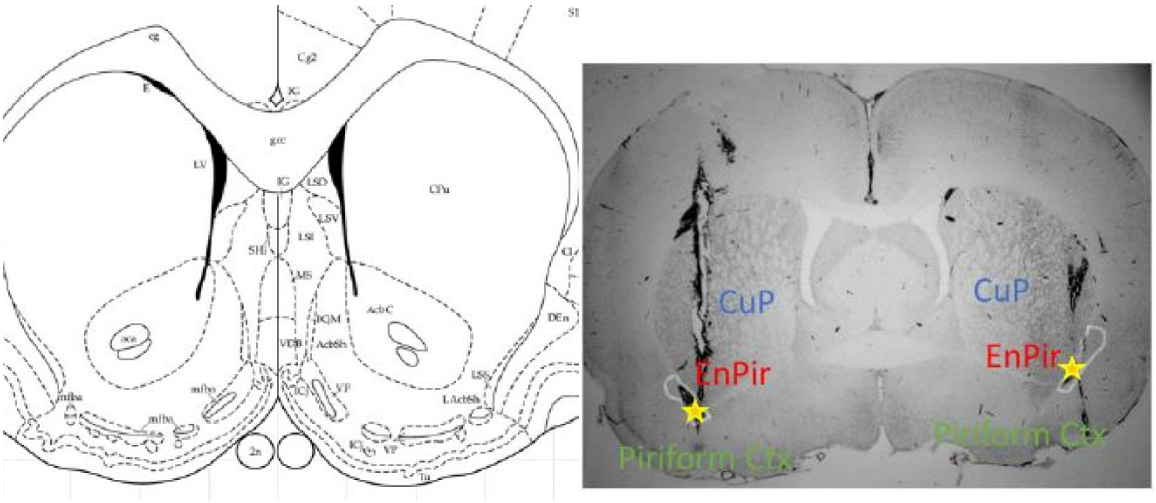


Figure 4. Endopiriform Nucleus in Sprague Dawley rats. Coronal section showing the caudate and putamen (CuP), endopiriform nucleus (ENc) and Piriform Cortex in a rat brain atlas (left) and bright field image of rat brain slice (right).

Other potential DBS sites have been proposed and studied experimentally or in clinical trials for epilepsy therapy, including the subiculum (Huang and van Luijtelaar, 2013), the anterior thalamic nuclei (Fisher et al., 2010), and the centro-medial thalamic nucleus (Good et al., 2009). The ENc may be a more effective site for modulation of limbic system seizures than are these other sites; the proposed study will determine whether ENc stimulation can effectively reduce seizures in a chronic rodent epilepsy model. Further studies would be required to compare ENc anti-seizure efficacy with that of stimulating other sites experimentally. Comparative clinical trials will be necessary to determine

which sites for DBS are most effective in different epilepsy syndromes.

High frequency stimulation of the piriform cortex is known to be a good method for seizure induction (Sato et al., 1990) and the piriform cortex may be critically involved in epileptogenesis (McIntyre and Plant, 1989) initiating in the endopiriform nucleus for kindled seizures (Hoffman and Haberly, 1996). The importance of the endopiriform nucleus in ictogenesis may be from its strong anatomical connections within the temporal lobe (Behan and Haberly, 1999). In kindled animals, there is a loss of inhibitory interneuron functional diversity in the piriform cortex (Gavrilovici et al., 2012), which is indicative that this area may be affected by seizures, or may even be a seizure focus. Stimulation of the central piriform cortex with low frequency stimulation has been shown to suppress seizures in kindled animals (Zhu-Ge et al., 2007).

In this chapter, we will show results testing the effects of electrical stimulation of the ENc to modulate epileptiform activity (Figure 4). The ENc is a good anatomical target because it has widespread connections to the temporal lobe, is well-defined, and is in a relatively safe location for electrode implantation because it is far from major blood vessels.

### 2.1.2 Ventral Hippocampal Commissure

We will try to optimize electrical stimulation parameters applied to the ventral hippocampal commissure (VHC). Previous studies have shown that 1Hz VHC stimulation suppresses seizures (Rashid et al., 2012). The VHC is an axonal fiber tract that traverses both hemispheres in the rodent brain (Figure 5). The VHC innervates the

Ventral CA3 region of the Hippocampus.

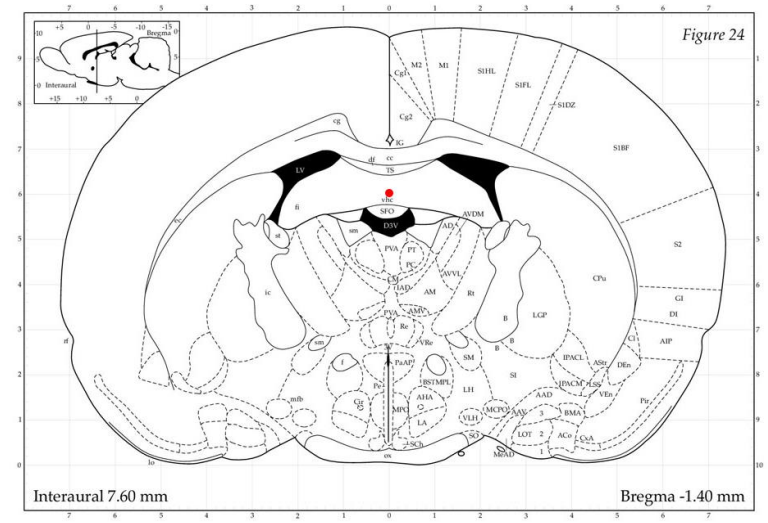


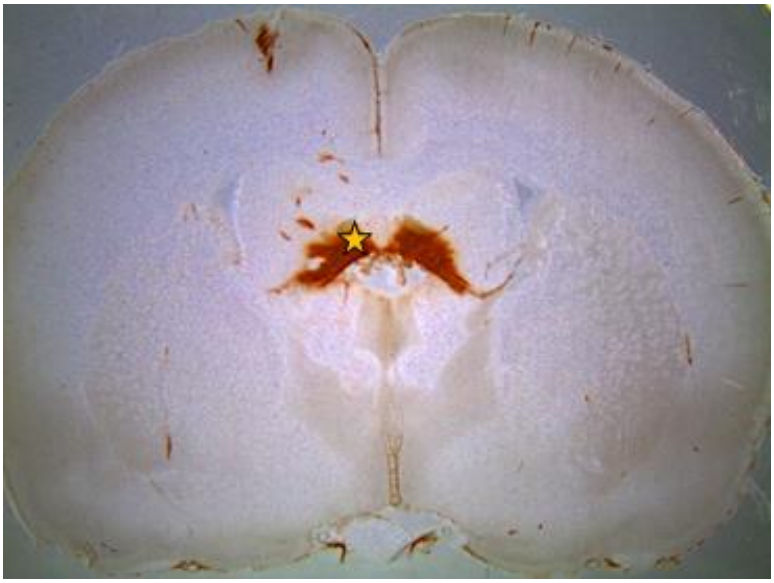

Figure 5. Ventral Hippocampal Commissure in Sprague Dawley rats.Top, rat brain atlas showing the VHC indicated by red dot. Bottom, coronal section showing placement of stimulation electrode (Microprobes, Inc.) into the ventral hippocampal commissure.

The VHC is not evolutionary conserved, therefore the structure does not exist in humans.

In a pilot trial Koubeissi et al. tested the effect of low frequency stimulation of the human Dorsal Hippocampal Commissure which is located at the apex of the posterior aspect of the fornix. Their results show a 92% reduction in seizures in a small cohort of patients (Koubeissi et al., 2013). Because of the strong effects of VHC stimulation in rodents, the fact the results translated well to humans, and the relatively low complexity and high success rate of targeting the VHC in rodent surgery, we decided to use the VHC stimulation as the experimental model to develop closed-loop controllers.

## *2.2 Methods*

### 2.2.1 Rodent Surgery

We have developed an acute model of ENc stimulation in our preliminary studies. Adult male Sprague-Dawley rats, weighing 350-400g, are anesthetized with inhalation isoflurane (4%) and then mounted in a stereotaxic frame. Then 2% lidocaine is injected into the scalp, the scalp retracted, and five burr holes are made for implantation of two recording electrodes, two stimulating electrodes, and a nanoinjector.

#### *2.2.1.1 ENc experiments*

Two EEG screw electrodes are placed symmetrically over each hemisphere. Electroencephalogram (EEG) activity from the cortex is recorded by differential recording between the two sites. Two stimulating electrodes are bilaterally implanted into the ENc (AP +.72, L -4.4, V -7.4) (shown in Figure 2), and two recording electrodes

(125$\mu$m in diameter) are inserted into the hippocampus CA1 area (AP -5.2, L -4.5, V -3.4).

#### *2.2.1.2 VHC experiments*

Adult male Sprague-Dawley rats, weighing 250-350g, are anesthetized with inhalation isoflurane (2%) and then mounted in a stereotaxic frame. Then 0.5% lidocaine is injected into the scalp, the scalp retracted, and five burr holes are made for implantation of one recording electrode, one stimulating electrode, and a nanoinjector, as shown in Figure 4. A single stimulating electrode is implanted into the Ventral Hippocampal Commissure (VHC, AP -1.4, L -1.2, V -2.5) (shown in Figure 5), and one recording electrode is inserted into the hippocampus CA1 area (AP -5.2, L -4.5, V -3.4).

### 2.2.2 Seizure Model

In all seizure experiments, we injected 0.5 µl 0f 4-aminopyridine (25 mM in artificial cerebrospinal fluid), in to the CA3 region of the hippocampus over 10-20 minutes. The full seizure state emerges following an intermediate step characterized by giant depolarizing events which are the LFP correlate of inter-ictal spikes recorded in the EEG. Our seizure model is extensively used in the lab of our collaborators (Yang and Rothman, 2001; Yang et al., 2003) and by other groups internationally. An example of a seizure recorded in CA1 is shown in Figure 6.

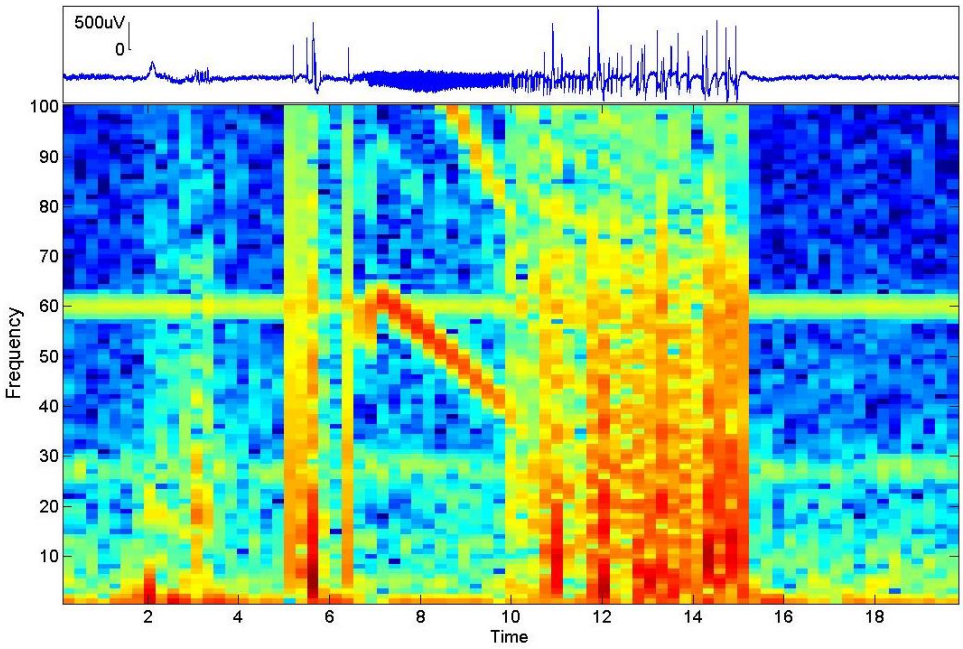


Figure 6. Seizure recorded in CA1 region of the hippocampus in an acute model of Epilepsy. Top: raw signal shows the seizure lasts for approximately ten seconds. Bottom: Spectrogram of seizure showing clear 'chirp' during the tonic phase.

The molecule 4-AP is a neuronal potassium channel blocker. Its effects on neural dynamics are two-fold: first, action potentials are elongated due to slower outward flux of potassium ions, and second there is an increase in the basal membrane potential of neurons, making them more excitable.

### 2.2.3 Data Acquisition

To record EEG, for detecting seizures to determine when to deliver stimuli, we use an open source real-time Linux-based platform, the Real-Time eXperimental Interface (RTXI) (http://rtxi.org) that is very flexible for developing closed-loop experiments. Using this system, seizure detection was done in real-time by threshold crossing of EEG

using an adaptive threshold. Once detected, the stimulus was applied at 130 Hz for 10-30 seconds, approximately the same duration as the tonic phase observed in unstimulated seizures.

### 2.2.4 Histology

Following euthanasia, the brains of the rats were extracted and placed into a 4% Formaldehyde solution then stored in a refrigerator for 3 days. Each brain was then blocked into sections containing the trajectory of the electrode, and the blocked sections were placed in a sucrose solution overnight. Each brain section was frozen on a cryotome and 60-micron coronal brain slices were cut from each section. All brain slices were mounted onto glass slides and imaged using bright-field light microscope.

## *2.3 Results*

### 2.3.1 Functional connection between the ENc and CA1 of the hippocampus.

During surgery, evoked potentials in the hippocampus are continuously measured while inserting the stimulating electrode into the ENc. A significant evoked potential in the hippocampus is seen in the ipsilateral hippocampus when the electrode is in the ENc, as shown in Figure 7; this helps verify that the target has been reached. Stimulation of on ENc also induces an evoked potential in the contra-lateral CA1; however, the response is delayed because the signal needs to travel across hemispheres via intra-hemispheric

axonal projections.

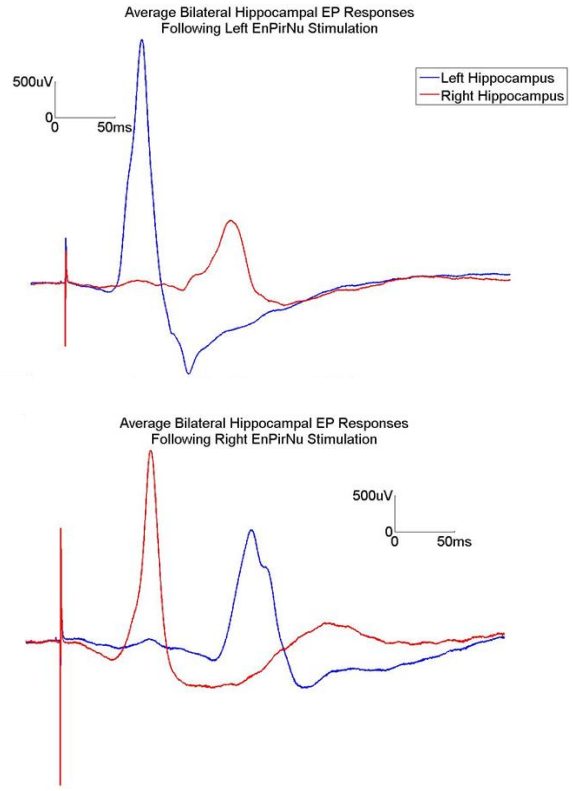


Figure 7. Unilateral stimulation results in bilateral evoked response in CA1 of the hippocampus.

For all experiments the amplitude of the electrical stimulation applied to the ENc was determined by the minimum stimulation amplitude that evoked the maximum evoked potential.

### 2.3.2 Open-loop stimulation of the ENc does not suppress seizures

We tested the efficacy of acute bilateral ENc stimulation to suppress seizures induced by 4-AP. In all experiments, we first measured the time epoch of two spontaneous recurrent seizures. This time was used to determine the duration of stimulation. In the first

experiment, we applied open-loop low frequency stimulation (1Hz) for the duration of two full seizures. We found that low frequency stimulation was unable to suppress seizures (Figure 8).

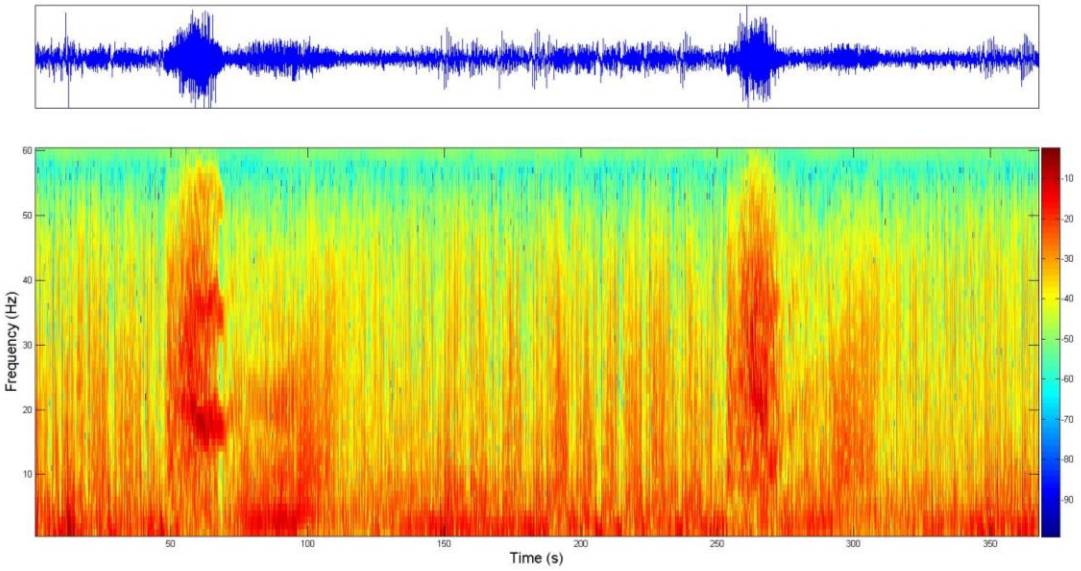


Figure 8. Bilateral low frequency stimulation of the ENu does not suppress seizures. Top, LFP time series recorded in dorsal CA1 of the hippocampus. Bottom, spectrogram showing strong broad band activity during seizures.

Since, low frequency stimulation was unable to suppress seizures, we decided to increase the amount of energy applied to the ENc. We tested the effects of acute high-frequency stimulation (130Hz) bilaterally. Our results indicate that open-loop high-frequency stimulation is unable to suppress seizures (Figure 9).

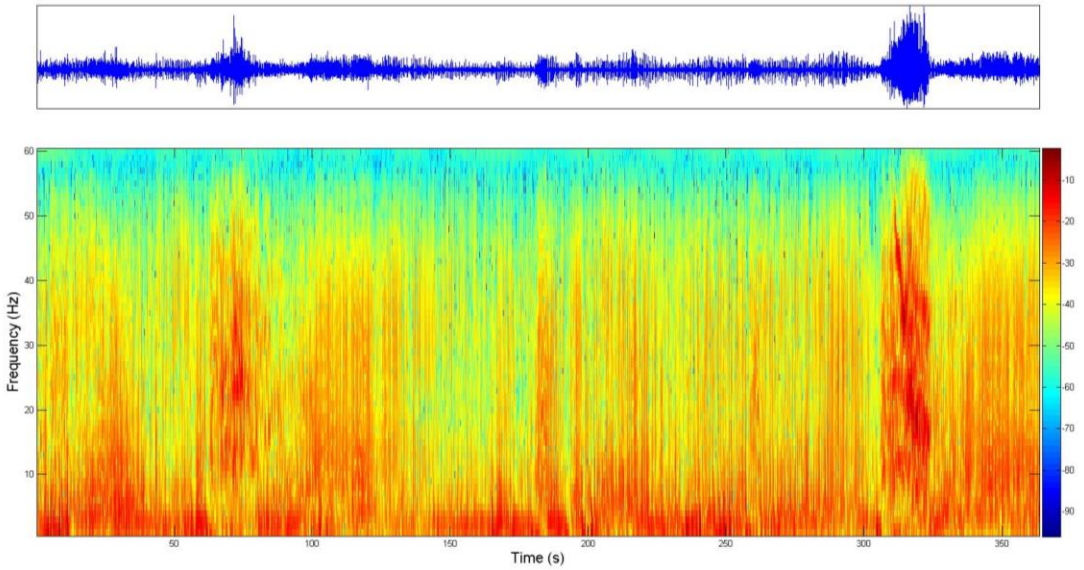


Figure 9. Bilateral high frequency stimulation does not suppress seizures. Top, LFP time series recorded in dorsal CA1 of the hippocampus. Bottom, spectrogram showing strong broad band activity during seizures.

It is clear from our acute experiments that neither low nor high frequency stimulation could suppress seizures. There are a few possible explanations for this. The ENc is not directly connected to CA1. It is possible that the information in the form of electrical stimulation transmitted via the Entorhinal Cortex is somehow dissipated before reaching the CA1 region. A second possible explanation for this is prolonged chronic stimulation of the ENc is required to effectively modulate the CA1 region. The seizure suppressive effects of ENc stimulation may not be evident in an acute preparation.

### 2.3.3 Effects of VHC stimulation on Hippocampal Dynamics

Simultaneous stimulation and recordings were done to determine proper placement of the stimulation electrode. Stimulation of the VHC results in a characteristic evoked response

in CA1 with approximately a 10ms delay (Figure 10). At low stimulation amplitudes, only a field excitatory post synaptic potential (fEPSP, or evoked potential) is seen. As the stimulation amplitude is increased, a population spike starts to emerge. A large population spike is observed at very large stimulation amplitudes. The minimum stimulation amplitude that results in an evoked response was used for all seizure suppression and closed-loop control experiments.

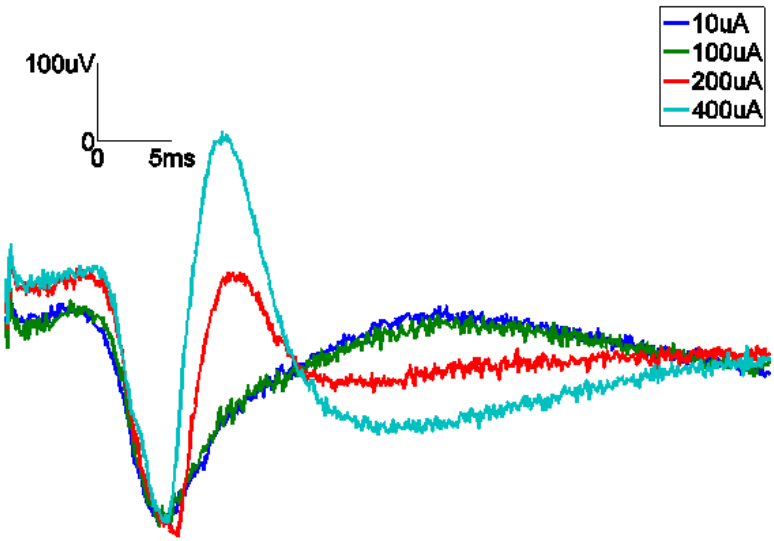


Figure 10. Evoked response to stimulation confirms stimulation electrode placement. Each color corresponds to a different stimulation amplitude.

Ramping the stimulation frequency from 1-120 Hz in a drug naive animal shows that stimulation at certain frequencies induces epileptiform activity while at others it suppresses activity. Low frequency stimulation below 1 Hz and high frequency stimulation above 80 Hz suppresses activity, while stimulation at 16-60Hz increases activity, as shown in Figure 11.

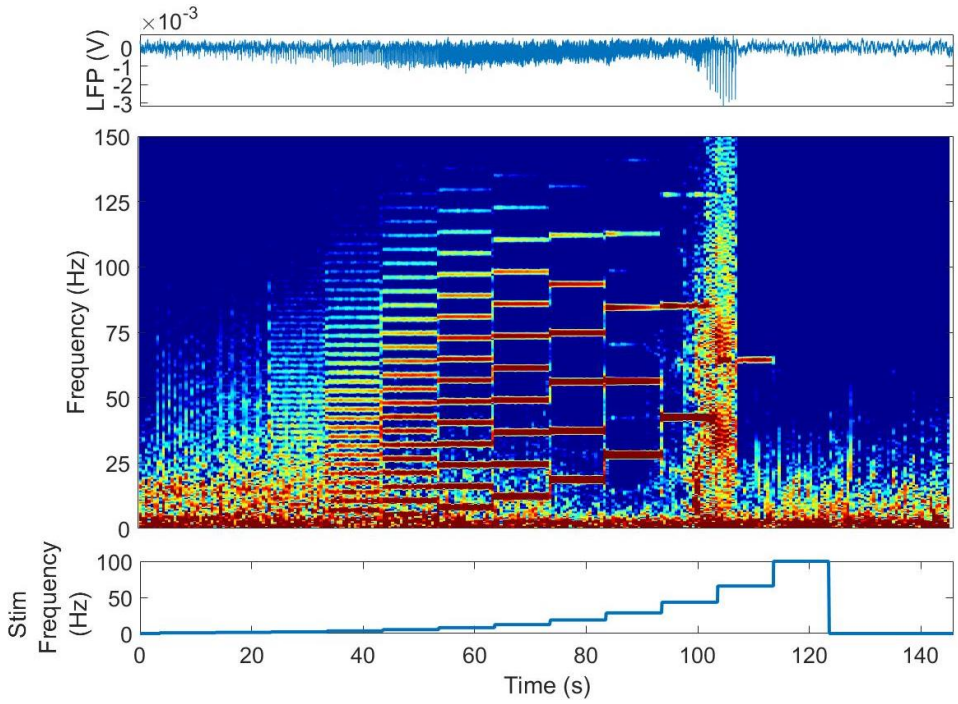


Figure 11. Hippocampal activity as a function of frequency of stimulation of the ventral hippocampal commissure. Voltage trace (top) and power spectrum (bottom) show neural activity in hippocampus during stimulation of the ventral hippocampal commissure. Stimulation artifacts are removed using an adaptive FIR filter to reveal neural activity. In spectrogram, time is the x-axis, frequency is the y-axis and color represents the power in the signal at each frequency band where reds and yellows represent high power and blues represent low power. Horizontal lines in the spectrogram occur due to periodic evoked responses occurring at the stimulation frequency and its harmonics. In drug naive animal, stimulation at low frequency decreases high frequency activity seen in the hippocampus, while 20-40Hz stimulation induces seizures, and is again suppressed by high frequency stimulation above 60Hz.

### 2.3.4 Low and high frequency stimulation of the Ventral Hippocampal Commissure suppresses seizures

We first tested to see if we could reproduce the results in Rashid et al (Rashid et al., 2012). The experiment (Figure 12) was set up as such: 1) baseline measure of LFP

activity, 2) stimulation, and 3) after stimulation. We used a metric called suppression ratio (SR) to assess the efficacy of low frequency stimulation (1Hz). The SR is a measure of broadband LFP power during stimulation compared with baseline. During the stimulation phase of the experiment the SR was approximately 54% leading to strong suppression of LFP activity. While a seizure still occurred during stimulation, we believe this was because the experiment was done in an acute animal model. Prolonged stimulation at 1Hz would most likely lead to stronger effects as seen in Rashid et al. (Rashid et al., 2012).

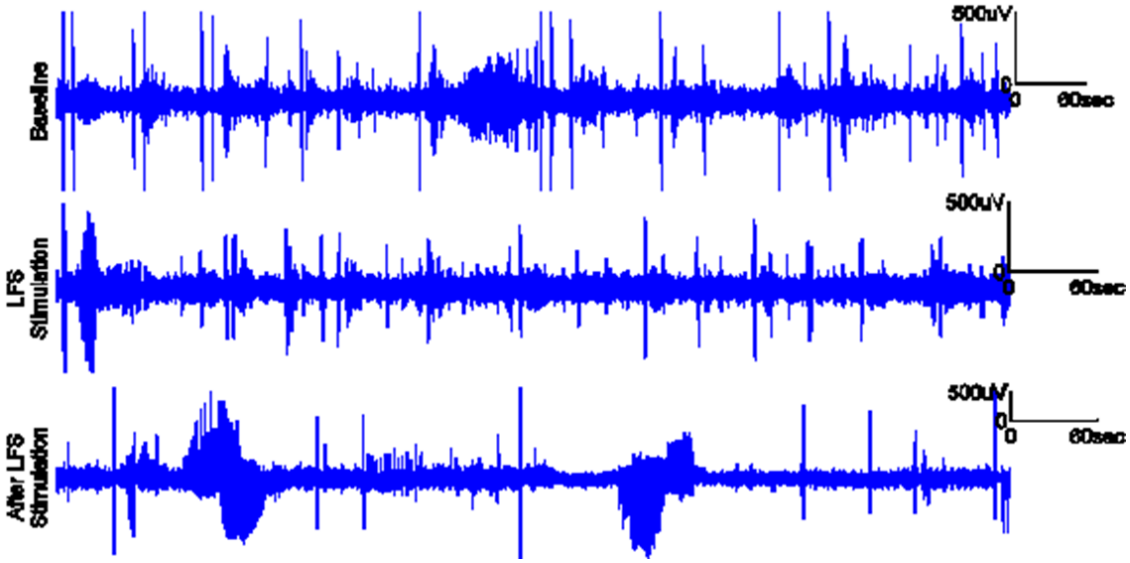


Figure 12. Low frequency stimulation of the VHC reduces suppression ratio compared to no stimulation. Top, baseline measurement of LFP in CA1 of the hippocampus. Middle, low frequency stimulation (LFS) suppresses LFP activity except for seizure at the onset of stimulation. Bottom, following LFS, seizure events begin to start again.

We also tested the efficacy of high frequency stimulation (HFS) of the VHC to suppress seizures. Similarly, to the previous experiment, we measured baseline LFP activity, then

applied HFS, and finally measured post-stim LFP. Before stimulation, it is clear that the seizure frequency is high (Figure 13). HFS leads to immediate broadband LFP suppression. When stimulation is turned off, seizure events start up again.

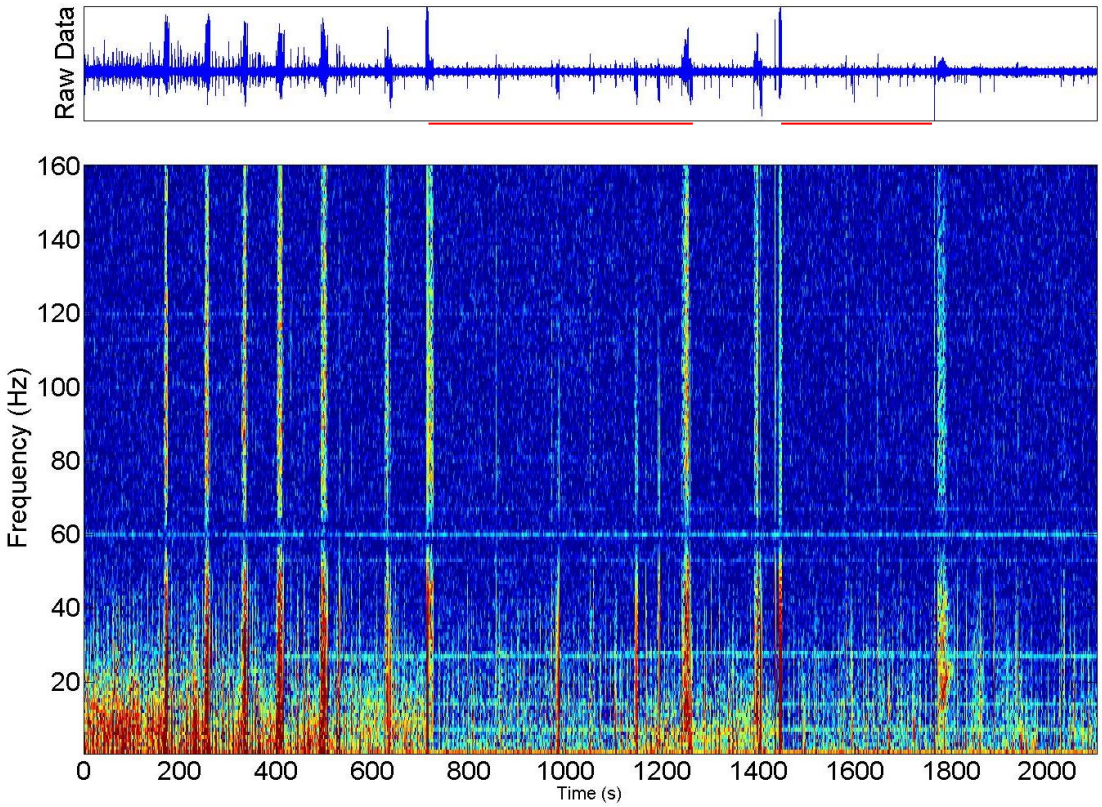


Figure 13. High frequency stimulation of the VHC suppresses seizures. Top, raw LFP trace in the experiment. Excursions from the baseline indicate seizure events. Bottom, spectrogram of seizure events clearly shows the broad band power distribution of seizures. During stimulation, all activity is suppressed to an extent. Following stimulation, seizure events return.

Our results indicate that both low frequency and high frequency stimulation of the VHC has strong effects on 4-AP induced seizures. We decided to use this platform to develop and test closed-loop algorithms which will be presented in chapters 3 and 5.

## *2.4 Discussion*

In this chapter, we tested the efficacy of stimulating a grey matter and white matter region for seizure suppression. Our results were clear, low and high frequency ENc stimulation, could not suppress seizures in an acute seizure model. We tested both unilateral and bilateral stimulation, but neither had any effect on seizure frequency. While there is a clear functional connection between the ENc and the hippocampus, applying acute stimulation was not enough to suppress seizure activity.

We shifted our efforts to using the VHC as a target area for seizure suppression. Our results reflect data produced from other labs: low frequency stimulation of the VHC can reduce seizure frequency. In preliminary experiments, we also tested the effects of high frequency VHC stimulation and found that it was highly effective at suppressing seizures. Other brain regions, both grey matter and white matter, were targeted for seizure control, but ultimately the VHC was the best option for moving forward with closed-loop experiments which will be discussed in the following chapters.

### 2.4.1 Chronic stimulation of the ENc

Bilateral LFS or HFS of the ENc did not suppress 4-AP induced seizures. We believe that the primary reason for this is because the ENc is a small nucleus that is not directly connected to CA1 of the hippocampus. Therefore, in order to have a strong and lasting effect on hippocampal dynamics, prolonged stimulation would be required. The acute preparation is not the appropriate experimental platform to test the effects of ENc stimulation for seizure control. Recently, preliminary data from Koubeissi et al., showed

strong seizure suppression following deep piriform cortex stimulation after two weeks. Their data clearly showed that the effects of deep piriform cortex stimulation did not manifest until after a few days.

### 2.4.2 VHC-Hippocampus as a robust experimental preparation for closed-loop control experiments

The VHC – Hippocampus system is a clearly delineated network. Stimulation of the VHC activates the ventral CA3 region of the hippocampus, and activity propagates throughout the hippocampus. Placing a recording electrode in dorsal CA1 separates the stimulation electrode from the recording electrode by approximately 4 millimeters. This way artifacts generated from the stimulation electrode are not as pronounced as it would be for direct hippocampal stimulation. Furthermore, both the VHC and the dorsal CA1 are well defined deep brain targets that are good stereotaxic targets. The surgical approach is simple and feasible which increases the chances of having successful experiments. This platform is ideal for developing and testing closed-loop algorithms because the anatomical regions do not vary much from animal to animal and because inputs to the hippocampus from the VHC are direct.

# Chapter 3 Real-time stimulus artifact removal for closed-loop electrical brain stimulation

*Excerpts, including text and figures from (Mendrela et al., 2016) have been included in this chapter.

## *3.1 Introduction*

Neuromodulation therapies have been developed to treat many different neurological diseases including Parkinson's disease (Krack et al., 1997), epilepsy (Fisher et al., 2010; Morrell and Group, 2011), stroke (Edwardson et al., 2013). Recently, there is increased interest in developing adaptive therapies that optimize stimulation parameters to maximize the therapeutic outcome (Nagaraj et al., 2016). New implantable hardware technologies can simultaneously record while delivering electrical stimulation. Closed-loop adaptive therapies are under development that measure a biomarker of the disease and adjust stimulation parameters to minimize a biomarker automatically (Little et al., 2013).

Closed-loop neuromodulation requires continuous feature/biomarker detection for precise therapy administration. With electrical and magnetic therapies, the stimulation pulses interfere with the recording amplifiers. It is difficult to simultaneously assess biomarkers while delivering high frequency therapies, therefore it is necessary to remove stimulation artifacts. While there are many stimulus artifact removal methods there is little research directly comparing the efficacy and efficiency of each filter. In this chapter, we will review and compare methods of stimulus artifact removal, with particular emphasis on methods that can be implemented in real-time on a low powered implantable device.

Stimulation artifacts recorded on implanted depth electrodes, electrocorticogram (ECoG), or electroencephalogram (EEG) arrays occur following transient stimulation

using current, voltage, or magnetic waveforms within a conductive medium. The industry standard is to blank the stimulus artifact by holding the electrode potential to zero values during a stimulation epoch (Frei et al., 2007). If the stimulus amplitude or electrode impedance is very high, recovery from the artifact can be much longer than the stimulus itself. Blanking methods can prevent the amplifiers from railing allowing the signal to be recovered faster. As long as the stimulus artifact duration does not change, this method is very effective, but it comes at a cost of losing all information during the stimulation window.

If the stimulus artifact does not saturate the amplifiers, the stimulus artifact waveform can be subtracted from the signal to recover the underlying neuronal signal. Stimulus artifacts are patient or subject specific, due to the subtleties of the electrode geometry and of the orientation of the stimulus electrode to the recording electrode, therefore the stimulus artifact must be fit to each recording. If the stimulus artifact is stereotypic within an experiment, a template removal method, where an average stimulation waveform is estimated and subtracted during each stimulation epoch (Hashimoto et al., 2002), is very effective, as long as the stimulation waveform does not change.

Comb filters subtract the signal recorded from the last stimulus input from the current stimulus response. In a way, this is an adaptive filter that only remembers the last response. Adaptive Finite Impulse Response (FIR) filters fit the filter coefficients stimulation artifacts on-line to recover underlying neuronal signals (Mendrela et al.,

2016). These adaptive filters fit filter coefficients using a least-squares algorithm to minimize the prediction error of the recorded signal. Adaptive filters require no *a priori* knowledge about the stimulus artifact waveform and are more robust to changes in stimulus artifacts during the experiments. Adaptive filters are computationally only a little more intensive than a standard FIR filter and are efficient enough to be implemented in an implantable device.

A Kalman filter can be used to remove stimulus artifacts. It requires developing a computational model of the stimulus artifact and the neuronal signal. It is more computationally intensive, but has the further advantage that it can be used to remove measurement noise from the measured signal and estimate the true state of the LFP.

In this chapter, we will compare the efficacy of time-invariant and time-varying filters in removing stimulation artifacts without introducing new artifacts. We test five different filters for real-time stimulus artifact removal: blanking, comb, template, adaptive FIR, and Kalman. The efficacy of each stimulation artifact removal algorithm will be tested on a computational model, where the underlying signal and stimulus artifacts are known. To simulate neuronal activity, we used a stochastic oscillator which produces a time series that qualitatively resembles local field potential activity. Stimulation artifacts were generated by band pass filtering a rectangular biphasic waveform while non-stationarities were simulated by changing the stimulus artifact twice during the simulation. Finally, we demonstrate the filters on a sample recording from *in vivo* rat hippocampal recording while applying electrical stimulation, and we will show

results using a bidirectional neural interface hardware that removes stimulation artifacts in real-time.

## *3.2 Methods: Digital Implementation*

A stimulation artifact removal filter estimates the underlying neuronal signal, preserving it, while removing the stimulation artifact. Accurate estimation, however, is needed to enable a closed-loop control scheme to optimize stimulation waveform removal. Figure 14 shows a canonical closed-loop control scheme for DBS with stimulus artifact removal.

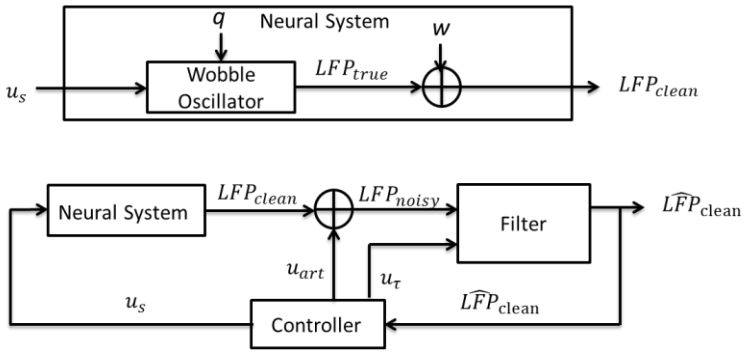


Figure 14. Block diagram of closed-loop Deep Brain Stimulation system with stimulus artifact removal. Top, the neural system consists of the Wobble Oscillator with process noise $q$ resulting in the true LFP output ($LFP_{true}$). The addition of $\frac{1}{f}$ measurement noise, $w$, is the output of the neural system ($LFP_{clean}$). Stimulation command $u_s$ leads to a delayed evoked response in the neural system. Bottom, neural system in a closed-loop control scheme. When a stimulation pulse is applied ($u_{art}$), the stimulation artifact waveform is added to the neural signal, and $u_s$ elicits a delayed evoked response in the Neural System. The timing of each stimulation $u_\tau$ is sent to the Filter. Measured signals from the dynamic system are filtered using a stimulation artifact removal algorithm. The filtered signals can be sent to a controller to determine the appropriate input to the system.

The stimulus results in a modulation of the neural response, as well as an artifact in the measured signal. The stimulus artifact is caused by the electrode response and the

hardware filters. We modeled the stimulation artifact as additive and non-stationary in amplitude for the simulation. Our motivation for using non-stationary artifacts was to simulate changes in artifacts during an experiment that could be caused by shifting of the electrodes, accumulation of the glial scar tissue around the electrode, or damage to the electrode.

In general, we estimate the underlying neural system response to stimulus input from the $LFP_{noisy}$ signal by subtracting a filtered version of the stimulus input $u(t)$:

$$LFP_{clean}(t) = LFP_{noisy}(t) - Filter(u(t)) \quad (3.1)$$

The estimation error is the difference between the actual LFP response and the estimated response:

$$LFP_{error}(t) = LFP_{response}(t) - LFP_{clean}(t) \quad (3.2)$$

### 3.2.1 Stimulation artifact removal methods

Blanking with linear interpolation

Blanking is a filtering technique that simply removes the data, for a pre-determined time window, during each stimulation pulse. The technique is often implemented in hardware where the filter prevents the stimulation artifact from saturating the amplifier by grounding the recording electrode and digitally holding the signal until the amplifier has recovered. If the stimulation artifacts are being removed offline, then the signal can be linearly interpolated between the voltage measured at the onset and offset of the blanking window. The formula for this technique is as follows:

$$LFP_{clean}(t) = \begin{cases} LFP(t_{Prev\ Stim\ Time}), if\ (t - Curr\ Stim\ Time < Blanking\ widow) \\ LFP(t), otherwise \end{cases} \quad (3.3)$$

The advantage of this technique is that the amplifier gain can be set to match the neural signal and blanking the stimulus artifact when the amplifier would be saturated. Additionally, this technique is both computationally and memory efficient. A drawback with this technique is that the neural signal is lost during the blanking window.

Comb filter

The comb filter is a time-varying filter that subtracts the previous stimulation artifact from the current stimulation artifact. The formula for this technique, for a recording with constant ISI is as follows:

$$LFP_{clean}(t) = LFP_{Noisy}(t) - LFP_{Noisy}(t - \tau) \quad (3.4)$$

This technique works best if the inter-stimulation interval (ISI) is fixed. The comb filter effectively cancels common sources between stimulation artifacts while also adapting rapidly, delayed by one stimulation pulse, to changes in the stimulation artifact shape. This filter is computationally efficient. However, the algorithm may require a significant amount of memory may be required if the ISI is long. Another problem is that subtraction of the delayed signal from the current sample corrupts the data between stimulation pulses. If the ISI is not fixed this filter can be adapted to use the response from the last stimulus pulse, decayed over time with time constant $\tau$:

$$LFP_{clean}(t) = LFP_{noisy}(t) - LFP_{noisy}(t - t_{Curr\ Stim\ Time} + t_{Prev\ Stim\ Time}) * \tau \quad (3.5)$$

This formulation makes the memory requirement dependent only on the length of

the artifact, not the time between artifacts, but at the cost of additional operations. Regardless of the formulation of this filter it has difficulties with overlapping stimulus artifacts. In high-frequency stimulation paradigms, this may be impossible to avoid and will result in decreased accuracy.

Template filter

Unlike the comb filter, which uses the previous stimulation artifact, the template filter subtracts an averaged stimulation artifact:

$$Template(k) = \frac{1}{N}\sum_{i=0}^{N} LFP_{noisy}(stim(i) + k) \quad (3.6)$$

where $N$ is the length of the stimulation artifact template. At each stimulus onset, this Template then subtracted at each stimulation time.

$$LFP_{clean} = \sum_{i=0}^{N} LFP_{noisy}(stim(i) : k) - Template(k) \quad (3.7)$$

Given the estimated stimulation artifact is averaged over many stimulation pulses, this filter will result in an accurate estimate of the artifact. If the stimulus artifact does not change, the template removal is very good at recovering the underlying signal. However, if the stimulation artifact changes then the template will introduce errors.

Optimal Filters

Optimal filters seek to minimize the difference between the filtered signal and the true signal. Optimal filter coefficients can be calculated, analytically, by solving a set of linear equations known as the Weiner-Hopf equations (Hayes, 2009). However, for time-invariant systems a recursive coefficient estimation method can be implemented. We will discuss and test two optimal filters, the Adaptive FIR filter and the Linear Kalman Filter.

*Adaptive FIR filter*

The adaptive FIR filter uses a Least-Mean Squares algorithm to optimize filter coefficients to minimize prediction errors (Haykin and Widrow, 2003). The FIR filters the history stimulus waveform using a set of coefficients $w$

$$MA_{model}(t) = \sum_{i=0}^{k} w(i) * u(t-i) \quad (3.8)$$

$$LFP_{clean}(t) = LFP_{Noisy(t)} - MA_{model}(t) \quad (3.9)$$

The Least-Mean Squares algorithm updates the coefficients to minimize the cleaned signal:

$$w(i) = w(i) + \mu * LFP_{clean}(t) * u(t-i) \quad (3.10)$$

Where $\mu$ determines the learning rate of the filter. The advantage of the adaptive FIR filter is that it requires no *apriori* knowledge about the stimulus artifact, except the approximate duration of the stimulus artifact to determine the order of the filter $k$. In fact, if the input waveform is just 1's and 0's, to indicate onset of stimulus times, this method is just a running average of the stimulus artifacts. This method generalizes well to any arbitrary stimulus input, provided the response is linear. The smaller $\mu$ is the longer the average and the slower the filter will adapt. Selecting $\mu$ is a bit tricky because it depends on the size of the signals being measured. If $\mu$ is too large, then the filter can become unstable. The higher the model order, the more computational time and memory is needed. However, this is a relatively efficient filter. If the input consists of only 0's and 1's, then $LFP_{clean}(t) = LFP_{Noisy}(t)$ for any signal beyond the order of the filter after

each stimulus. Like the template, this method will not distort the signal between stimulus pulses. However, it is necessary to make the order high enough to assure that the stimulus response has returned to zero or there will be a discontinuity in the filtered signal at k points after each stimulus pulse.

*Kalman Filter*

The three previous techniques simply subtract a stimulation artifact estimation to remove the stimulation artifact. The Kalman Filter finds an optimal weighting for raw data and model predictions to make the best estimate of the underlying neuronal signal. This method requires a priori knowledge of the stimulation artifact and LFP dynamics, therefore two different models need to be determined. Assuming that the response is linear and time-invariant (LTI) we can use the Ho-Kalman subspace system identification algorithm (Miller and de Callafon, 2012) to estimate a model of both the stimulation artifact and the LFP dynamics. The Ho-Kalman algorithm is detailed in Appendix C.

The Kalman Filter assimilates model predictions with observed data. The first step in designing the Kalman Filter is to establish the dynamical systems model:

$$x(t+1) = Ax(t) + Bu(t) + w(t) \tag{3.11}$$

$$y(t) = Cx(t) + Du(t) + v(t)$$

Where, $A$, $B$, $C$ are block-matrices that contain the coefficients for both the evoked response and stimulation artifact models.

$$A = \begin{bmatrix} A_{ER} & 0 \\ 0 & A_{Art} \end{bmatrix}, B = \begin{bmatrix} B_{ER} \\ B_{Art} \end{bmatrix}, \text{C} = [C_{ER} \quad C_{Art}], \text{ and } D \sim 0 \tag{3.12}$$

Where $A_{ER} \in R^{NxN}$, $A_{Art} \in R^{MxM}$, $B \in R^{NMx1}$ , and $C \in R^{1xNM}$

To estimate the LFP during the stimulation artifact we used the dot product of only the coefficients pertaining to the evoked response model:

$$LFP_{clean} = C(1:N) * x_t(1:N) + D * u_t \quad (3.13)$$

where $N$ is the order of the evoked response model. Because we are interested in only the prediction of the LFP and not the stimulation artifact, the output for the state space system reflects only the coefficients and states pertaining to the evoked response model.

The complete Linear Kalman Filter is as follows:

The first step is to predict the output $(y(t))$ of the model

$$y(t) = Cx'(t) + Du(t) + v(t) \quad (3.14)$$

where $x'(t)$ is the *apriori* estimate of the state of the system.

The Kalman gain $(K)$ is calculated using the *apriori* estimate of the system state covariance matrix $(P')$ and the measurement noise $(R)$.

$$K(t) = P'(t)C(CP'(t)C^T + R)^{-1} \quad (3.15)$$

During the state update step, the Kalman gain determines how much emphasis needs to be placed on the difference between the predicted model output and the raw data.

$$x(t) = x'(t) + K(t)(y(t) - LFP_{noisy}) \quad (3.16)$$

Similarly, the state covariance matrix is updated as follows:

$$P(t) = (I - KC)P'(t) \quad (3.17)$$

Lastly, time update equations for the states and covariance matrix are:

$$P'(t+1) = AP(t)A + V \quad (3.18)$$

$$x'(t+1) = Ax(t) + Bu(t)$$

### 3.2.2 Computational model and simulations

We used a stochastic damped oscillator model to produce time series data with characteristics similar to neural local field potential activity. Stimulation artifacts are modeled as band-pass filtered rectangular biphasic waveform with cutoff frequencies of 1 Hz and 500 Hz. The response to stimulation is modeled as a $\alpha$ function, and was injected into the input variable of the model. Sustained $\frac{1}{f}$(pink) noise is injected into the model, for all simulations, to model measurement noise.

The formulation of the computational model is as follows:

$$y(n) = x(n) - (ax(n-1) - 2ay(n-1)\cos\big(w(n)\big) - a^2 y(n-2) + \zeta(n) \tag{3.19}$$

where $y(n)$ is the simulated time series data, $x(n)$ is the input response to stimulation, $a$ is the damping coefficient, $w$ is the mean reverting angular frequency of the oscillator determined by an Ornstein-Uhlenbeck process, and $\zeta$ is the measurement noise. We tested the efficacy of each algorithm to remove stimulation artifacts from periodic and random stimulation. The large range of ISIs that are possible with random stimulation tests whether the algorithm can remove artifacts without removing features of the system response. To determine the efficacy of each algorithm we compared the filtered signal with the true neural signal (without measurement noise $\zeta$) and the raw neural signal (with measurement noise).

All simulations and analysis were completed in Matlab 2016a. The time step of the model used was 0.1 ms. Power spectral density (PSD) was estimated using the Welch

method (*pwelch* function in Matlab) and spectrogram was estimated using a Hamming window of 512 samples with an overlap of 256 samples (spectrogram function in Matlab).

### 3.2.3 Experimental methods

We also tested the stimulation artifact removal algorithms on experimental data recorded from anesthetized rats. Experimental protocols are described in chapter 2 section 2.2.1.2. Male Sprague-Dawley rats weighing 225-300g were used for all experiments. Animals were anesthetized using 1.5-2% Isoflurane. A bipolar stimulation electrode from FHC Inc (CBASC75). was stereotaxically placed into the Ventral Hippocampal Commissure.

LFP was amplified using a Grass P15 amplifier. Data was digitized using a PCIe 6950 National Instruments digital acquisition card at 5k Hz. Data was recorded use an open source real-time Linux-based platform, the Real-Time eXperimental Interface (RTXI) (http://rtxi.org) that is very flexible for developing closed-loop experiments (Netoff et al., 2005; Miranda-Domínguez et al., 2010; Stigen et al., 2011). It consists of a computer running real-time Linux interfaced to the recording and stimulation amplifiers. Biphasic rectangular waveforms were generated using a custom built module in RTXI and output through the DAQ card to an A-M Systems Model 2200 Stimulus Isolation Unit. Stimulation amplitude for all experiments was at 500uA.

## *3.3 Results: Digital Implementation*

### 3.3.1 Real time stimulation artifact removal in a computational model

Here we test the efficiency and accuracy of five different algorithms for digitally removing stimulation artifacts. We first test artifact removal in a simulated LFP recording with known neuronal response and stimulation artifact and the error between the estimated neural signal and the actual signal can be measured. To simulate the LFP we use a stochastic process to generate signals we call a "Wobble Oscillator" to which we added a dynamic input, to simulated the neuronal response, a stimulus artifact and measurement noise. The stimulus artifact model was changed midway through the modeled experiment to simulate non-stationarities of the recording system. We applied periodic stimulation and a Poisson train. We test stimulation artifact removal by: blanking, comb filter, template, adaptive FIR, and Kalman filter. The inter-stimulus interval $ISI$ in Poisson stimulation is random which results in some pulses landing during the evoked response of the system. We used Poisson stimulation to ascertain the efficacy of each algorithm in removing the algorithm while minimizing loss of neural information from the evoked response. Figure 15 shows examples of each stimulation artifact removal algorithm following Poisson stimulation in the computational model.

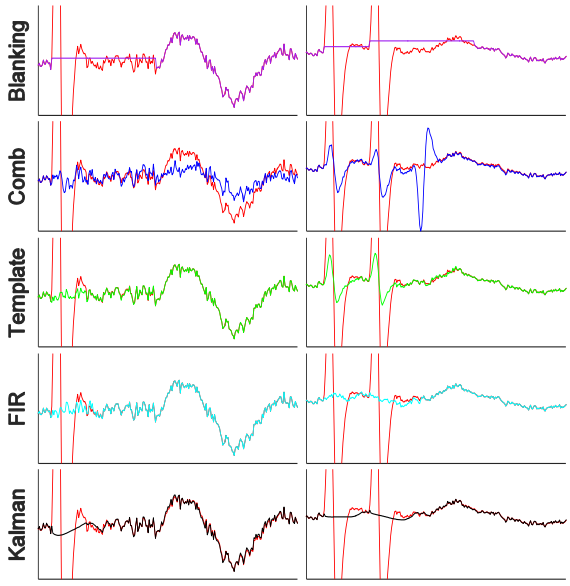


Figure 15. Example traces of stimulation artifact removal filters during two situations: non-interfering stimulation artifact (Left) and interfering stimulation artifact (Right). The high frequency deviations in the neural signal are a result of the added measurement noise $\zeta$. The noisy data with stimulation artifacts are indicated by the red traces. Filtered data are color coated by the type of filter that was used.

Non-interfering stimulation artifacts are removed with all filters while interfering stimulation artifacts best removed with the Adaptive FIR and Kalman filters. Presence of the stimulation artifacts still occur with the Template filter while the Comb filter sees the propagation of error in neural signal recovery.

#### *3.3.1.1 Spectral Analysis*

We calculated power spectral density (PSD) estimates using Matlab's *pwelch* to get the average spectral response following stimulation artifact removal. The sharp peak at approximately 100 Hz in Figure 16 corresponds to the evoked responses. The effect of the stimulation artifact is clear in the unfiltered PSD. There is a strong departure from the

clean PSD between 100 Hz and 1700 Hz. All the filters reduce the spectral artifact; however, only the template filtering method recovers the clean PSD completely for both periodic and Poisson stimulation.

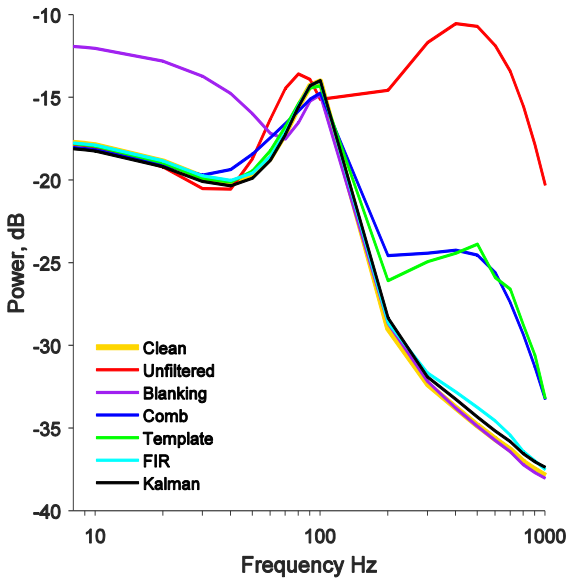


Figure 16. Power spectral density estimates for Poisson stimulation. The clean signal corresponds to the output of the computational model without measurement noise $\zeta$. The Unfiltered signal contains both measurement noise and the stimulation artifact. Outputs for each filter are indicated by the different color PSD estimates.

The template, FIR, and Kalman filters reduced the power between 200 and 1000Hz towards the clean signal the best when compared with the clean signal PSD (Figure 16). Both the comb and template filters reduces the power within that frequency range, the bump within the frequency range is still present indicating that part of the stimulation artifact is still present. One item of note is that while the blanking filter successfully removes the stimulation artifact, it also introduces strong low frequency power. This is most likely due to the sample-and-hold policy over the duration of the stimulation artifact. When considering the appropriate stimulation, artifact removing filter for a given

application it is imperative to makes sure the filter does not produce artifacts while removing the stimulation artifact.

#### *3.3.1.2 Performance Measures*

To measure how well each filter removed the stimulation artifact while preserving the neural signal, we calculated the root mean squared error (RMSE). The RMSE calculated was the residual between the true signal ($LFP_{true}$) and the filtered signal ($\widehat{LFP}_{clean}$). Our motivation for calculating RMSE between the true signal and filtered signal was because the purpose of Kalman Filter is to extract the true state of the system. Optimal state estimation via Kalman Filtering is an integral component in Optimal Control algorithms (Friedland, 2012; Kirk, 2012), therefore the ability to do stimulation artifact removal and state estimation simultaneously is imperative for some control applications.

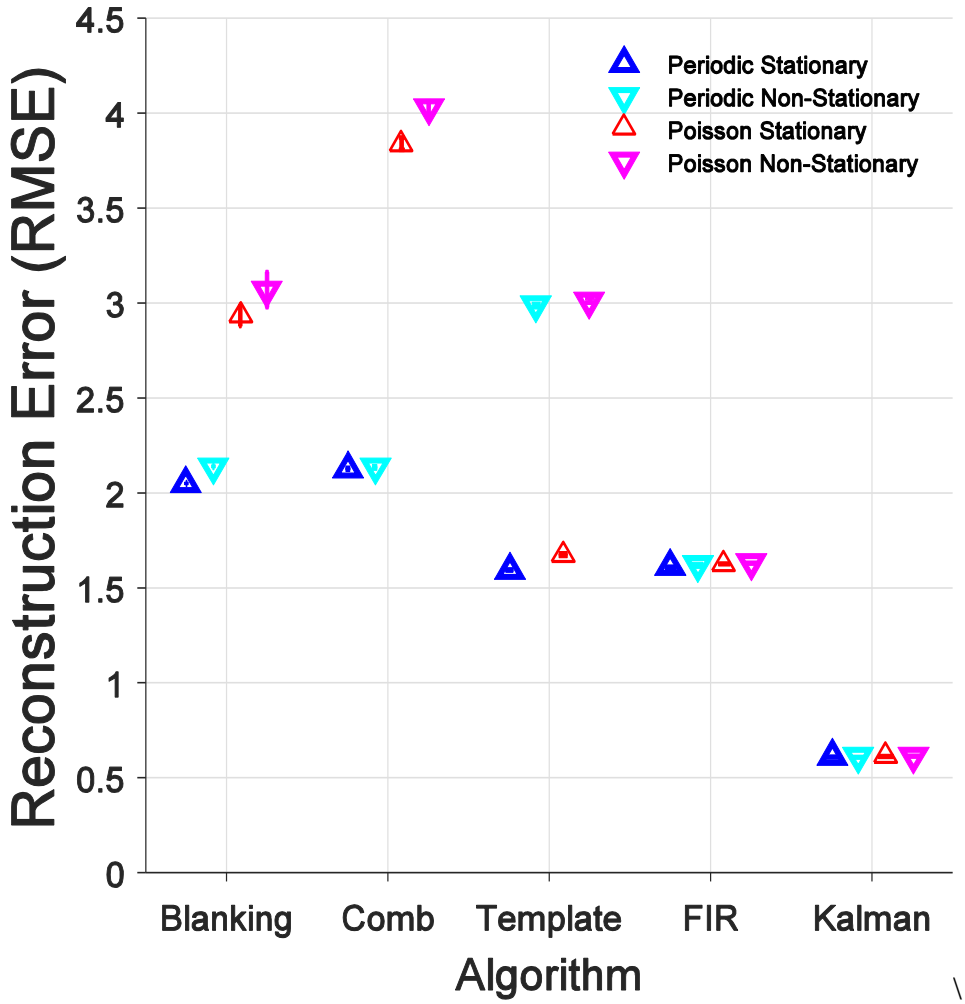


Figure 17. Mean RMSE across 10 simulations of the Wobble Oscillator and 95% confidence intervals. Stimulation paradigms are identified by the color and orientation of the triangle data points. RMSE value between the true signal ($LFP_{true}$) and the filtered signal ($\widehat{LFP}_{clean}$) across five different filters.

Our results indicate that the Kalman Filter outperforms all other filters when comparing the true LFP with the filtered LFP. This makes sense because the Kalman Filter is the only method designed to recover the true state. The Kalman Filter is robust to the type of stimulation paradigm and non-stationarities in the stimulation artifact amplitude, and therefore confers a low RMSE value for the true signal comparison. The

adaptive FIR filter comes in second place in performance across stimulation methods, and has a much lower computational cost than the Kalman Filter. However; the adaptive FIR filter is not designed to extract state information, and therefore cannot be used for Optimal Control algorithms.

An important consideration of algorithm performance with regards to hardware implementation is the computational cost of the calculations. We measured the performance time for each algorithm using an Intel i7 processor with 12GB RAM on a Windows 7 OS (Figure 18). The Kalman Filter was the most sophisticated algorithm tested, and its complexity is reflected in the highest computational cost. The template and blanking filters have the lowest cost because the methods use only simple calculations. Additionally, the computational cost of the adaptive FIR maybe low enough to be implemented in DBS device technologies. A summary of the performance of each algorithm and recommendations are provided in Table 1. In order to simplify our recommendations we used a scale (Low, Medium, and High) to indicate performance and complexity of each filter.

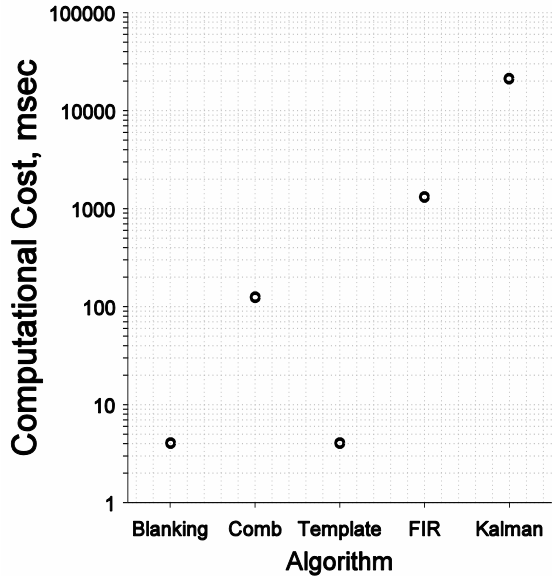


Figure 18. Computational cost for each algorithm averaged over 50 simulations including both periodic and Poisson time stimulation. Comb is the second fastest filter, but has the largest RMSE of all the filters. Template algorithm performs the quickest, but a template needs to be determined prior to running the filter. FIR takes the approximately the same amount of time to run, however, both have lower RMSE values than a comb filter and, occasionally, the template filter along with not requiring a template of the stimulus artifact prior to running the filter.. Kalman takes the longest to run but has a similar RMSE value as the FIR and IIR filters.

Table 1. Summary of Stimulation Artifact Filter Performance

| | **Blanking** | **Comb Filter** | **Template** | **FIR** | **Kalman** |
|---|---|---|---|---|---|
| **RMSE (Stationary)** | High | High | Medium | Medium | Low |
| **RMSE (Non-stationary)** | High | High | High | Medium | Low |
| **Computational Cost** | Medium | Medium | Low | Medium | High |

### 3.3.2 Offline implementation using In vivo data

We tested the efficacy of each filter *in vivo*. Biphaic pulses were applied in periodic fashion to the VHC and fEPSP's were measured in CA1 of the Hippocampus. Our results show that the Kalman Filter provides the greatest reduction in signal corruption from the stimulation artifact (Figure 19, right). The effects of the filter is more clearly seen in the spectrogram (Figure 19, left). The raw signal is contaminated with strong bands corresponding to the 100Hz stimulation and harmonics. All filters can reduce the stimulation artifact to an extent; however, the blanking, comb, and template filters are unable to remove the entire artifact as indicated by the light band at 100Hz. The adaptive FIR performs poorly at first due to the learning process. Eventually, the FIR filter learns the model for the stimulation artifact and removes majority of the artifact.

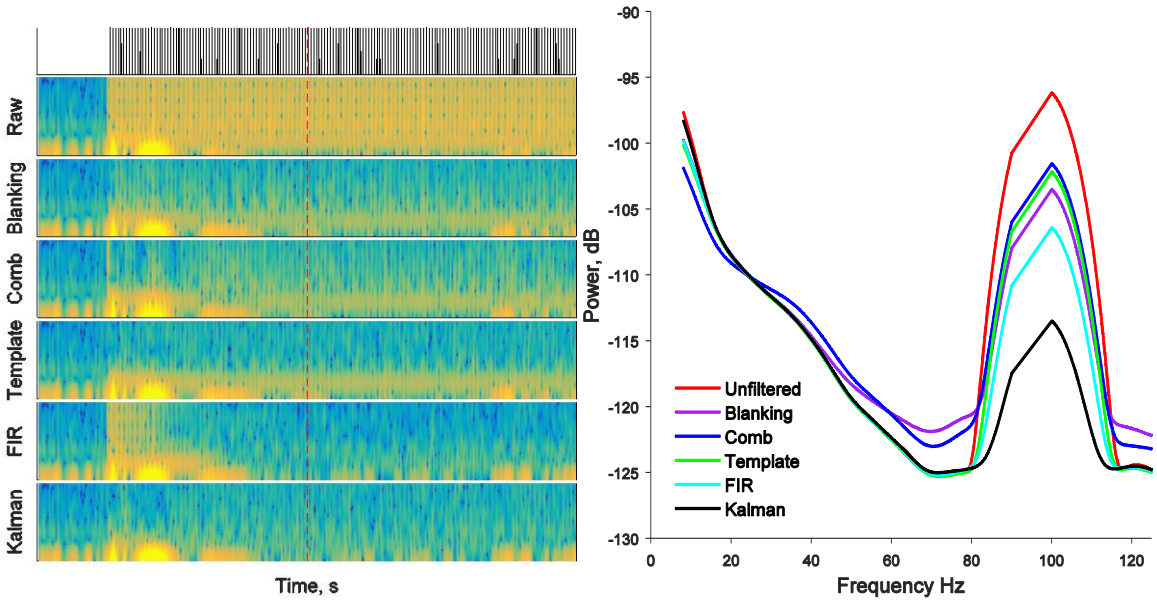


Figure 19. Efficacy of different stimulation artifact filters following periodic stimulation *in vivo*. Left, spectrogram showing the evolution of the frequency content in the signal as a function of time. Lighter colors indicate higher power in the frequency band. Right,

PSD using data following the red dotted line in the spectrogram. The outputs for each filter are color coded.

## *3.4 Methods: Analog Implementation*

### 3.4.1 Bidirectional neural interface for real-time stimulation artifact removal

Stimulus artifacts can be removed in real-time after digitizing the data, but collaborators of ours also developed a stimulus artifact filter that filters the data prior to digitization. In this section, we will also show the results from our collaboration with an engineering group from the University of Michigan testing the efficacy of a novel neural interface in removing stimulation artifacts *in vivo*. The device implemented the Adaptive FIR filter using a field programmable gate array (FPGA) that was programmed using LabView. Filter parameters and coefficient updates were executed digitally whereas the real-time signal subtraction step was performed prior to signal digitization. This method is superior to pure digital implementation because it can remove stimulation artifacts that can saturate amplifiers during the digitization step. The full details of the hardware implementation can be found in (Mendrela et al., 2016).

### 3.4.2 Experimental methods

Animal experimental protocols are described in chapter 2 section 2.2.1.2 and in Chapter 3 section 3.2.3. We used LabVIEW software for all data acquisition, stimulation, and FPGA device programming. Biphasic waveforms were used for all

stimulation experiments. In the naïve animal, we used 12Hz stimulation, and in the seizure model we used 120Hz stimulation.

## *3.5 Results: Analog Implementation*

In a separate study, we used a bidirectional neural interface with a field programmable gate array (FPGA) to implement the Adaptive FIR filter for real-time stimulation artifact removal *in vivo*. Stimulation was applied to the VHC at 12Hz. Our results show that an Adaptive FIR filter implemented in hardware can successfully remove stimulation artifacts prior to digitization (Figure 20). The stimulation artifact magnitude is highest at the beginning of stimulation but eventually decreases as the model of the stimulation artifact begins to capture the real artifact waveform. Total learning time in this example was approximately 15s. This is because the learning rate was set fairly low, increasing the learning rate would decrease convergence time.

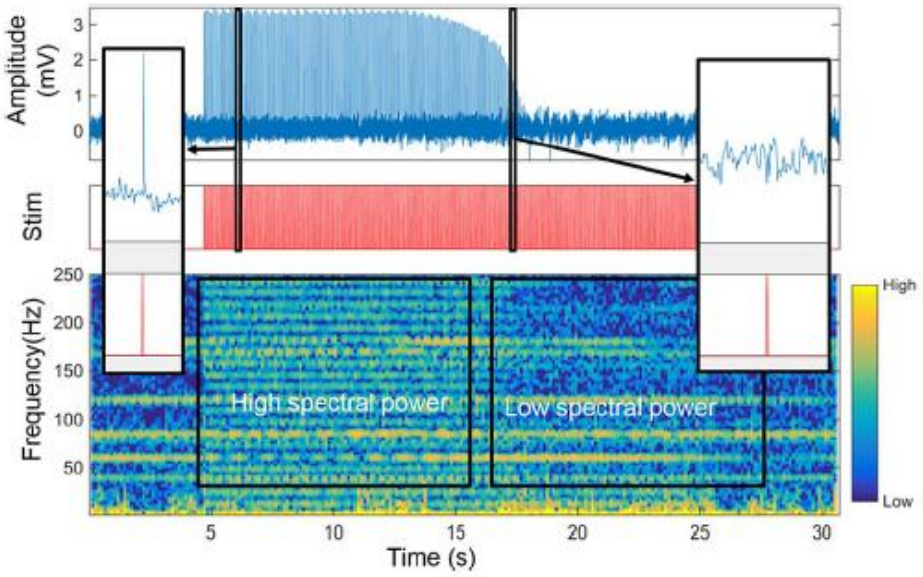


Figure 20. Real-time stimulation artifact removal using a bidirectional neural interface. Top, LFP hippocampal recording showing high amplitude stimulation artifacts that subside as the Adaptive FIR learns the stimulation artifact model. Insets show the magnitude of the stimulation artifact before learning (left) and after learning (right). Middle, stimulation command indicates when the stimulation is being applied. Bottom, time frequency representation of the LFP. During stimulation there strong signal corruption across frequency bands. After learning, the power in the higher frequencies diminishes reflecting the removal of the stimulation artifact.

In another experiment we tested whether the Adaptive FIR filter could remove stimulation artifacts during a seizure. In this experiment, seizures were induced using the protocol described in Chapter 2 section 2.2.2. Without filtering, the stimulation artifact occludes the LFP trace and most of the seizure (Figure 21, left). The corruption of the artifact is clearly visible in the spectrogram as well at approximately 120Hz which corresponds to the stimulation frequency in the experiment. In another seizure, we turned on the filter and within 5 seconds, the Adaptive FIR removed the stimulation artifacts (Figure 21, right). Some artifacts leak through during the learning process (10-13

seconds) and this is reflected in the strong broadband frequency power in the spectrogram. Our results show that the Adaptive FIR is robust to non-stationarites in the LFP like seizures: it can remove the artifact noise while preserving the time domain and spectral features of seizures.

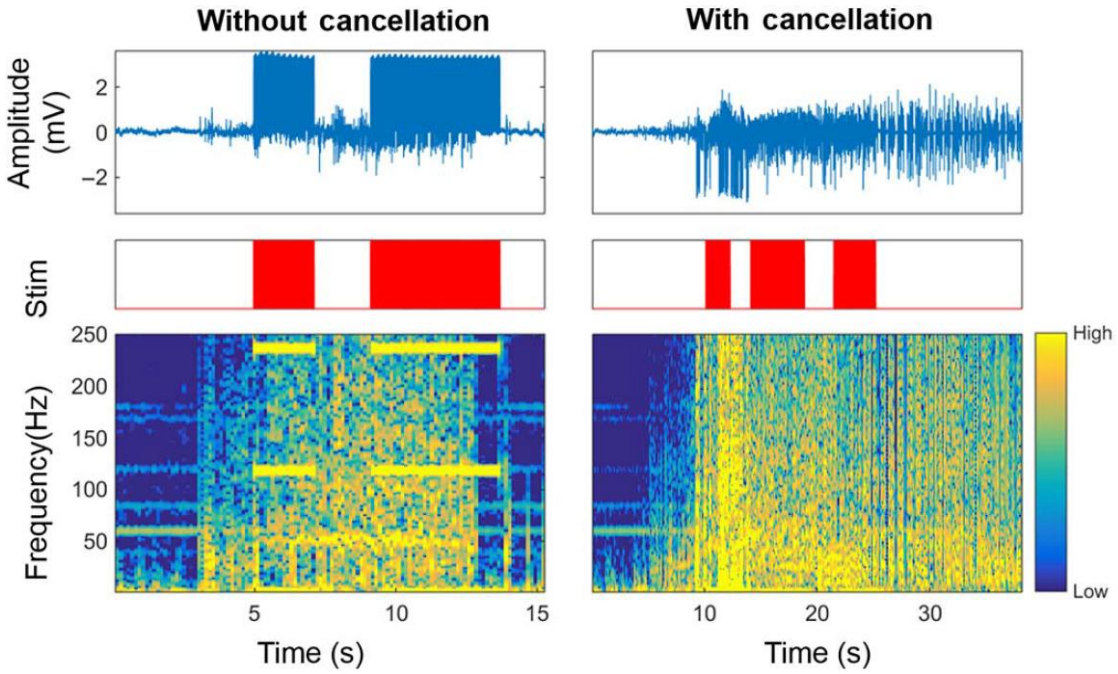


Figure 21. Seizures with and without stimulation artifact removal using Adaptive FIR filter. Top figures, raw LFP data from the rat hippocampus. Middle, stimulation command signal indicates when the pulse is administered. Bottom, spectrogram showing stimulation artifact corruption at specific frequency bands.

## *3.6 Discussion*

We examined the efficacy of five different algorithms to remove stimulation artifacts in a computational model and *in vivo* data. Our results indicate that the optimal

filter for DBS artifact removal is dependent on four key factors: 1) the stimulation protocol (i.e. deterministic vs. random stimulation), 2) robustness of the hardware being used (stationary or non-stationary artifacts), 3) ease of implementation and complexity of the filter, and 4) computational cost of the filter. Our results showed that the adaptive FIR and Kalman Filters had the lowest error rates. If the researcher is interested in a low complexity filter, then the Adaptive FIR filter is the best tool. If stimulation artifact removal is one component of a sophisticated control scheme that requires accurate state estimation, then a Kalman Filter may be more appropriate. Additionally, if the algorithm needs to be implemented in hardware with size and memory limitations then it may be preferable to use the Adaptive FIR over the Kalman Filter. In the case where ease of implementation is far more important than performance then Template or Blanking filters could be used. Comb filters should not be used for most applications unless the stimulation paradigm is periodic in nature.

The Adaptive FIR filter is used widely in many engineering applications. In this chapter, we presented data using this filter in a bidirectional neural interface. The simplicity of the Adaptive FIR and medium computational cost makes it amenable to implementation in an embedded device. The ability to remove the stimulation artifact prior to digitization is very important. High impedance electrodes and short distances between the stimulation and recording channels can lead to very large stimulation artifacts. If the amplifier is saturated, then it is impossible to accurately build a model of the stimulation artifact in the digital space. Our data shows that by using an FPGA we

can tune an Adaptive FIR filter in discrete time to remove stimulation artifacts in continuous time. Furthermore, the results provide a proof-of-concept in using adaptive filters to remove stimulation artifacts in real-time using a neural interface. Next generation technologies from medical device companies should implement these methods within the device hardware.

The use of Optimal Control theory in neuroscience applications has increased over the last few years (Shanechi et al. 2016; Shanechi et al. 2013; Yeo et al. 2016). A key assumption in optimal control is that there is continuous reliable state estimation of the system that needs to be controlled. System states are typically estimated using a Linear Kalman Filter, but other variants are available for more complex and non-linear time series. The Kalman Filter stimulation artifact removal algorithm can accomplish two things at once: remove stimulation artifacts while estimating system states in real-time. Controllers can be designed around this approach, reducing the likelihood of state estimation errors. Ultimately, the choice of which stimulation artifact removal algorithm to use is at the discretion of the experimenter or clinician. This chapter clearly presents the advantages and disadvantages for each method in an effort to facilitate the use of closed-loop algorithms for neuromodulation.

# Chapter 4 Optimization of stimulation parameters for seizure control using model-free reinforcement learning

*Excerpts, including text and figures from (Nagaraj et al., 2016) have been included in this chapter.

## *4.1 Introduction*

Limitations of current DBS methods include a large parameter space and highly variable patient response over time. While current state-of-the-art technologies are effective in controlling seizures for some patients, there is still a need for improvement. To address these limitations many groups have made advances in closed-loop neuromodulation. Closed-loop feedback for seizure control has been tested using three different methodologies. First, the neural signal is used to determine when to apply a stimulus. Stimulus administration in response to seizure onset has successfully suppressed seizures using electrical stimulation (Bikson et al., 2001; Schiller and Bankirer, 2007; Good et al., 2009; Nelson et al., 2011) and optogenetic inhibition (Krook-Magnuson et al., 2013). Secondly, feedback can be used for stimulus optimization over trials in order to maximize the therapeutic effect. Machine-learning algorithms have been used to optimize stimulus parameters by measuring the efficacy of different stimuli and their effects on seizure duration and frequency (Panuccio et al., 2013). Lastly, closed-loop feedback can lead to the modulation of a stimulus parameter in real-time based on physiological measures. Seizures have been suppressed by modulating DC fields proportionally to the activity measured (Gluckman et al., 2001) and by applying precisely timed transcranial electrical current stimulation at certain phases of spike-and-wave ictal behavior (Berenyi et al., 2012).

We propose that a reinforcement learning algorithm that is constantly vigilant and can detect subtle changes in therapeutic efficacy may result in better patient outcomes.

Furthermore, a reinforcement learning algorithm may be able to detect changes in the patient's needs and adapt accordingly. In this chapter we will present a novel approach to controlling seizures, based on the temporal difference reinforcement learning algorithm in a computational model of epilepsy called Epileptor. The algorithm, through an iterative process, determines a policy so as to maximize reward while minimizing a cost function (Sutton and Barto, 1998) (Figure 24). In this instance, the policy is a mapping between different physiological states in an epileptic system and specific stimulus parameter combinations. The reward signal is inversely proportional to an epileptic biomarker, therefore greater reward means less epileptiform activity. The cost function tries to minimize the total amount of stimulation energy for any state-action pair. We define energy as the product of stimulation amplitude squared and the total duration of stimulation in a simulation.

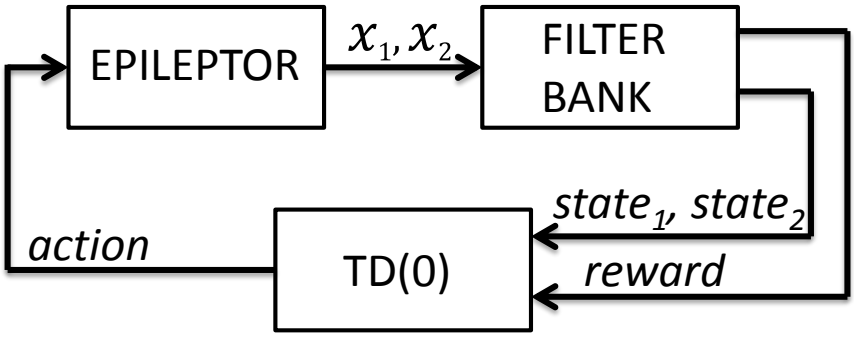


Figure 22. A schematic representation of the seizure control paradigm with the computational model of seizures. Local field potentials (LFP) generated by the Epileptor model ($-x_2+x_1$) are filtered (high and low pass filter) to estimate the current state of the model. A reward signal is also generated from the LFP such that it decreases with seizures and increases with time following a seizure. The reinforcement learning algorithm (TD(0)) integrates state and reward information to determine the optimal action (stimulation frequency) that will maximize the reward provided the current state. The action selected is the stimulation frequency to apply to the Epileptor model.

We use a reduced temporal difference reinforcement learning model, TD(0), that is able to learn the optimal stimulus parameters by exploring the possible stimulation parameters and observing the results. We will show that the learning algorithm converges to a low energy solution.

## *4.2 Methods*

### 4.2.1 Epileptor Model

To the best of our knowledge the Epileptor is the only computational model that captures the transition from non-seizing (inter-ictal) to seizing (ictal) local field potential (LFP) activity. We tested various stimulation paradigms on the Epileptor mean-field model (Jirsa et al., 2014; Proix et al., 2014). The Epileptor model reproduces many of the invariant seizure characteristics documented across species including fast oscillations, spike and wave events and logarithmic increase in inter-spike interval as the seizure approaches termination (Jirsa et al., 2014).

This model provides a platform in which to test closed-loop adaptive algorithms. An example of the transition between inter-ictal to ictal can be seen in Figure 23 (top) and is confirmed by a clear increase in power between 1Hz and 25Hz, Figure 23 (bottom). For our simulations, we use a bandpass filtered signal between 2-15Hz to calculate the reward which captures the fast oscillations in the LFP.

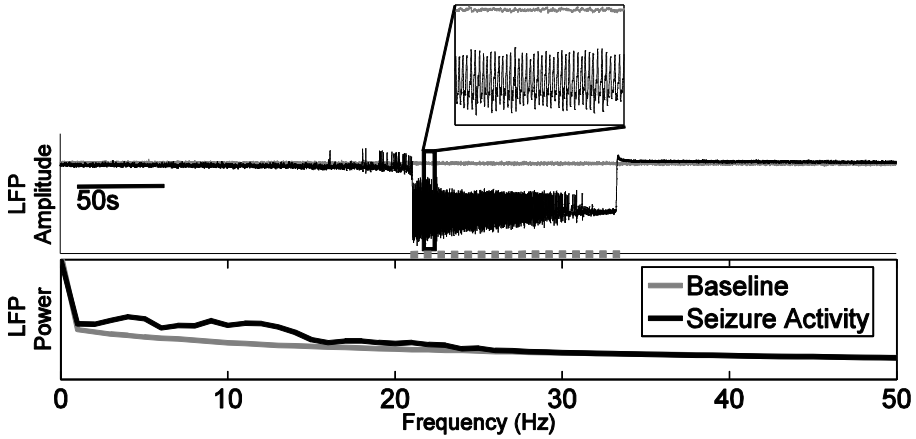


Figure 23. Example of the LFP generated by the Epileptor model. Top, sample trace of the Epileptor model during a non-seizing state (gray) and seizing state (black). Seizures are clearly indicated by prominent DC shift in the LFP. Duration of the seizure, from onset to offset, is indicated by the grey dotted line. Bottom, power spectral density estimate of non-seizing vs seizure activity. The y-axis shows power at different frequency bands on a log scale.

The parameters in the model were empirically fit to generate the different behaviors seen in physiological recordings. The full equations are included in section 6.4Appendix B. The derivation of the equations for the Epileptor model are described in two papers from the Jirsa lab (Jirsa et al., 2014; Proix et al., 2014). The model is represented as a system of 5 ordinary differential equations with three different time scales. The fastest time scales $(x_1, y_1)$ generate fast oscillations, the medium time scales $(x_2, y_2)$ generate spike and wave events, and the long time scale slow permittivity variable $(z)$ determines the seizure duration and frequency. The LFP is determined by adding $-x_1$ and $x_2$. They also developed a reduced 2-dimensional model (Proix et al., 2014) to characterize seizure onset and offset behaviors using only the fast $x_1$ and slow $z$ parameters, as described in 0.

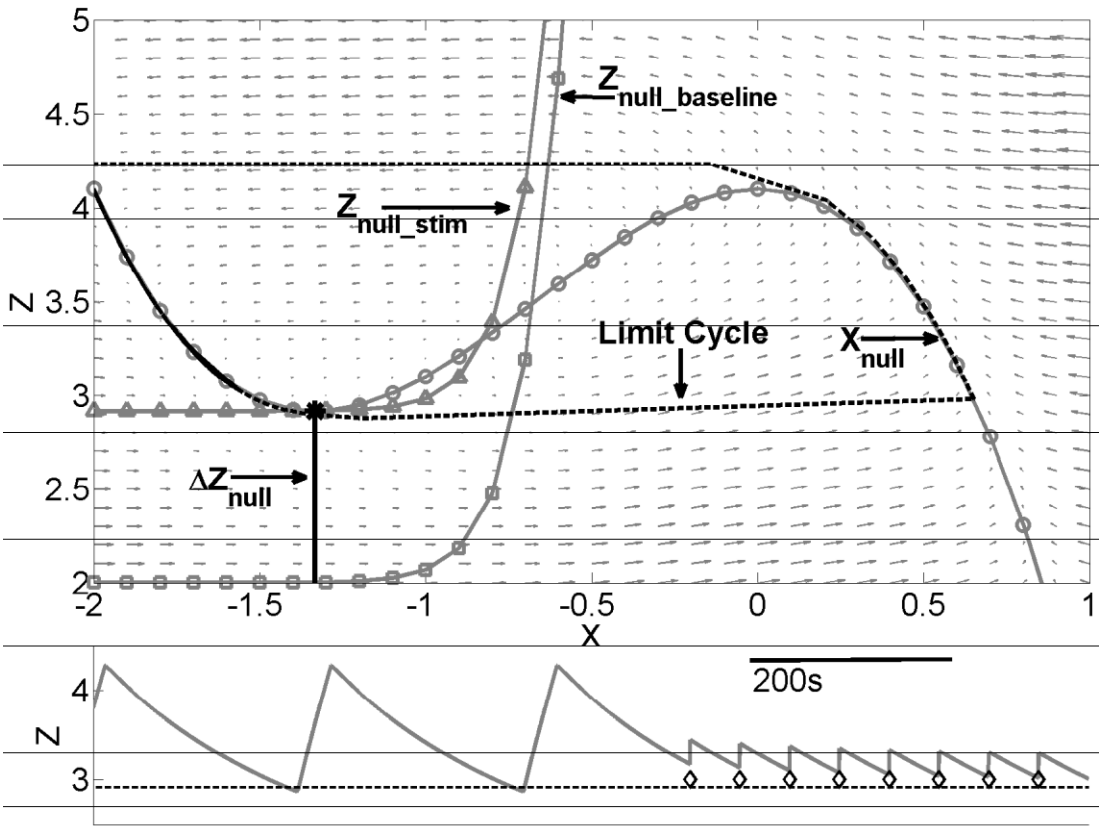


Figure 24. Prediction of minimum energy stimulus parameters to suppress seizures using a state space analysis of the Epileptor model. Top, vector field of the reduced Epileptor model with $x$-nullcline (gray-circles) and baseline $z$-nullcline (gray-square) and $z$-nullcline with stimulation (gray-triangle). The black solid line between the $x$-nullcline and $z$-nullcline (with stimulation) indicates the distance the baseline $z$-nullcline needs to move in order to generate a stable fixed point (large black asterisk) which leads to seizure suppression. The limit cycle is represented by the trapezoidal shaped dotted black line. The bifurcation, from inter-ictal to ictal, occurs when the $z$ value crosses ≈ 2.91. Bottom, example trace of the $z$ variable before and during periodic stimulation (diamonds). The y-axis values correspond to the changing $z$ values as the system moves along the limit cycle in the top figure. The $z$ variable does not descend below the seizure threshold (black dashed line) during stimulation.

The reduced model does not produce the high frequency burst activity seen in the full model but it does reproduce the duration of the seizure and intervals between the seizures. The limit cycle in (top) captures the reduced system cycling through inter-ictal

(left) and ictal (right) phases. We used the reduced model to visualize how stimulation affects the trajectory of two state variables responsible for ictogenesis. The goal of the state space analysis is to determine the minimum stimulation required to cause a bifurcation that terminates seizures. Both models were integrated with an Euler-Maruyama integrator with a time step of 1 millisecond. Longer time steps lead to numerical errors, while smaller time steps did not improve resolution of salient features of the Epileptor.

To simulate electrical stimulation, we applied pulses of one millisecond width to the $z$-variable. Periodic stimulation of the $z$-variable using sufficient energy leads to seizures control (bottom, Figure 24). Changing the $z$ time constant $\tau_0$ modulates the duration of the seizure and the inter-seizure intervals (ISIs). For a given $\tau_0$ the model will produce seizures that occur at almost regular intervals.

### 4.2.2 Determining minimum stimulation frequency to suppress seizures

To determine the minimum stimulus frequency while holding stimulation pulse amplitude to achieve seizure suppression in the Epileptor model, we conducted a state-space analysis. The dynamics of the seizure duration and interval can be understood by analyzing the $x$ and $z$-nullclines of the reduced model. With no stimulation input, the $x$ and the $z$-nullclines have a single unstable equilibrium point at the crossing, allowing for a stable limit cycle that determines the seizure duration and inter-seizure intervals. Positive stimulation current pushes the $z$-nullcline up. With sufficient current, the $x$ and the $z$-nullclines go through a bifurcation resulting in a stable fixed point and seizures

stop. The minimum pulsatile amplitude necessary to stop the seizures can be determined by the distance of the $z$-nullcline to the knee of the $x$-nullcline ($\Delta Z_{null}$), as shown in top of Figure 24.

The $z$-variable relaxes slowly back to its resting point after each pulse at the $z$ time constant $\tau_0$, as shown in the bottom of Figure 24. Therefore, to assure seizure control throughout the simulation, the minimum $z$ value between pulses must be greater than $\Delta Z_{null} + x_0$. This will keep the $z$ variable above the bifurcation point Fig. 3 (bottom).

Provided the stimulus amplitude ($S_a$), stimulus interval ($S_{IPI}$), and the time constant of $z$ ($\tau_0$) it is possible to calculate the approximate minimum distance $z$ needs to travel to control seizures:

$$z_{min} \approx x_0 + \frac{S_a}{1-e^{\left(-\frac{S_{IPI}}{\tau_0}\right)}} \tag{4.1}$$

This $z_{min}$ value is the minimum distance required to suppress seizures and is $z_{min} = 2.91$. The parameter $x_0$ determines at what value the floor of the z-nullcline is at (in Figure 24 it is at 2.0). If $x_0$ is set at ≈ 2.91 then there will be no seizure events. This is because a stable fixed point forms at the intersection of the $z$ and $x_1$ nullclines. For all simulations, the $x_0$ variable was set at 2.0. This, however, is the minimum pulsatile perturbation to suppress seizures, not all epileptiform activity. Increasing the stimulation above this can produce significant improvements by suppressing high frequency activity seen in the full model when the system is close to the bifurcation point.

Alternatively, we can calculate the stimulus amplitude for any arbitrary stimulus interval.

$$S_a = \Delta Z_{null}\left(1 - e^{\left(-\frac{S_{IPI}}{\tau_0}\right)}\right) \quad (4.2)$$

Where, $S_a$ is the minimum stimulus amplitude to suppress seizures given the stimulus intervals ($S_{IPI}$) and $z$ variable time constant ($\tau_0$). The simulations presented in the chapter use integer value stimulus intervals. A map of analytically calculated minimum pulsatile stimulation amplitude necessary for seizure suppression given different integer stimulation frequencies is shown in Figure 25. Given a specific stimulation amplitude-frequency pair, we ran simulations to compare the minimum frequency (while holding stimulation amplitude) to the analytically calculated stimulation frequency for the given amplitude.

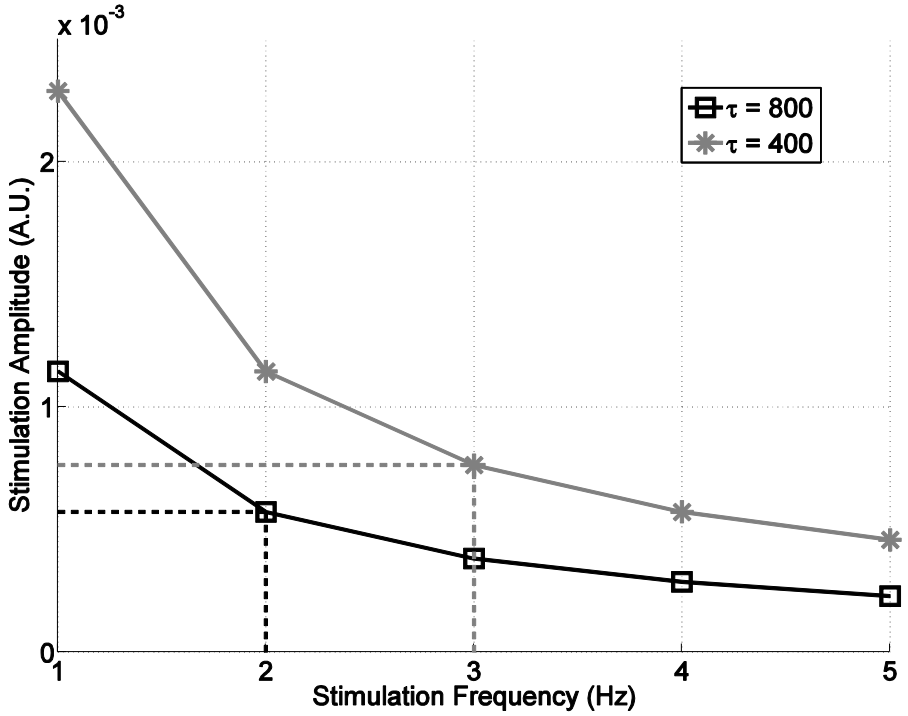


Figure 25. Map indicating minimum stimulus frequency and amplitude to control seizures given the time constant for the $z$ variable (legend). Stimulus parameter combinations below the solid lines will fail to suppress seizures

### 4.2.3 Temporal difference reinforcement learning algorithm implementation

Reinforcement learning has a long and intricate history and has been reviewed in detail elsewhere (Sutton and Barto, 1998). These algorithms have been implemented in many different neural-engineering problems, including seizure control (Bush and Pineau, 2009; Pineau et al.). However, in these cases, the investigators develop a stimulation policy with a reinforcement learning algorithm using off line training data. Optimal parameters were determined through dynamic programming principles using the full data set, a process known as batch learning. Our approach substitutes batch learning for an iterative learning process where optimal stimulation parameters are determined as data is collected in real-time. The decision to use an online learning scheme, as opposed to offline training, was motivated by the fact that the effects of stimulation for each experimental preparation (and patient) seems to be different depending on subtle details as the positioning of the stimulation electrode. Thus, optimizing parameters based on data collected from one patient or experiment and applied to another may not result in the best control.

The full temporal difference learning algorithm uses memory in the form of eligibility traces (Sutton and Barto, 1998) to identify sequences of actions necessary to achieve a high reward state. However, in simulations we found little benefit from using eligibility traces, presumably because the occurrence of seizures is stochastic, and therefore was not used. This way, only the current state-action pair is updated. In this

study we show how a learning algorithm can determine the optimal stimulation parameters for a seizure model using domain knowledge. Furthermore, the algorithm can respond to non-stationarities of an epileptic system and adapt accordingly.

#### *4.2.3.1 Reduced Temporal Difference Learning Algorithm: TD(0)*

The TD(0) algorithm generates a map relating the state of the system, to the expected reward given a selected action: $Q(s,a)$. Where $Q$ is the expected reward, proportional to the time since the last seizure, $a$ is the selected action from a set of stimulation frequencies, and $s$ is the state of the Epileptor (i.e. seizing or not seizing). When the simulation is initialized the algorithm is naïve, and has no information about the expected reward for each state-action pair. To initialize $Q(s,a)$ we use the average reward over a few seizures without stimulation, $\mu_0$, plus additive zero mean white noise with small variance $\sigma^2 \ll 1$.

Once the algorithm is turned on, it chooses the next action a' expected to produce the highest reward, $Q(s',a')$, given the current state $s'$. After the action is executed $(a' \rightarrow a)$, the actual reward $(R)$ is measured. The error between the measured reward and expected reward is calculated as follows:

$$\delta = R - Q(s,a) + \gamma * Q(s',a') \quad (4.3)$$

Where $\gamma$ is the delay discounting factor that accounts for how much the algorithm values future expected reward. We set to $\gamma$ to zero $(\gamma = 0)$ to favor more greedy behavior to speed up convergence. Initially, because the matrix $Q$ is arbitrarily set at a high value,

the error, between the measured reward and the expected reward will be great. Using this error the map is updated to increase the accuracy of the predicted reward value with the following update equation:

$$Q(s,a) = Q(s,a) + \alpha * \delta \quad (4.4)$$

Where $0 < \alpha \ll 1$ determines how quickly $Q(s,a)$ is updated given the error to the measured reward.

#### *4.2.3.2 Action Selection*

If at every time step the action providing the greatest expected reward is selected, this algorithm is called "greedy". However, it may result in convergence to action selection policy at a local minimum rather than allowing for exploration to find a global minimum. Therefore, using a selection policy that explores some actions that may at first appear sub-optimal can help find the optimal solution more reliably than a purely greedy algorithm. An alternative to the purely greedy approach is the Softmax policy that selects actions with probability proportional to the expected reward. So, the action with the highest expected reward will have a higher probability of being selected than other actions. The probability of each action, a' selected at each state, s' can be calculated as:

$$P(a'|s') = \frac{e^{\frac{Q(s\prime,a\prime)}{\tau_s}}}{\sum_{i=1}^{n} e^{\frac{Q(s\prime,a_i\prime)}{\tau_s}}}. \quad (4.5)$$

The variable $\tau_s$ is the temperature value for the Softmax policy, which determines how sensitive the probability is to differences in expected reward, and n is the total number of actions.

#### *4.2.3.3 Learning Rate*

The degree to which the TD(0) algorithm weights new information compared to its current estimate of the reward for a selected state-action pair is determined by the parameter $\alpha$. When the model is completely naïve, using a large α results in updating the map rapidly. This method can lead to a quick convergence to an action that is effective, but the solution may not be the global minimum. We selected $\alpha$ to be long enough to average information over the duration of one inter seizure interval. The algorithm updates the action in windows ($w = 15s$). Given the duration of an $ISI$, we set $\alpha$ so that the time constant is about one seizure interval:

$$\alpha \approx 1 - e^{(-w/ISI)} \tag{4.6}$$

A full description of the algorithm is given in section 6.4Appendix B.

### 4.2.4 Ictal and inter-ictal state estimation from LFP data

The Epileptor model has a continuous state-space and a reward signal must be derived from the measured signals. Domain knowledge (i.e. spectral characteristics, Figure 23) was used to establish the reward function and discretize the state-space to make the model more amenable to the simplest Temporal Difference algorithms. In the Epileptor model, the ictal vs inter-ictal dynamics can be separated in state space by plotting the $x_1$ variable against $x_2$ as shown in Figure 26C. The inter-ictal activity is in the left side of the state space and during ictal activity it is on the right side. The goal is to generate some state space representation from the LFP that separates the inter-ictal from

the ictal activity.

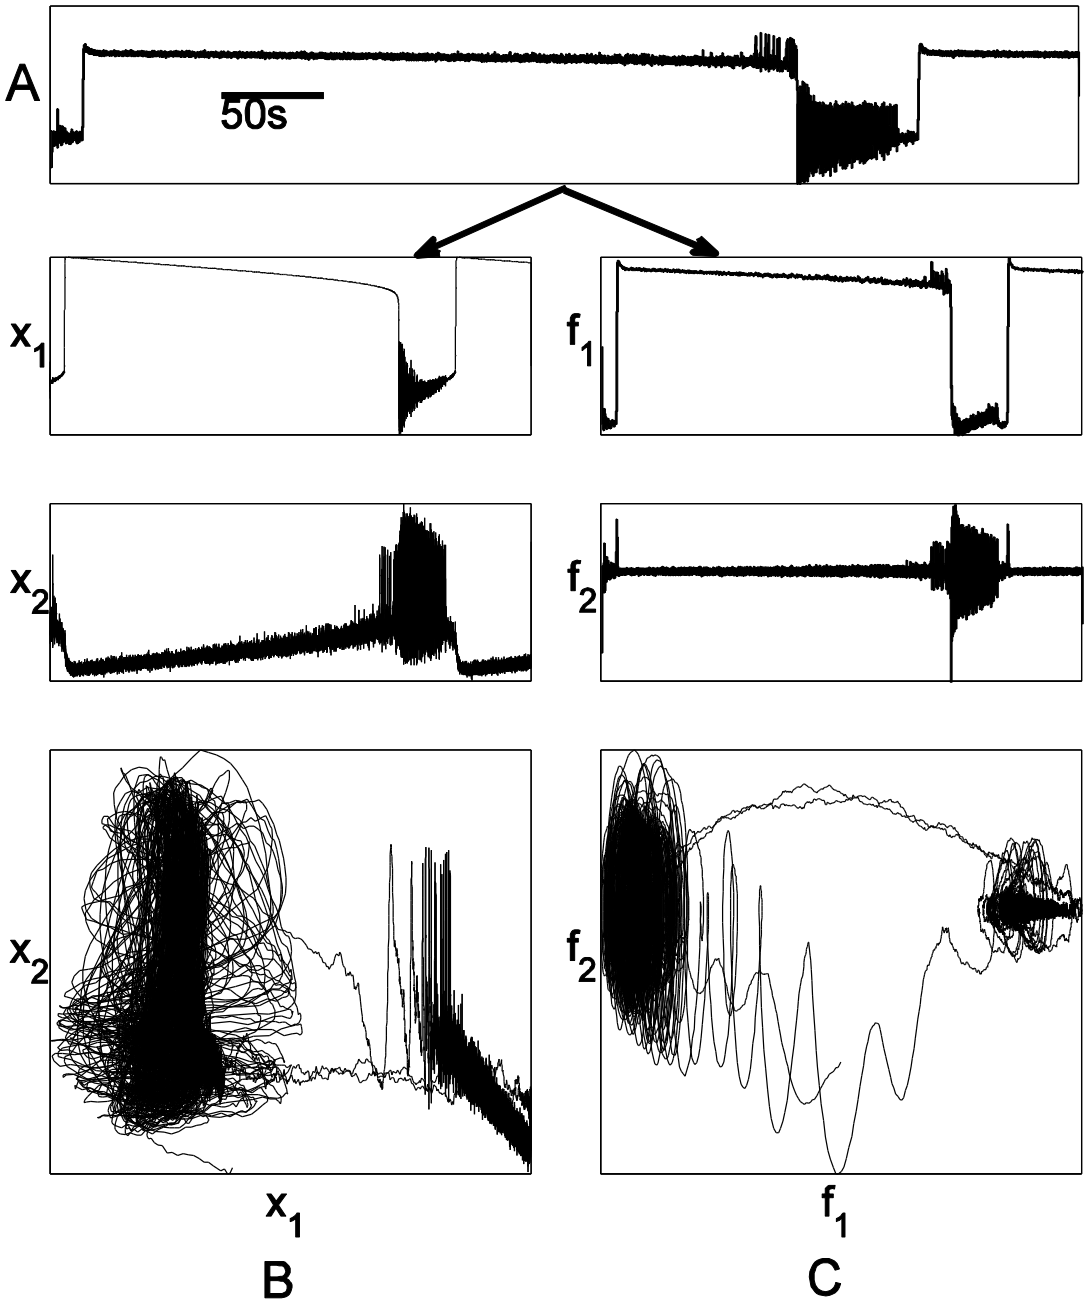


Figure 26. Decomposition of the Epileptor LFP using a filter bank results in feature space closely resembling the actual state-space of the Epileptor dynamic variable $x_1$and $x_2$. A top, raw LFP generated by the Epileptor model. B, State-space of the Epileptor model

parameters $x_1$and $x_2$. Top, the $x_1$variable time series, middle, the $x_2$ variable time series, bottom, the state space from $x_1$and $x_2$ clearly distinguishes inter-ictal (right) from ictal (left) epochs. C, Estimation of the state-space using a filter bank. Top, low-pass filtered LFP data, middle, high-pass filtered LFP data, bottom, feature space clearly shows an inter-ictal cluster (right) and ictal cluster (left). Cutoff frequency was 0.5Hz.

By filtering the data with a high pass filter and using another low-pass filter, the two signals resulted in a state space representation very similar to that seen by plotting the $x_1$ variable against $x_2$. Data was filtered using first order Butterworth high pass (Figure 26A top) and low pass filters (Figure 26A bottom) each with 0.5Hz cutoff frequency. Filter coefficients were calculated using the MATLAB *butter* command. Filter outputs were normalized to each axis in the two-dimensional state-space. This state space is partitioned into ictal and inter-ictal states. A state index is assigned at each time step depending on where state value falls in the inter-ictal or the ictal partition. Finer resolution partitions are possible, but more partitions increased the exploration time.

### 4.2.5 Reward

The reward for each $Q(s,a)$ is calculated at each iteration by taking the negative log of the band-pass filtered local field potential ($LFP_{bp}$, 2-15Hz). The reward for each time step is calculated as follows:

$$\tau_{fiter} = \left(\frac{1}{1-e^{\left(\frac{-1}{\tau_{ISI}}\right)}}\right) \tag{4.7}$$

$$\overline{LFP}_t^{bp} = \frac{(LFP_t^{bp} * LFP_t^{bp})}{\tau_{filter}} + (\tau_{filter} * \overline{LFP}_{t-1}^{bp}) \tag{4.8}$$

$$R_t = -\log\left(\overline{LFP}_t^{bp} - Cost\right) \quad (4.9)$$

Where $\tau_{filter}$ is the decay rate for the smoothing filter, $\overline{(LFP)}_t^{bp}$ is the smoothed power of the LFP at time $t$, and $R_t$ is the reward at time $t$. The tau-filter smooth's the noisy reward signal to remove noise in the raw reward signal. Noise in the raw reward signal can lead to misleading updates of the state-action function $Q(s, a)$.

Each stimulation parameter combination has an associated cost. The cost function is as follows:

$$Cost = \left(S_a^2 * \left(\frac{1}{S_{IPI}}\right)\right) * CW \quad (4.10)$$

Where $CW$ is the cost weight, $S_{IPI}$ is the stimulation inter-pulse interval, $S_a$ is the stimulation amplitude, and $\mu_0$ is the mean reward during the simulation when stimulation pulses were not applied.

## *4.3 Results*

### 4.3.1 Open-loop stimulation controls stationary seizures

We first tested open-loop stimulation on seizure control in the Epileptor model. Seizures were suppressed with sufficiently high stimulation frequency and amplitude. The minimum stimulation energy required to suppress seizures depended on the time constant of the $z$-variable of the model; the faster the time constant, the more stimulation energy is required to suppress seizures. Examples of stimulation applied to the Epileptor

model when set at two different time constants, $\tau_0 = 800$ and $\tau_0 = 400$ seconds, are shown in Figure 27. The stimulation frequency analytically calculated to suppress seizures (1.04Hz) when the Epileptor's time constant is set at 800 seconds suppresses all seizure like activity (Figure 27, top), while it is necessary to stimulate at 2.08Hz to achieve control when the system has a time constant of 400 seconds (Figure 27, middle). The high frequency activity is distinguishable from a seizure because the system did not undergo a saddle-node bifurcation at the seizure onset. Seizures in the Epileptor model are characterized by a prominent DC offset, which is a key feature of saddle-node bifurcations (Jirsa et al., 2014). Applying the minimum stimulus frequency to suppress seizures for $\tau_0 = 800$ in a simulation with $\tau_0 = 400$ seizure dynamics does not suppress the seizure (Figure 27, bottom).

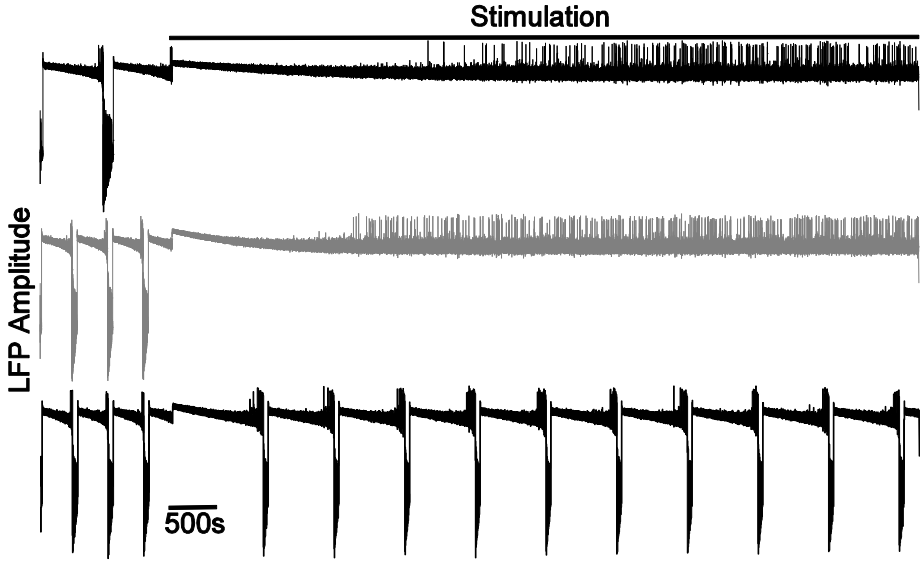


Figure 27. Open-loop stimulation can control seizures if sufficient stimulation energy is applied. Top, applying stimulus pulses into the z variable at the minimum calculated stimulus frequency (1.04Hz) with slow seizure dynamics ($\tau_0 = 800$) controls seizures. Middle, similarly, seizures can be controlled if the seizure dynamics are faster ($\tau_0 =$

400) when stimulation is administered at the minimum calculated frequency (2.08Hz). Bottom, insufficient stimulus energy cannot control seizures. In this example 1.04Hz was applied to seizures with $\tau_0 = 400$. Black bar indicates duration of stimulus. Note that the time scale has been condensed to see how dynamics evolve on a longer time scale, and the seizures appear as abrupt downward deflections (DC offset) in the LFP.

The map in Figure 25 shows the necessary stimulation parameters to achieve seizure control for different time constants of the model as determined by analysis of the reduced model. Parameters used in Figure 27are indicated in the map shown in Figure 25. While the parameters necessary to suppress seizures are calculated from the reduced model, high frequency activity may persist even though the seizures have been stopped. Increasing the stimulation frequency above this minimum threshold of seizure suppression may provide further improvement by reducing this "epileptiform" behavior. Therefore, we consider the optimal parameter combination for epileptogenic feature control to be slightly higher than the minimum necessary to suppress seizures.

### 4.3.2 Closed-loop adaptive stimulation results in seizure control while minimizing energy

Having established the optimal stimulation parameters, we next tested whether the TD(0) algorithm, starting with no a-priori knowledge, can converge to an effective stimulus parameter solution. The goal of the TD(0) algorithm is to maximize reward through selection of actions, in this case, stimulation frequencies at each state (ictal vs. inter-ictal). The reward trace increases as activity decreases, therefore suppressing seizures results in an increased reward. A small cost is attributed to each stimulation pulse which helps refine the stimulation frequency once a range of stimulation actions

that suppress seizures are found. At the onset of the simulation the model is naïve and set to random initial conditions. We set the learning rate to reflect the time course of a seizure (equation 6). This results in rapid learning with fairly stable behavior once good control is achieved.

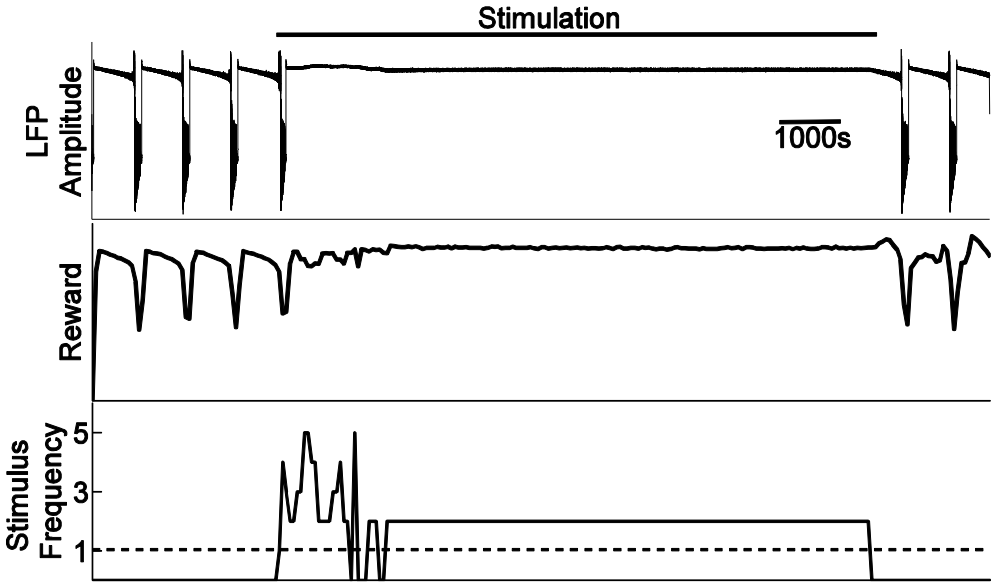


Figure 28. TD(0) algorithm converges to optimal solution (2Hz) when seizure dynamics are stationary. Top, raw LFP trace. Middle, reward trace associated with the LFP. Bottom, Stimulation frequency selected at each action selection interval. The TD(0) algorithm explores the state-action space at first because the initial expected values for each state-action is high. Dotted black line represents the minimum stimulus frequency required for control (1.04Hz). Black bar indicates duration of stimulus.

An example of the TD(0) algorithm identifying optimal seizure control parameters is shown in Figure 28. Simulations were run for 15000 seconds on a Windows 7 with an Intel (i7) 3.5 GHz processor. The TD(0) algorithm computation time was 160ms. High frequency activity and seizures correspond to downward deflections of the reward signal. After probing the state-action space following stimulation onset (black bar) the policy

converges to the optimal state-action solution that suppresses all epileptiform activity. Since the action choices are discrete frequencies, the optimal frequency closest to the theoretical optimum was 2Hz. While the stimulation energy is higher, the total epileptiform activity present during the simulation is lower than seen in Figure 27 (top). Once the optimal stimulus frequency is identified, the Softmax policy for action selection using a very low temperature maintained the stimulation parameters throughout the remainder of the simulation. Turning off stimulation leads to reemergence of seizures, shown at the end of the simulation.

Increasing the temperature value ($\tau_s = 1$) in the Softmax action selection algorithm alters the action selection policy (Figure 29). Instead of converging on a single stimulation frequency to suppress the seizures, it converges on a distribution from which the TD(0) selects stimulation frequencies probabilistically based on expected value. In this case the average stimulation frequency was 2.05Hz.

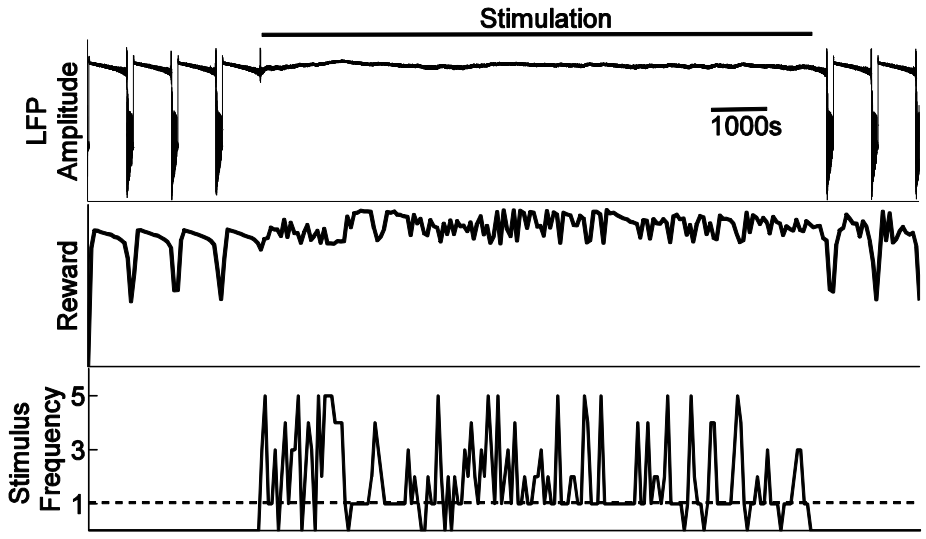


Figure 29. TD(0) algorithm action selection with high Softmax temperature value. Softmax policy action selection is more explorative when temperature values are high ($\tau_s = 1$). Top, raw LFP trace. Middle, reward trace associated with the LFP. Bottom, Stimulation frequency selected at each action selection interval. Stimulus frequencies selected vary between 0 and 5 hertz, but the average stimulus frequency is approximately 2.05Hz. Dotted black line represents the minimum stimulus frequency required for control. Black bar indicates duration of stimulus.

The effects of temperature in the Softmax algorithm on the number of escaped seizures, convergence time and stimulation frequency is shown in Figure 30. At very low temperature values ($\tau_s = 0.01$) the TD(0) algorithm behaves in a greedy fashion and converges on a solution quickly (Figure 28).

As the temperature value increases, the algorithm spends more time searching the parameter space, which ensures a global minimum at the cost of time of convergence. The mean convergence time increases to over 2000 seconds when the temperature value $\tau_s = 0.2$. However, at very high temperatures we found a surprising behavior. The algorithm does not settle on any single stimulus frequency, but instead uses a distribution of stimulus parameters and from which is selects probabilistically to deliver stimuli and

optimizes the distribution (Figure 29) resulting in a good convergence rate to the globally optimal solution. Ultimately, there is a trade off in the choice of temperature value for the Softmax action selection policy. Lower energy solutions are typically less robust. For instance, the lowest energy solutions occur when the temperature value is between 0.1 and 0.2 (Figure 31), but this results in the highest number of seizures after onset of the controller.

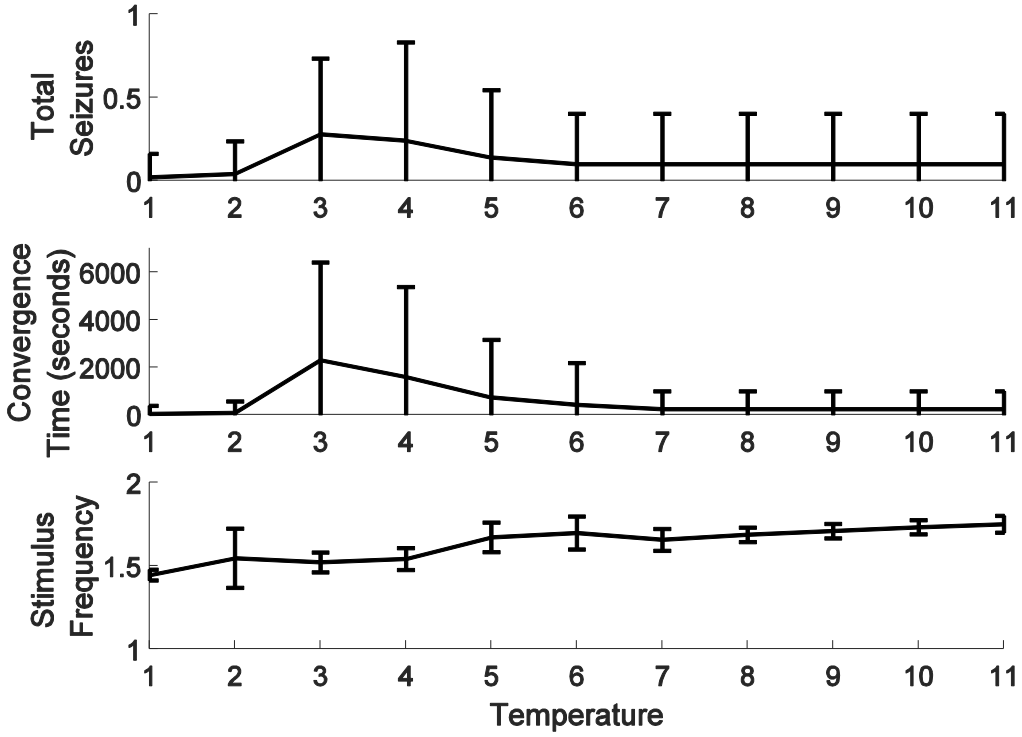


Figure 30. Performance of reinforcement learning algorithm at different temperature values of Softmax action selection for large ISI ($\tau_0$=800). Top, mean total number of seizures that escape "therapy" when the reinforcement learning algorithm is choosing stimulation actions. Middle, mean time to converge (seconds) on a policy that controls seizures. Bottom, average stimulus frequency during the simulation.

### 4.3.3 TD(0) converges to optimal solution when seizure dynamics are

## faster

We tested whether the TD(0) algorithm could adapt and maintain stimulation at optimal stimulation frequency when the parameter $\tau_0$ of the Epileptor model was changed. In these simulations, the seizure frequency increased, thereby requiring a net increase in stimulation energy, as shown in Figure 31. During the learning process, two seizures escape requiring further refinement in action selection. Towards the end of stimulation, the algorithm converges on 3Hz as the optimal stimulation frequency for controlling seizures with a fast time constant ($\tau = 400$).

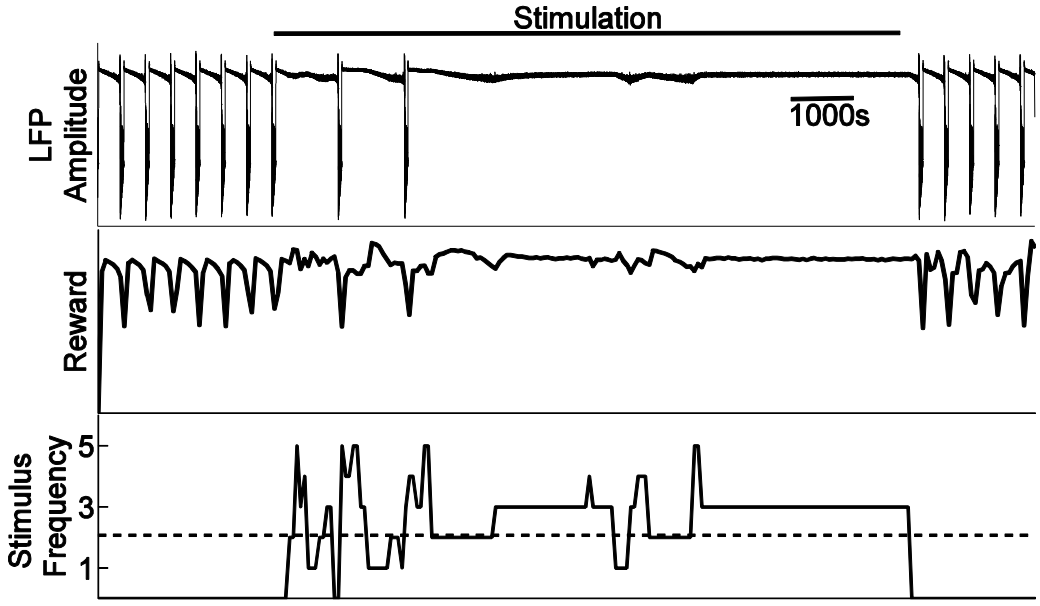


Figure 31. TD(0) algorithm converges to optimal solution (3Hz) when seizure dynamics are faster ($\tau = 400$). Top, raw LFP trace. middle, reward trace associated with the LFP. Bottom, stimulation frequency selected at each action selection interval. The TD(0) algorithm needs to explore the state-action space for a longer period of time when the ISI is smaller. Dotted black line represents the minimum stimulus frequency required for control (2.08Hz). Black bar indicates duration of stimulus.

## *4.4 Discussion*

This chapter presents a reinforcement learning algorithm to determine the optimal stimulation parameters for seizure control. Specifically, we use a reduced TD(0) algorithm to determine low energy stimulation parameters and tested the approach in a computational model of epilepsy called Epileptor.

### 4.4.1 State-space approximation

The TD(0) algorithm determines the optimal stimulation parameters given an estimate of the current state of the system. The optimal action for the ictal-state may be different than that found during the inter-ictal state. This allows the algorithm to optimize stimulation parameters to prevent seizures as well as find the optimal parameters to terminate seizures, which may be different. To separate seizing vs. non-seizing states we used a two state measurement based on the power from low and high pass filtered raw LFP data. Seizure control is achieved when a therapy brings the system to a state where seizures are unlikely to occur. Seizure termination on the other hand can be defined as the stimulus induced suppression of seizures following seizure initiation. While seizure termination was not the central focus of this chapter, the TD(0) algorithm nonetheless finds the optimal solutions based on the state of the Epileptor (ictal vs. inter-ictal), and ultimately optimizes for both seizure suppression and seizure termination.

With finer state space partitioning, or the addition of additional feature dimensions, it may be possible to identify pre-ictal states in which stimulation could prevent seizures. These states would be different from inter-ictal states, where stimulation has very limited

benefit on seizure suppression. Further division of the state-space however, results in additional state-action pairs, which increases the computational cost. Without any underlying model of the system, this can result in long searches before finding an efficient policy. Therefore, in this chapter we focused only on the simple state space partition of ictal and non-ictal but finer partitioning with rapid identification of a good control policy remains a problem that needs to be further investigated.

A limitation to this work is how well the design choices (i.e. reward function and state space discretion) would carry over to other computational and experimental models. Our design approach was specific for the Epileptor model, and implementation of the TD(0) algorithm in other contexts may require a different reward function and state-space discretization.

### 4.4.2 Stimulation variables

In this chapter, TD(0) was used to optimize stimulus pulses applied to the $z$-variable, the slow permittivity variable in the Epileptor model. The $z$-variable may be a physiological correlate for K+ dynamics (Jirsa et al., 2014), or could potentially even model adenosine dynamics (Van Gompel et al., 2014). Experimental evidence shows that electrical stimulation directly modulates adenosine concentrations (Van Gompel et al., 2014) and fast stimulation causes an increase in extracellular potassium concentration (Bikson et al., 2001).

The TD(0) algorithm can also be used to optimize stimulation applied to the $x_1$ and $x_2$ state variables, from which the LFP is directly estimated. In modeling experiments,

not shown in this chapter, the reinforcement learning algorithm was able to control seizures by applying stimuli to these variables. Since the time-constants of these variables are much faster than the $z$- variable, higher frequency stimulation was needed to achieve similar seizure control.

In this study, we held stimulus amplitude constant while optimizing over the stimulus frequency. The map in Figure 25 indicates that the opposite is also possible: stimulus frequency could be fixed while optimizing the stimulus amplitude. We used a discrete pulse because it was the simplest and most commonly used stimulation waveform class in neuroscience. It is possible that other waveform classes could result in a stable fixed point with less energy, but exploration into this is outside the scope of this chapter. Ultimately, the Epileptor model is sensitive to total energy, therefore optimizing over stimulus frequency and amplitude can be done, but the number of state-action pairs grows rapidly with the number of parameters to be optimized resulting in long convergence times to a solution.

In the Epileptor model the amount of activity was monotonically dependent on the amount of energy applied to the z-variable. The only cost to high frequency stimulation was from the stimulation energy subtracted from the reward function. However, in experiments, we have found that the effect of stimulation is not monotonic and there are bands of stimulation frequencies that suppress seizures while other frequencies induce seizures. The TD(0) algorithm will find the global minimum for each state-action pair, but at the cost of occasionally inducing seizures that occur while testing stimulation

frequencies that are ictogenic. Black box modeling, specifically input-output models could be empirically determined to predict the effects of periodic pulsatile stimuli at different frequencies. Frequencies that predict emergence of epileptogenic biomarkers could then be cut out of the optimization process, thereby reducing the parameter space.

### 4.4.3 Action Selection

Under the TD(0) framework, action selection is a tradeoff between exploration and exploitation. If the temperature value of the Softmax policy (equation 5) is very low, the TD(0) quickly explores the state-action pairs and then converges on the optimal solution from this estimation. This is clear in the simulation presented in Figure 28. Sometimes the quick solution may not be the global maximum, therefore making greedy choices more likely to be less cost effective in the long run. This approach focuses on finding a single action that it considers as optimal.

High temperature values on the other hand lead to a policy that does not converge to a single action. Instead, actions are selected probabilistically from distribution of actions, as seen in Figure 29. There is no one action, rather the algorithm optimizes the selection probability across all the actions resulting in an average stimulus frequency close to optimal to obtain seizure control.

The action selection paradigms explored in the chapter are discrete in nature (i.e. from a set of stimulation frequencies). This limits the ability of the algorithm to find an optimal solution, especially if it is not in the defined set. One way to circumvent this is to implement other reinforcement learning methods which can be applied to a continuous

action space.

### 4.4.4 Algorithm robustness

To the best of our knowledge there are no other computational models that capture the transition from inter-ictal to ictal activity. While we have not tested if our method is robust to changes in the class of models, we have tested robustness by evaluating the algorithm's performance as parameters of the Eplieptor model were varied. Specifically, we decided to evaluate the algorithm by changing the inter-seizure intervals and seizure durations. These parameters are modulated by the $\tau_0$ parameter. Reducing the value of the $\tau_0$ parameter resulted in smaller inter-seizure intervals and shorter seizure durations Figure 31. The TD(0) was able to determine the optimal stimulation frequency even when seizure rates increased dramatically. These simulations indicate that our method is robust to changes in the seizure dynamics.

## *4.5 Conclusion*

The results presented in this chapter show that a TD(0) algorithm can effectively converge to the optimal stimulus parameter. Here optimization was performed only over stimulation frequency. This approach can be used optimize over both stimulation frequency and amplitude at the cost of increased computation and an increase in convergence time. The method presented is computationally efficient and could potentially be programmed into a small bidirectional neural interface. Future studies will require testing of the algorithm in an in vivo epilepsy model.

# Chapter 5 Model based control of local field potential biomarkers

## *5.1 Introduction*

Open loop neural stimulation, where periodic pulses of voltage or current are delivered to through an electrode to a site in the brain, has been used to treat Parkinson's disease, Epilepsy and a growing number of other diseases. New implantable neural stimulation devices have enabled simultaneous recording and stimulation (Morrell, 2006; Stypulkowski et al., 2014) enabling closed-loop therapies where stimulation is triggered off biomarkers, such as strong beta oscillations for treatment of Parkinson's Disease, and epileptiform activity to suppress seizures. It is hypothesized that stimulation either suppresses these pathological behaviors, or mitigates their effects on other networks to restore their function (Agnesi et al., 2013). However, these therapies, even in reactive neural stimulation, are delivering stereotyped stimulus patterns.

The central hypothesis of this work is that subject specific computational model can be used to calculate stimulus waveforms from the current state of the neural activity to optimally suppress pathological neural activity. If spontaneous neural activity like bursts which have spectral profiles similar to epileptogenic discharges can be suppressed using such a closed-loop controller, it may be used to improve efficacy of neuromodulation therapies for seizure disorders. Here we test a Linear Quadratic Gaussian (LQG) algorithm for delivering state dependent current stimulation to the brain that is optimized to the subject's physiological responses.

Both open-loop stimulation and closed-loop stimulation, currently used for clinical treatment of Parkinson's Disease and epilepsy, deliver pulsatile waveforms. These

waveforms are designed based on stimulation hardware capabilities and safety limits, but they are not physiologically based. Closed-loop stimulation has been used to deliver stimulation by triggering off the detection of epileptiform activity as is the case in NeuroPace's Responsive Stimulator (Morrell, 2006). It can also be used to modulate the stimulus amplitude based on a physiological signal, such as patient orientation for spinal cord stimulation for pain (Schade et al., 2011), or beta oscillation amplitude for Parkinson's disease (Little et al., 2013).

Closed-loop algorithms may also be used to learn and optimize stimulation parameters to assist clinicians in programming neuromodulation devices (Panuccio et al., 2013; Nagaraj et al., 2016). With a fast closed-loop algorithm, stimulation can be phase locked to physiological signals to suppress neural oscillations (Holt et al., 2016; Cagnan et al., 2017). Optimization algorithms have been used to identify stimulus waveforms that can synchronize neural populations (Wongsarnpigoon and Grill, 2010; Nabi et al., 2013a; Nabi et al., 2013b; Wilson et al., 2015; Brocker et al., 2017). The next generation closed-loop device may be able to measure physiological signals, analyze them in real time and deliver state dependent stimulation waveforms to optimally suppress pathological biomarkers with less energy and side effects.

In the mid-20$^{th}$ century, engineers developed closed-loop algorithms to suppress unwanted oscillations or to dampen a disturbance in buildings, airframes and chemical plants (Friedland, 2012; Kirk, 2012). More recently, Optimal Control theory methods have been applied in neuroscience for efficient brain-computer and brain-machine

interface technologies (Shanechi et al.; 2016; Shanechi et al.; 2013; Yeo et al.. 2016). Fundamental to Optimal Control design is a model that can accurately predict the response of the system to an input. The response of the nervous system to electrical stimulation is complex and is not easily predicted based on first principles, but a predictive model can still be developed using a 'black-box' system identification approach, by applying a stimulation, measuring the response and then fitting a model to the data. A variety of different 'black-box' algorithms have been developed that can extract model coefficients from input-output data (Ljung, 1999). The benefit of using Optimal Control methods over classical feedback control is the ability to achieve the control objective while simultaneously minimizing constrains imposed to improve efficiency of the controller. In this study, we will use a Linear Quadratic Gaussian (LQG) controller to determine electrical stimulation waveforms to suppress spontaneous neural activity in rat.

Using a real-time computer, we are able to test and validate closed-loop algorithms *in-vivo* that cannot yet be implemented in currently available implantable hardware. The Real-Time eXperimental Interface (RTXI), an open-source program, enables closed-loop experimental designs (http://rtxi.org). We can use this software platform to implement and test some of the optimal control approaches to dampening neural activity. With a standard desktop computer, this system can sample data and produce outputs at 2000 samples a second. Standard linear algebra software libraries can be used to calculate the optimal stimulus to suppress activity, and control a voltage regulated current source in

real-time with jitter less than 10$\mu$sec. We hypothesize that if neural activity is predictable on time horizons of milliseconds and sufficient electrical stimulation amplitude is used, then we may be able to design a feedback controller to suppress the neural activity using optimal control strategies.

There are several reasons why optimal control strategies could fail to control spontaneous neural activity. The first is that electrical stimulation excites neural tissue, so it is not clear that electrical stimulation alone could result in suppression of neural activity, except through excitation of inhibitory neural populations or inhibition of excitatory neurons through depolarization block (Johnson et al., 2008). Second, neural response to stimulation is highly nonlinear with respect to stimulation parameter combinations (Millard et al., 2013). Lastly, the neural response to stimulation is not stationary (Marmarelis, 2012); a fixed model of the neural response may not provide sufficient predictive power over time to allow for control.

Although neural systems clearly exhibit non-linearities and stochastic properties, there are still linear, deterministic characteristics that can be exploited using linear control methods. For small perturbations, the neural response may be linear enough to be amenable to optimal control approaches (Freeman et al., 2010). However, LFP has been shown to have strong linear properties (Buzsaki, 2006). It has also been shown that linear features of LFP signals have more predication power of seizure activity than non-linear features (Mormann et al., 2005; Jerger et al., 2003). Nevertheless, we are not aware anyone who has tested closed-loop optimal control to suppress spontaneous neural

activity with electrical stimulation in the literature.

To achieve good suppression of spontaneous neural activity, it is necessary to have a stimulation site that causes widespread neural activation. Suppression of seizures have been demonstrated with high frequency electrical stimulation applied to the anterior nucleus of the thalamus (Fisher et al., 2010), and low frequency stimulation of the hippocampus (Koubeissi et al., 2013), as well as the ventral-hippocampal commissure (VHC) (Rashid et al., 2012). Because VHC innervates a large portion of the hippocampus, giving large spatial control with a small stimulation target, we chose this target for controlling neural activity in the hippocampus. Additional explanation and motivation for using the VHC is described in Chapter 2 sections 2.2.1.2 and 2.4.2.

In this study, we will show that LFP biomarkers can be suppressed using electrical stimulation using a Linear Quadratic Gaussian (LQG) controller. The LQG controller consists of a Kalman Filter to estimate state and a Linear Quadratic Regulator (LQR) for feedback control. In the control scheme (Figure 32), the KF assimilates LFP prediction from the model with LFP data from the recording electrode. The LQR determines the stimulation input required to suppress the activity with minimum energy.

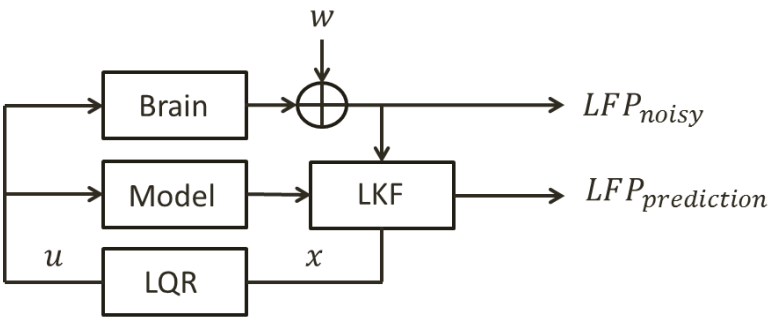


Figure 32. Optimal Control signal flow diagram. The Linear Kalman Filter (LKF) predicts the true LFP ($LFP_{prediction}$) by combining Noisy LFP ($LFP_{noisy}$) data from the rat brain and the Model predictions. Estimates of the system states ($x$) are sent to the Linear Quadratic Regulator (LQR) from the LKF to determine the appropriate stimulation input. Deep Brain Stimulation ($u$) is applied to the VHC which excites the hippocampus. The stimulation input is also provided as input to the state space model.

## *5.2 Methods*

### 5.2.1 Experimental Methods

Surgical methods have been described in Chapter 2 section 2.2.1.2. We use an open source real-time Linux-based platform, the Real-Time eXperimental Interface (RTXI) (http://rtxi.org) that is very flexible for developing closed-loop experiments. We stimulated the ventral hippocampal commissure (Figure 34) with Gaussian white noise (GWN) using a concentric stimulation electrode while recording LFP activity in the hippocampus. The GWN was generated using an Ornstein-Uhlenbeck process with theta (decay time constant) set to half of the sampling rate. Continuous voltage outputs from RTXI (described in previous chapter) were sent to an A-M systems Model 2200 stimulation isolation unit to convert the voltage to a current output. process with theta (decay time constant) set to half of the sampling rate. Continuous voltage outputs from

RTXI (described in previous chapter) were sent to an A-M systems Model 2200 stimulation isolation unit to convert the voltage to a current output.

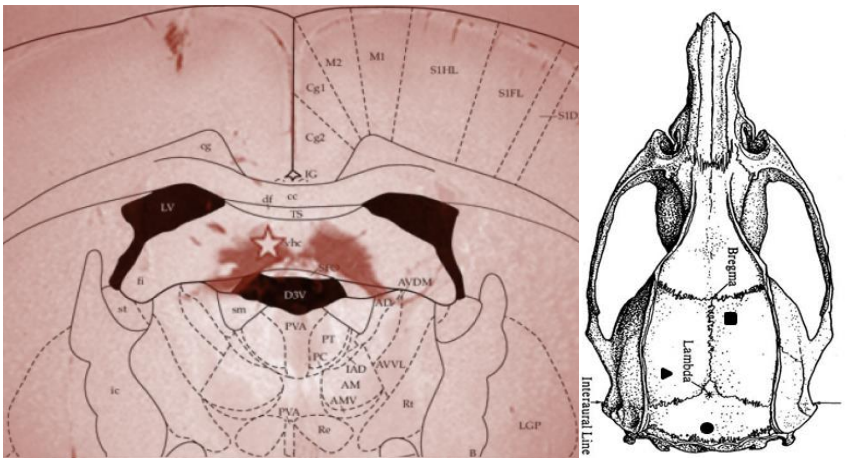


Figure 33. Ventral Hippocampal Commissure as the deep brain target for electrical stimulation. Left, rat brain atlas overlaid with coronal section showing placement of stimulation electrode (Microprobes, Inc.) into the VHC (Star, dark red region). Right, diagram showing electrode locations: Stimulation electrode (square), recording electrode (triangle), and ground (circle).

### 5.2.2 Black-box system identification

#### *5.2.2.1 Ho-Kalman Subspace System Identification algorithm*

A dynamical systems model of LFP activity was determined using the *Ho-Kalman* algorithm (Miller and de Callafon, 2012). This algorithm calculates model coefficients using the evoked potential of the system. The details of the algorithm can be found in Appendix C.

A dynamical system can be represented in a state-space form as follows:

$$x(t+1) = Ax(t) + Bu(t) + w(t) \quad (5.1)$$

$$y(t) = Cx(t) + Du(t) + v(t)$$

where $x(t) \in R^n$ is a vector of dimension $n$ that contains all unobservable states of the system. The $A$ matrix is the state transition matrix which contains the dynamics of the system, the $B$ matrix transforms determines how inputs affect the next state, the $C$ matrix determines which unobservable states are ultimately observed, and lastly the $D$ matrix is a constant gain term. All coefficients are determined by the Ho-Kalman algorithm. The output of the system $y(t)$ is the model prediction of the LFP activity. Additive process noise $(w)$ and measurement noise $(v)$ are modeled as a Gaussian white noise. Figure 34 shows an example of the impulse response from a model determined in one of the experiments. Low order models produce terrible impulse response waveforms that do not reflect the dynamics of the evoked response. Increasing the model order produces impulse response waveforms that capture the dynamics of the evoked potential.

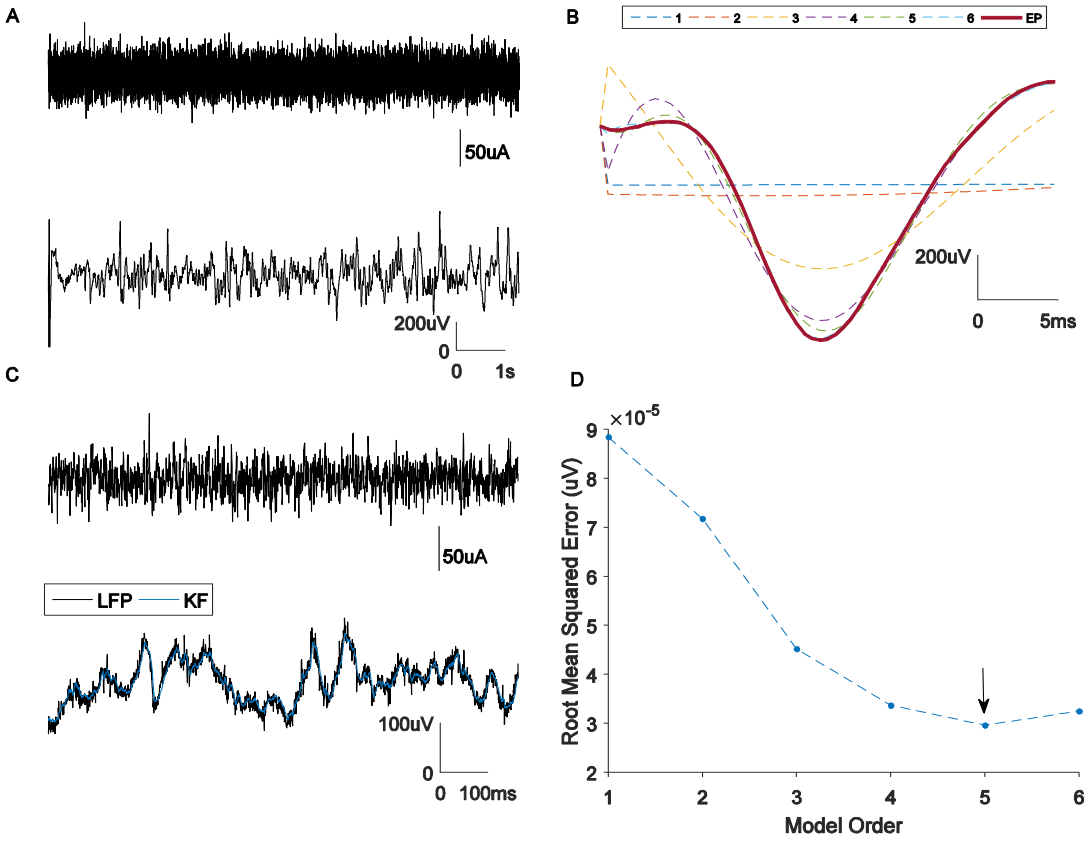


Figure 34. Local Field Potential model identification. A, (top), stimulation to the VHC using GWN. A, (bottom), LFP activity during stimulation. B, Impulse response fits to the LFP Evoked Response using Ho-Kalman SID method on a training data set. The dark red trace is the evoked response recorded in the hippocampus following stimulation. There is a strong depolarization that occurs following 5 ms, and the response peak is at approximately ~13ms. Increasing model order results in better impulse response fits to the evoked response. C Kalman Filter predictions on cross-validation data set. Top, stimulation input to the VHC through a concentric bipolar electrode. Bottom, Kalman Filter prediction follows the LFP smoothly throughout the length of the data set. D, RMSE values for six different models realized using Ho-Kalman SID algorithm in one experimental subject. In this example, the 5th order model has the lowest RMSE value.

### 5.2.3 Model cross-validation

We used two qualitative and one quantitative measure to determine the best model for each experiment. Qualitative measures for the Ho-Kalman model identification

experiments included assessing impulse response fit to the evoked response (Figure 34, B) and the Kalman Filter prediction on the cross validation data (Figure 34, C). The best model would have an impulse response waveform similar to the evoked response and a Kalman Filter prediction that was smoothly following the raw LFP time-series.

We used the root mean squared error (RMSE) as our quantitative metric to determine the best model for the closed-loop control experiments (Figure 34, D). The RMSE was calculated as follows:

$$RMSE = \sqrt{\frac{(y - LFP_{noisy})^2}{n}} \quad (5.2)$$

where $n$ is the total length of the data set, $y$ is the model prediction and $LFP_{noisy}$ is the raw data. Low RMSE values correspond to model predictions that are more accurate. The model with the lowest RMSE value was used in all experiments and for the design of the Regulator.

### 5.2.4 Linear Quadratic Regulator Design

#### *5.2.4.1 Controller Design*

We designed the Regulator to minimize LFP activity while minimizing stimulation energy. The cost function was modeled as a quadratic function:

$$J(u) = z^T Q z + u^T R u \quad (5.3)$$

$$z = y - y_d \quad (5.4)$$

Where $z$ is the difference between the model prediction ($y$) and the desired control target

$(y_d)$, Q is the state weight matrix that determines how much to penalize state values, $u$ is the stimulation input, and $R$, cost for stimulation, is the input weight that determines how much to penalize stimulation. Higher $R$ weights minimizes the stimulation amplitude from the Regulator, and lower $R$ values increases the stimulation amplitude. The cost for stimulation varied for each experimental subject because of variability in the electrode location and electrode durability.

The feedback gain matrix for the Regulator is determined using dynamic programming. We want to find the minimum stimulation that overall minimizes the cost function $J(u)$. First we define a value function using the cost function above that accounts for the cost incurred and the minimum cost-to-go for the LQR.

$$V(z) = \min_{u_0,..u_{N-1}} \sum_{\tau=0}^{N} (z_\tau^T Q z_\tau + u_\tau^T R u_\tau) \quad (5.5)$$

Once we have our value function, we use the Matlab Discrete Algebraic Riccati Equation function *dare* to solve for the Regulator feedback gain matrix. In summary, the Regulator takes in state information from the LKF and determines the appropriate stimulation using the feedback gain matrix. A full description for designing a LQG controller can be found in (Bertsekas, 1995)

#### *5.2.4.2 Biomarker*

Previous experiments have indicated that 4-AP induced seizures can have broadband LFP profiles. For these experiments, we bandpass filtered the raw LFP between 1-100Hz. Bandpass filtering for the LFP also improved model identification using the Ho-Kalman algorithm because the evoked potential from the filtered data had a

smoother profile than unfiltered raw LFP data. Furthermore, the 1Hz highpass cutoff eliminated all drifting in the evoked response that could make it difficult to fit a linear model.

Power spectral densities (PSDs) of the raw LFP data were calculated using Matlab's Welch estimation method *pwelch* ($Window = 2^{12}\ and\ overlap = \frac{1}{2}Window$). In order to propagate errors to determine confidence intervals (CIs) for the percent modulation metric, we used a square root transform on the PSD values. This method changes the distribution of the PSD estimates for each frequency from a Chi-squared distribution to a Gaussian distribution. The standard error of the mean for each frequency can be calculated using the multiple PSD estimates. We then use the Gaussian propagation equations detailed in (Taylor, 1997) to construct the CIs for percent modulation. Our main metric for controller efficacy was percent modulation of the local field potential which is mathematically formulated as follows:

$$PM_{StoB} = \frac{Pxx_S - Pxx_B}{Pxx_B} * 100, \text{ and } PM_{BtoP} = \frac{Pxx_B - Pxx_P}{Pxx_B} * 100 \quad (5.6)$$

where $PM$ is the percent modulation, $S$ is the stimulation epoch, $B$is the baseline or pre-stimulation epoch, $P$ is the post-stimulation epoch, and $Pxx$ is the PSD for the signal in a given epoch.

## *5.3 Results*

### 5.3.1 Stimulation via LQG control leads to broad-band suppression of LFP activity

To test the efficacy of the LQG controller we compared the power of the LFP before and during stimulation. Our approach was as follows: 1. Baseline: LFP recordings in the hippocampus, 2. Stimulation: continuous stimulation determined by the regulator was applied through a concentric electrode, and 3. Post-stimulation recording to determine whether LFP power recovered to pre-stimulation levels. The duration of each phase was two minutes. Figure 35 (left) shows and example of an experiment using closed-loop stimulation. In this example, the LFP activity is much stronger before stimulation (black trace) than during stimulation (red trace) as indicated by the strong bursts. Notice that during closed-loop stimulation the input (grey trace) closely reflects the trajectory of the LFP (red trace); however, this similarity is missing when stimulation is applied in open-loop (blue trace).

Following stimulation, the number of bursts in the LFP increases towards baseline levels as indicated in the PSD (Figure 35, right). Our results indicate broadband power suppression in the closed-loop experiment, and replaying the same stimulus in open-loop did not suppress LFP activity (Figure 35, right). Comparison between baseline and post-stimulation epochs shows little difference in the broad band PSDs (Figure 35, right).

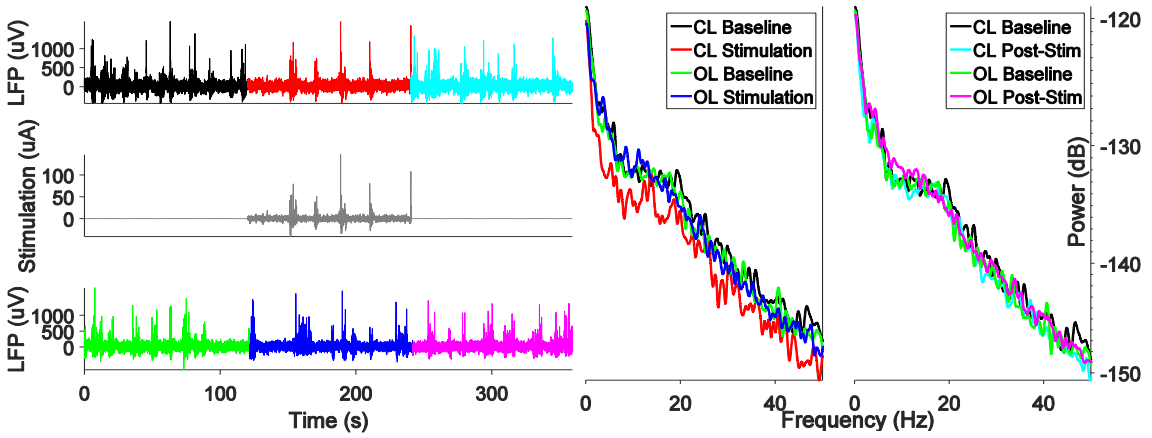


Figure 35. Experimental approach to testing model based closed-loop control. Left, LFP recorded from the CA1 region of the hippocampus (top) during closed-loop stimulation (middle) and LFP from open-loop stimulation (bottom). Right, PSD estimates comparing stimulation and baseline LFP power and baseline with post-stimulation power. Comparisons are done for closed-loop and open-loop stimulation. The different colors correspond to different experimental phases.

### 5.3.2 LQG performance: Responders vs Non-Responders

We tested our method on ten animals; however, the results from only six are presented in this chapter in order to compare the results across experiments using the same system identification algorithm (Ho-Kalman). For the other four animals we used another popular state-space system identification method called *n4sid*. Although we used the same control objective across animals, the performance of the LQG controller was not the same for each experiment. Our results from six animals show performance can be differentiated between Responders and Non-Responders (Figure 36).

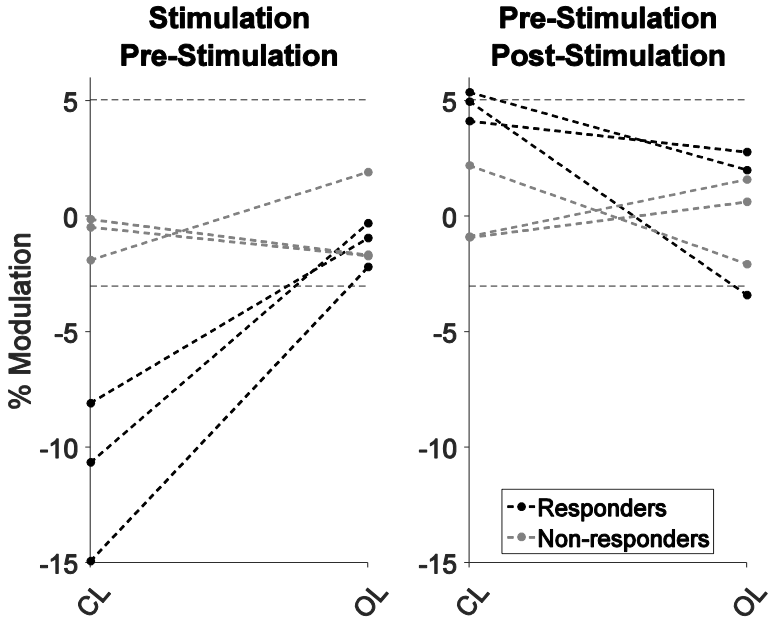


Figure 36. Performance of LQG controller are grouped into two categories: Responders and Non-Responders. Left, percent modulation of the LFP power between the stimulation and pre-stimulation (baseline) experimental phases. Right, comparison between the pre-stimulation and post-stimulation experimental phases. Data is categorized by closed-loop (CL) and open-loop (OL) experiments. Dotted lines correspond to the null hypothesis distribution at two standard deviations.

In three of our experiments, we found strong suppression of the LFP during the stimulation epoch using closed-loop stimulation compared to baseline (Figure 36, left). The Non-Responders on the other had little to no suppression of LFP during closed-loop stimulation (Figure 36, left). In all experiments, open-loop stimulation had negligible effects. We also compared the LFP power before and after stimulation to determine whether suppression of LFP persists beyond the stimulation period. Modulation at 0% would indicate that the LFP had fully recovered. In the Responder group, the percent modulation between pre-stimulation and post-stimulation during closed-loop tended to be higher than 0% (Figure 36, right). The effects were not as pronounced in the open-loop

stimulation case. The Non-Responders did not have any discernable trend in LFP modulation between closed-loop and open-loop stimulation. To determine whether the modulation was significant we determined the distribution spread for our null hypothesis experiments which included the open-loop comparisons between stimulation and baseline as well as baseline and post-stimulation. Our data indicates the Non-Responder performance metrics indistinguishable from open-lop stimulation and only the Responders had significant LFP modulation (Figure 36, left).

### 5.3.3 Model accuracy on out of sample data and controller efficacy

To determine what distinguished Responders and Non-Responders we investigated prediction accuracy of the models (Figure 37). Essentially, we are asking how well the model can predict future outputs using past state values and the current input. Good models will have higher prediction accuracies at longer prediction horizons. The three responders had strong prediction accuracy out to 3.5ms prediction horizon. At the 7ms horizon one of the models still had strong prediction accuracy. All models performed poorly at the 200ms prediction horizon.

What is interesting in this data is that the prediction accuracy decay rate as a function of prediction horizon time is not the same for each model. For example, the model with the highest prediction accuracy in the responder group at the one sample point prediction horizon decayed faster than the model with the lowest one sample point prediction accuracy. In the following sections, we will determine the relationship between the magnitude of suppression and prediction accuracy.

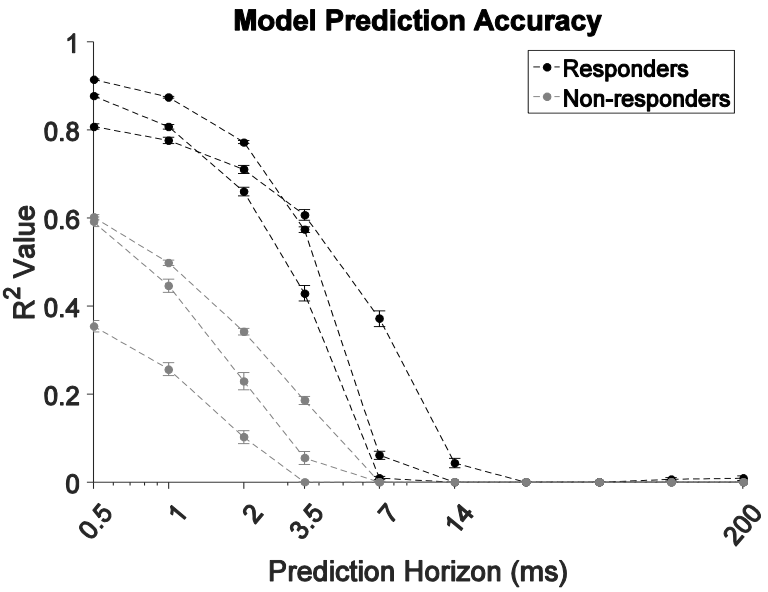


Figure 37. Accuracy of different models across different prediction horizons classified by responders and non-responders. Each data point is the mean prediction accuracy across multiple segments of cross validation data. The bars correspond to the standard error of the mean. Methods for calculating the prediction horizon are in 6.4Appendix C.

### 5.3.4 Broadband power in the LFP prior to stimulation and controller efficacy

Another predictive factor of the magnitude of LFP modulation was the broadband power in the LFP prior to stimulation. The responder group had stronger spontaneous neural activity with bursting activity, as seen in Figure 35, left. The bursting activity has a strong broadband spectral profile and contributed to the increase in the PSD (Figure 36). It is unclear why bursting was present in the responder group experiments and not the non-responder group. Since the experimental objective is to minimize LFP activity, there needs to be sufficient neural activity (i.e. LFP bursts) in order to see significant effects. Our data clearly shows that experiments with low burst frequency prior to baseline had

no significant effect in minimizing LFP activity.

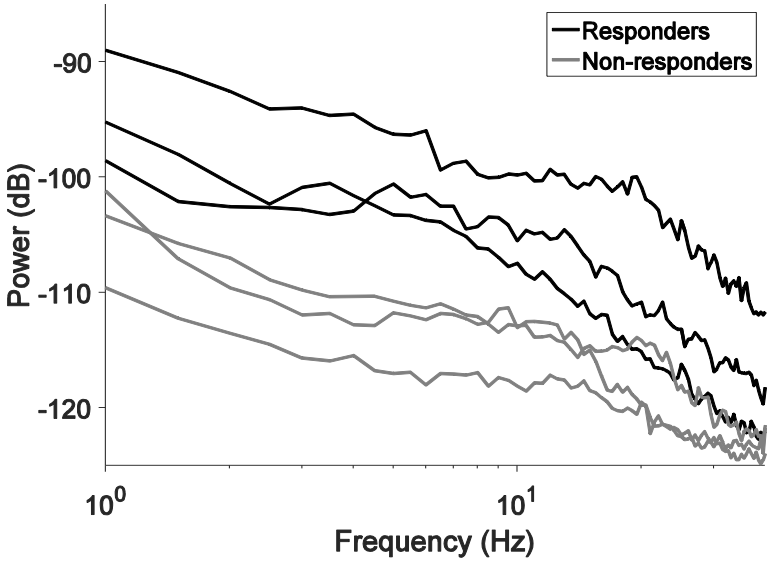


Figure 38. LFP power in Responder groups was much higher than in Non-Responder groups.

### 5.3.5 Robust models confer stronger LFP modulation

Our results showed both high frequency of bursts prior to stimulation and good prediction accuracy at 3.5ms prediction horizon was required to suppress LFP successfully (Figure 39). This makes sense, as a high model $R^2$ value indicates good state estimates which in turn determine the appropriate inputs to minimize the value function. Closed-loop stimulation in the responder group was significantly greater than both open-loop stimulation and closed-loop stimulation in the non-responder groups.

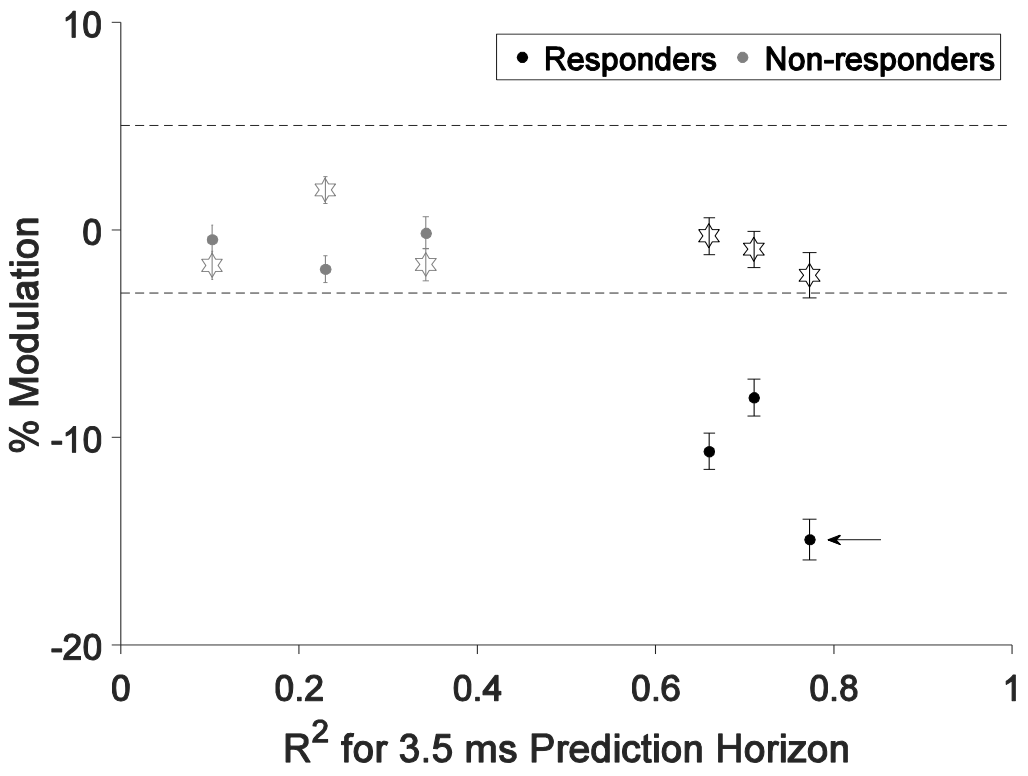


Figure 39. High predictability determines extent of LFP suppression. Circles indicate closed-loop experiments, and the stars are for open-loop experiments. The bars correspond to 95% confidence intervals. The black arrow indicates the percent modulation for the experiment example shown in Figure 35. Dotted lines correspond to the null hypothesis distribution at two standard deviations.

Next, we wanted to ascertain what prediction horizon was most correlated with suppression. Regression of the suppression with prediction accuracy at different prediction horizons is shown in Table 2. At all prediction horizons, there is a negative correlation between prediction accuracy and % modulation. The greatest correlation between prediction and modulation amplitude was at a prediction horizon of 1 ms (Figure 37).

Table 2: Model accuracy across prediction horizons.

| **Prediction Horizons (ms)** | ***0.5*** | ***1*** | ***2*** | ***3.5*** | ***7*** | ***14*** |
|---|---|---|---|---|---|---|
| $R^2$**value** | 0.81 | 0.85 | 0.83 | 0.75 | 0.09 | 0.03 |
| $p - value$ | **0.03** | **0.02** | **0.02** | **0.04** | 0.58 | 0.76 |

## *5.4 Discussion*

Our results show that model based feedback controller can use state dependent stimulation to efficiently suppress LFP activity. When the stimulation is applied in open-loop the effects of stimulation on LFP suppression are minimized. Our data indicates that the LQG controller is highly effective when the 'black-box' model is accurate at multiple prediction horizons and there is sufficient background LFP activity to suppress. In experiments where the LQG was unsuccessful at suppressing LFP activity we found that there was a lack of LFP burst of activity and/or model prediction accuracy was weak. In negative control experiments, we applied the state dependent stimulation in an open-loop experiment. We found significantly different within subject effects between closed-loop and open-loop stimulation in four different experiments. Our results provide evidence that model based suppression of LFP activity is possible, and that Optimal Control algorithms could be used to design patient/subject specific therapies.

### 5.4.1 Model Design

Our regression analysis indicated that model prediction horizon had strong effect on the

efficacy of the Regulator. Robust suppression of LFP activity requires accurate state estimation. There are many 'black-box' algorithms could potentially be used for LFP data. We decided to go with Ho-Kalman system identification method because the state space model representation is required for designing the Kalman Filter and Regulator. Furthermore, the estimation of the state-space model using the Ho-Kalman algorithm measures correlations across longer time scales than typical difference equation methods like auto-regressive exogenous (ARX) models.

In the responder group, we showed that the model accurately predicted the LFP at multiple time horizons, but in the non-responder group, prediction accuracy dropped quickly as the time horizon increased. One reason for the difference in predictability may be that the system identification data was highly non-linear which resulted in the design of linear state space models unable to capture dynamics of the evoked response. Non-linearities in the data sets tend to arise when the stimulation amplitude is too high. In order to avoid this, lower stimulation amplitudes need to be used; however, if the stimulation amplitude is too low then a evoked response will not occur.

This work also tested different system identification algorithms to estimate state-space systems. In three other experiments, we used an algorithm called "n4sid" and in one experiment the auto-regressive exogenous input (ARX) algorithm to design state-space models of LFP activity. Although the results from these experiments were inconclusive due to low sample size, these methods could design models with good prediction accuracies at multiple time horizons.

Further analysis is needed for the combined white noise LFP data set. There are a many different system identification approaches that can be used, and we have not yet scratched the surface to determine the appropriate method for models for feedback controllers. Furthermore, the use of adaptive models would be preferred to static models generated during system identification. An iterative system identification approach, where model coefficients are updated continuously or at specific intervals would be able to account for non-stationarities in the neural signal over long recordings. Also, use of an adaptive model scheme removes the need for having to do system identification prior to experimentation.

Yet another approach would be to use nonlinear system identification algorithms. These methods implement a non-linear transformation to the inputs prior to calculating the output variable. Lastly, adaptive models could be used to account for non-stationary conditions in the LFP. In addition to state variables corresponding to the LFP, the Kalman Filter can estimate optimal model parameters. This approach would increase computational cost, but would increase prediction accuracy.

### 5.4.2 Regulator Design

We designed the Linear Quadratic Regulator with two free parameters: tracking target and stimulation cost. For all experiments the tracking target was set to zero, which means the control policy will force the state variables to have trajectories with a mean of zero. In other words, the regulator does not introduce DC offsets. The stimulation cost differed across experimental subjects. This was predominantly due to the physical location of the

stimulation electrode. The VHC is a bundle of axons that transverses the midline. The optimal position of the stimulation electrode would be in the middle of the VHC for maximum activation. The stimulation amplitude threshold for generating an evoked response was inversely proportional to the position of the electrode with respect to the VHC. Smaller amplitudes were required when the electrode position was in the optimal location.

The stimulation cost was chosen in a way so that the stimulation magnitude of the regulator during a simulation run would be comparable to the stimulation magnitude of the GWN during system identification. Prior to executing the experiments, we calculated the eigenvalues of the full closed-loop system (including the feedback gain matrix). Eigenvalues greater than or equal to one indicated that the system was unstable. In these instances, we decreased the stimulation cost which would bring the full system into a stable region.

### 5.4.3 Control Policy

In all experiments, we set the control policy to minimize the difference between the target output and the model prediction and minimize the stimulation amplitude. The system states reflect the raw signal, which in all experiments was bandpass filtered (1-100Hz) LFP data. The objective in our experiments was to access the efficacy of the LGQ controller to minimize broadband filtered LFP data. We discovered that the likelihood of suppressing LFP activity was proportional to the frequency of baseline bursting activity. High frequency bursts would increase the broadband LFP frequency profile, giving the

controller something to suppress. In experiments with few bursts we saw limited suppression of LFP activity. Our data indicates that there is a flooring effect that prevents the controller from suppressing LFP activity if there are a limited number of bursts.

# Chapter 6 Conclusions

The goal of this dissertation was to develop novel patient/subject closed-loop optimization algorithms for electrical stimulation with applications to seizure disorders. There are three fundamental approaches to optimizing electrical stimulation: 1) Determining optimal brain regions to stimulate for different seizure disorders, 2) Optimizing electrode performance, and 3) Development of personalized, precision electrical brain stimulation using closed-loop algorithms. The dissertation is focused on the first and third approaches, and I will discuss implication and future directions of this work.

## *6.1 Optimal stimulation targets*

There are three types of brain regions that can be stimulated for seizure control. First, modulatory nuclei which has been shown to minimize seizure activity (Fisher et al., 2010). Second, stimulation of the seizure focus using on-demand stimulation which is currently the state-of-the art for seizure control (Morrell and Group, 2011). Third, stimulation of fiber tracts which has shown to be highly effective in suppressing seizures in rodents (Rashid et al., 2012) and humans (Koubeissi et al., 2013).

We tested whether open-loop stimulation of the Endopiriform Nucleus (ENu) could suppress seizures. The motivation for testing this region was three-fold: first, it has been shown to be the site for driving seizure activity in the piriform cortex; second, the ENu sends afferents to the entorhinal cortex which is part of the gross hippocampal network responsible for seizure in patients with mesial temporal lobe epilepsy; third, it's location is in a low-surgical risk area, so intervention may be preferable to other regions. Our

results showed that neither low nor high frequency stimulation could suppress focal hippocampal seizures in an acute seizure model. It may be the case that modulation of the ENc for seizure control requires prolonged stimulation over many days if not weeks. This hypothesis was supported by preliminary data presented at the Annual American Epilepsy Foundation in 2015. Their data indicated that stimulation of the deep Piriform Cortex over two weeks within a small cohort of rodents resulted in seizure suppression.

Following these results, we shifted our deep brain target for seizure control to the Ventral Hippocampal Commissure (VHC) which had been shown to be a good stimulation region for seizure control (Rashid et al., 2012). Our first step was to reproduce the results in previous studies. High frequency stimulation in an acute focal hippocampal seizure model was very effective in suppressing seizures. We decided to use the VHC – Hippocampal network as an experimental platform to test closed-loop algorithms for two reasons: first, the VHC strongly modulates the hippocampus proper, and second, electrode placement in the VHC is not difficult and is highly reproducible. Minimizing variability in surgical outcome prior to the seizure control experiments was imperative so that we could accurately access the efficacy of complex closed-loop algorithms.

Additionally, preliminary results from (Rashid et al., 2012) and (Koubeissi et al., 2013) indicate that effective stimulation protocols tested in animals can also provide some therapy in humans.

## *6.2 Stimulation artifact removal for closed-loop electrical brain stimulation*

Real-time biomarker detection is predicated on the ability to extract features from clean neural signals. In closed-loop stimulation paradigms, the act of modulating the biomarker using stimulation introduces artifacts that corrupt the neural signal. It is vital to have robust filters that can remove artifacts while minimizing artifacts from the filtering process itself. Implemented correctly, these filters can clean corrupted neural data so that closed-loop algorithms can adjust stimulation parameter settings as a function of the biomarker of interest. We tested the performance of five different stimulation artifact filters in a computational model and then verified our results *in vivo*.

We used a "wobble oscillator" computational model to generate simulated local field potential (LFP) data with stimulus artifacts, evoked responses, and measurement noise. The root mean squared error (RMSE) between the noisy signal with the stimulation artifact and the true signal was used as a measure of filter performance. Our data shows that the Kalman and Adaptive FIR filters outperformed other approaches across all stimulation paradigms. However, the Kalman filter is computationally more intensive than the other methods by several orders of magnitude. Ultimately, the choice of which filter to use is dependent on the stimulation paradigm. For instance, if the experimenter or clinician is using periodic stimulation and has high confidence that the stimulation artifact will not change, then a simple template filter would be the recommended choice because is it simple and computationally efficient. On the other hand, if the experimenter

or clinician is using a sophisticated closed-loop algorithm with non-periodic stimulation, an Adaptive FIR or Kalman filter would be more appropriate.

Next, we tested five different digital stimulation artifact removal filters *in vivo*. Similar to the simulations studies, the performance of Adaptive FIR and Kalman filters exceeded that of the Blanking, Comb and Template filters. We also tested the efficacy of a bi-directional neural interface with a built in Adaptive FIR filter in collaboration with a group at the University of Michigan (Mendrela et al., 2016). The hardware successfully removed stimulation artifacts from high and low frequency stimulation pulse trains. We also showed that the Adaptive FIR could remove stimulation artifacts during seizures without corrupting the underlying neural signal.

Continuous monitoring of neural biomarkers is imperative for closed-loop therapies that administer therapy as a function of the magnitude of the biomarker. Without robust removal of stimulation artifacts, closed-loop algorithms may not converge to physiologically relevant stimulation solutions. There are many different filters experimentalists or clinicians can use for closed-loop therapies. The choice of filter largely will depend on the complexity of the therapy and hardware limitations of the implantable device.

## *6.3 Determining optimal stimulation parameters for seizure control using reinforcement learning*

One method to optimize stimulation parameters for seizure control is to use a reinforcement learning (RL) algorithm that iteratively tests different parameter

combination to maximize reward. In our simulation experiments we set the reward to be inversely proportional to the magnitude of epileptogenic activity. The algorithm tries to find the best stimulation frequency for each state that maximizes the expected reward. We established a two-dimensional state-space using simple linear filters that separated the seizing state from the non-seizing state. The state space generated from simple linear filters closely resembles that of the original state space of the unobservable state variables.

Our simulation results showed that the reinforcement learning algorithm can find the optimal stimulation frequency under different experimental conditions. The discrepancy between the minimum stimulation frequency determined analytically through dynamical systems analysis and the optimal stimulation frequency can be explained by the objective of the RL algorithm, which is to maximize reward not suppress seizures. Since the Epileptor LFP signal contains various types of epileptogenic activity including seizures, we tuned the RL so that it removes all epileptogenic activity.

The time constant for the Softmax action selection policy allows for tuning of the RL algorithm. A low Softmax value results in a greedy approach while a high SoftMax value forces the RL algorithm to explore the parameter space leading to a process of optimizing a distribution of stimulation frequencies. This tuning parameter could be very helpful in a clinical setting to determine the optimal stimulation frequency that suppresses an epileptogenic biomarker. First, the clinician can set the temperature to a high value to explore the parameter space while allowing a distribution of frequencies to be optimized.

Only a subset of the larger frequency parameter space would have a high probability of being chosen. Then, the clinician could lower the SoftMax and have the RL algorithm learn the optimal stimulation frequency over a smaller parameter space. This is just one example of how RL algorithms can be used in a clinic to help patients with seizure disorders.

## *6.4 Rejecting spontaneous neural activity through model based feedback control*

Another approach to personalize electrical brain stimulation is to use model based feedback control to suppress neural activity. Here an empirical model is fit to LFP response to electrical stimulation to make forward predictions. These predictions are assimilated with raw data using a Kalman Filter to estimate the neural state. The state estimates are then used by a Regulator to determine the appropriate feedback stimulation to suppress the neural activity. This approach generates state dependent arbitrary waveforms in closed-loop.

Multiple milestones were accomplished in this project. First, to the best of our knowledge, this is the first time a model-based feedback controller was used to modulate LFP activity using electrical stimulation. Our results clearly indicated that state-dependent stimulation is required to suppress LFP activity, and that the suppression in LFP was not due to stimulation energy. Another milestone from this project was validation of the predictive power of linear state space models. We used the Ho-Kalman system identification algorithm to estimate linear state-space models on six separate

experiments. In the responder group, we showed that the model accurately predicted the LFP at multiple time horizons clearly showing that linear models can capture most of the variance seen in LFP signals.

Further analysis is needed for the combined white noise LFP data set. There are a many different system identification approaches that can be used, and we have not yet scratched the surface to determine the appropriate method for models for feedback controllers. Furthermore, the use of adaptive models would be preferred to static models generated during system identification. An iterative system identification approach, like that of the Adaptive FIR stimulation artifact filter in Chapter 3, would be able to account for non-stationarities in the neural signal over long recordings. Also, use of an adaptive model scheme removes the need for having to do system identification prior to experimentation.

Lastly, this approach will need to be tested in seizing animals. We hypothesize that the LQR will drive the VHC with sufficient stimulation to minimize epileptogenic activity. The state space model would need to be trained on baseline inter-ictal data. The control policy would be designed to suppress large deviations from baseline activity, therefore suppressing fast oscillations and spike and wave events that occur prior to tonic-clonic seizure onset. This approach is different than on-demand approaches popularized by NeuroPace (Morrell and Group, 2011) and Optogenetic studies (Krook-Magnuson et al., 2013) because the objective is not to suppress seizures, but to prevent seizures from occurring by increasing stimulation during ictogenesis. This should in

principle send the system back to an inter-ictal state.

It may be possible for Optimal Control algorithms, like the one tested here, to be implemented within an implantable device with sufficient hardware architecture. Model based controllers are used routinely in many engineering applications. If we can estimate a robust LFP model, then in principle we should be able to design controllers to any specification as determined by a clinician. One could envision a clinician tuning the feedback controller for a programmable stimulator during a routine clinic visit. The tuning process would first validate that the LFP model is making accurate predictions, and second establish that the full closed-loop system is stable. While there are still many unanswered questions regarding the long-term efficacy of Optimal Control neuromodulation therapies, there is a possibility that these technologies will become mainstream in the near future.

# References

Abegg MH, Savic N, Ehrengruber MU, McKinney RA, Gähwiler BH (2004) Epileptiform activity in rat hippocampus strengthens excitatory synapses. J Physiol 554:439-448.

Agnesi F, Connolly AT, Baker KB, Vitek JL, Johnson MD (2013) Deep brain stimulation imposes complex informational lesions. PLoS One 8:e74462.

Armstrong C, Krook-Magnuson E, Oijala M, Soltesz I (2013) Closed-loop optogenetic intervention in mice. Nat Protoc 8:1475-1493.

Avoli M, Louvel J, Pumain R, Köhling R (2005) Cellular and molecular mechanisms of epilepsy in the human brain. Prog Neurobiol 77:166-200.

Behan M, Haberly LB (1999) Intrinsic and efferent connections of the endopiriform nucleus in rat. J Comp Neurol 408:532-548.

Berenyi A, Belluscio M, Mao D, Buzsaki G (2012) Closed-loop control of epilepsy by transcranial electrical stimulation. Science (New York, NY) 337:735.

Bertsekas DP (1995) Dynamic programming and optimal control: Athena Scientific Belmont, MA.

Bikson M, Lian J, Hahn PJ, Stacey WC, Sciortino C, Durand DM (2001) Suppression of epileptiform activity by high frequency sinusoidal fields in rat hippocampal slices. J Physiol 531:181-191.

Bower MR, Buckmaster PS (2008) Changes in granule cell firing rates precede locally recorded spontaneous seizures by minutes in an animal model of temporal lobe epilepsy. J Neurophysiol 99:2431-2442.

Bower MR, Stead M, Meyer FB, Marsh WR, Worrell GA (2012) Spatiotemporal neuronal correlates of seizure generation in focal epilepsy. Epilepsia 53:807-816.

Brocker DT, Swan BD, So RQ, Turner DA, Gross RE, Grill WM (2017) Optimized temporal pattern of brain stimulation designed by computational evolution. Sci Transl Med 9.

Bush K, Pineau J (2009) Manifold embeddings for model-based reinforcement learning under partial observability. *Advances in neural information processing systems*:pp. 189-197.

Buzsaki G (2006) Rhythms of the Brain: Oxford University Press.

Cagnan H, Pedrosa D, Little S, Pogosyan A, Cheeran B, Aziz T, Green A, Fitzgerald J, Foltynie T, Limousin P, Zrinzo L, Hariz M, Friston KJ, Denison T, Brown P (2017) Stimulating at the right time: phase-specific deep brain stimulation. Brain 140:132-145.

Chabardès S, Kahane P, Minotti L, Koudsie A, Hirsch E, Benabid AL (2002) Deep brain stimulation in epilepsy with particular reference to the subthalamic nucleus. Epileptic Disord 4 Suppl 3:S83-93.

CJL M, AD L (1994) Global comparative assessments in the health sector : disease burden, expenditures and intervention packages. Geneva : World Health Organization 196.

Debanne D, Thompson SM, Gähwiler BH (2006) A brief period of epileptiform activity strengthens excitatory synapses in the rat hippocampus in vitro. Epilepsia 47:247-256.

Demont-Guignard S, Benquet P, Gerber U, Biraben A, Martin B, Wendling F (2012) Distinct hyperexcitability mechanisms underlie fast ripples and epileptic spikes. Ann Neurol 71:342-352.

Edwardson MA, Lucas TH, Carey JR, Fetz EE (2013) New modalities of brain stimulation for

stroke rehabilitation. Exp Brain Res 224:335-358.

Einevoll GT, Franke F, Hagen E, Pouzat C, Harris KD (2012) Towards reliable spike-train recordings from thousands of neurons with multielectrodes. Curr Opin Neurobiol 22:11-17.

Fisher R et al. (2010) Electrical stimulation of the anterior nucleus of thalamus for treatment of refractory epilepsy. Epilepsia 51:899.

Fisher RS (2013) Deep brain stimulation for epilepsy. Handb Clin Neurol 116:217-234.

Frasca A, Aalbers M, Frigerio F, Fiordaliso F, Salio M, Gobbi M, Cagnotto A, Gardoni F, Battaglia GS, Hoogland G, Di Luca M, Vezzani A (2011) Misplaced NMDA receptors in epileptogenesis contribute to excitotoxicity. Neurobiol Dis 43:507-515.

Freeman DK, Rizzo JF, Fried SI (2010) Electric stimulation with sinusoids and white noise for neural prostheses. Front Neurosci 4:28.

Frei MG, Osorio I, Rise MT, Giftakis JE, Graves NM, Carlson DL, Jensen RM (2007) Channel-selective blanking for a medical device system. In: Google Patents.

Friedland B (2012) Control system design: an introduction to state-space methods: Courier Corporation.

Gavrilovici C, Pollock E, Everest M, Poulter MO (2012) The loss of interneuron functional diversity in the piriform cortex after induction of experimental epilepsy. Neurobiol Dis 48:317-328.

Gluckman BJ, Nguyen H, Weinstein SL, Schiff SJ (2001) Adaptive electric field control of epileptic seizures. The Journal of neuroscience : the official journal of the Society for

Neuroscience 21:590.

Good LB, Sabesan S, Marsh ST, Tsakalis K, Treiman D, Iasemidis L (2009) Control of synchronization of brain dynamics leads to control of epileptic seizures in rodents. International journal of neural systems 19:173.

Handforth A, DeGiorgio CM, Schachter SC, Uthman BM, Naritoku DK, Tecoma ES, Henry TR, Collins SD, Vaughn BV, Gilmartin RC, Labar DR, Morris GL, Salinsky MC, Osorio I, Ristanovic RK, Labiner DM, Jones JC, Murphy JV, Ney GC, Wheless JW (1998) Vagus nerve stimulation therapy for partial-onset seizures: a randomized active-control trial. Neurology 51:48-55.

Hashimoto T, Elder CM, Vitek JL (2002) A template subtraction method for stimulus artifact removal in high-frequency deep brain stimulation. J Neurosci Methods 113:181-186.

Hashizume K, Tanaka T (1998) Multiple subpial transection in kainic acid-induced focal cortical seizure. Epilepsy Res 32:389-399.

Hauser WA, Annegers JF, Kurland LT (1993) Incidence of epilepsy and unprovoked seizures in Rochester, Minnesota: 1935-1984. Epilepsia 34:453-468.

Hayes MH (2009) Statistical digital signal processing and modeling: John Wiley & Sons.

Haykin S, Widrow B (2003) Least-mean-square adaptive filters: John Wiley & Sons.

Heck CN et al. (2014) Two-year seizure reduction in adults with medically intractable partial onset epilepsy treated with responsive neurostimulation: Final results of the RNS System Pivotal trial. Epilepsia 55:432-441.

Hill MW, Wong M, Amarakone A, Rothman SM (2000) Rapid cooling aborts seizure-like

activity in rodent hippocampal-entorhinal slices. Epilepsia 41:1241-1248.

Hoffman WH, Haberly LB (1996) Kindling-induced epileptiform potentials in piriform cortex slices originate in the underlying endopiriform nucleus. J Neurophysiol 76:1430-1438.

Holt AB, Wilson D, Shinn M, Moehlis J, Netoff TI (2016) Phasic Burst Stimulation: A Closed-Loop Approach to Tuning Deep Brain Stimulation Parameters for Parkinson's Disease. PLoS Comput Biol 12:e1005011.

Huang L, van Luijtelaar G (2013) The effects of responsive and scheduled subicular high frequency stimulation in the intra-hippocampal kainic acid seizure model. Epilepsy Res.

Iasemidis LD (2003) Epileptic seizure prediction and control. IEEE transactions on bio-medical engineering 50:549.

Iasemidis LD, Sackellares JC, Zaveri HP, Williams WJ (1990) Phase space topography and the Lyapunov exponent of electrocorticograms in partial seizures. Brain topography 2:187-201.

Jerger KK, Netoff TI, Francis JT, Sauer T, Pecora L, Weinstein SL, Schiff SJ (2003) Comparison of Methods for Seizure Detection. In: Epilepsy as a Dynamic Disease (Milton J, Jung P, eds). New York, NY: Springer.

Jirsa VK, Stacey WC, Quilichini PP, Ivanov AI, Bernard C (2014) On the nature of seizure dynamics. Brain 137:2210-2230.

Johnson MD, Miocinovic S, McIntyre CC, Vitek JL (2008) Mechanisms and targets of deep brain stimulation in movement disorders. Neurotherapeutics 5:294-308.

Kalitzin S, Velis D, Suffczynski P, Parra J, da Silva FL (2005) Electrical brain-stimulation

paradigm for estimating the seizure onset site and the time to ictal transition in temporal lobe epilepsy. Clin Neurophysiol 116:718-728.

Karkar KM, Garcia PA, Bateman LM, Smyth MD, Barbaro NM, Berger M (2002) Focal cooling suppresses spontaneous epileptiform activity without changing the cortical motor threshold. Epilepsia 43:932-935.

Keller CJ, Truccolo W, Gale JT, Eskandar E, Thesen T, Carlson C, Devinsky O, Kuzniecky R, Doyle WK, Madsen JR, Schomer DL, Mehta AD, Brown EN, Hochberg LR, Ulbert I, Halgren E, Cash SS (2010) Heterogeneous neuronal firing patterns during interictal epileptiform discharges in the human cortex. Brain 133:1668-1681.

Kent AR, Grill WM (2012) Recording evoked potentials during deep brain stimulation: development and validation of instrumentation to suppress the stimulus artefact. J Neural Eng 9:036004.

Kirk DE (2012) Optimal control theory: an introduction: Courier Corporation.

Koubeissi MZ, Kahriman E, Syed TU, Miller J, Durand DM (2013) Low-frequency electrical stimulation of a fiber tract in temporal lobe epilepsy. Ann Neurol 74:223-231.

Krack P, Pollak P, Limousin P, Benazzouz A, Benabid AL (1997) Stimulation of subthalamic nucleus alleviates tremor in Parkinson's disease. Lancet 350:1675.

Krook-Magnuson E, Armstrong C, Oijala M, Soltesz I (2013) On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy. Nat Commun 4:1376.

Litt B, Lehnertz K (2002) Seizure prediction and the preseizure period. Current opinion in neurology 15:173.

Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, Foltynie T, Limousin P, Ashkan K, FitzGerald J, Green AL, Aziz TZ, Brown P (2013) Adaptive deep brain stimulation in advanced Parkinson disease. Ann Neurol 74:449-457.

Ljung L (1999) System identification: Wiley Online Library.

Lockman J, Fisher RS, Olson DM (2011) Detection of seizure-like movements using a wrist accelerometer. Epilepsy Behav 20:638-641.

Loddenkemper T, Pan A, Neme S, Baker KB, Rezai AR, Dinner DS, Montgomery EB, Lüders HO (2001) Deep brain stimulation in epilepsy. J Clin Neurophysiol 18:514-532.

Lopantsev V, Both M, Draguhn A (2009) Rapid plasticity at inhibitory and excitatory synapses in the hippocampus induced by ictal epileptiform discharges. Eur J Neurosci 29:1153-1164.

Marmarelis V (2012) Analysis of physiological systems: The white-noise approach: Springer Science & Business Media.

McIntyre CC, Grill WM, Sherman DL, Thakor NV (2004) Cellular effects of deep brain stimulation: model-based analysis of activation and inhibition. J Neurophysiol 91:1457-1469.

McIntyre DC, Plant JR (1989) Pyriform cortex involvement in kindling. Neurosci Biobehav Rev 13:277-280.

Mendrela AE, Cho J, Fredenburg JA, Nagaraj V, Netoff TI, Flynn MP, Yoon E (2016) A Bidirectional Neural Interface Circuit With Active Stimulation Artifact Cancellation and Cross-Channel Common-Mode Noise Suppression. IEEE Journal of Solid-State Circuits 51:955-965.

Millard DC, Wang Q, Gollnick CA, Stanley GB (2013) System identification of the nonlinear dynamics in the thalamocortical circuit in response to patterned thalamic microstimulation in vivo. J Neural Eng 10:066011.

Miller DN, de Callafon RA (2012) Identification of linear time-invariant systems via constrained step-based realization. IFAC Proceedings Volumes 45:1155-1160.

Mina F, Benquet P, Pasnicu A, Biraben A, Wendling F (2013) Modulation of epileptic activity by deep brain stimulation: a model-based study of frequency-dependent effects. Front Comput Neurosci 7:94.

Miranda-Domínguez O, Gonia J, Netoff TI (2010) Firing rate control of a neuron using a linear proportional-integral controller. J Neural Eng 7:066004.

Montaseri G, Yazdanpanah MJ, Pikovsky A, Rosenblum M (2013) Synchrony suppression in ensembles of coupled oscillators via adaptive vanishing feedback. Chaos 23:033122.

Mormann F, Andrzejak RG, Elger CE, Lehnertz K (2007) Seizure prediction: the long and winding road. Brain : a journal of neurology 130:314.

Mormann F, Kreuz T, Rieke C, Andrzejak RG, Kraskov A, David P, Elger CE, Lehnertz K (2005) On the predictability of epileptic seizures. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 116:569.

Morrell M (2006) Brain stimulation for epilepsy: can scheduled or responsive neurostimulation stop seizures? Current opinion in neurology 19:164.

Morrell MJ, Group RNSSiES (2011) Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. Neurology 77:1295.

Morris GL, Gloss D, Buchhalter J, Mack KJ, Nickels K, Harden C (2013) Evidence-based guideline update: vagus nerve stimulation for the treatment of epilepsy: report of the Guideline Development Subcommittee of the American Academy of Neurology. Neurology 81:1453-1459.

Müller L, Tokay T, Porath K, Köhling R, Kirschstein T (2013) Enhanced NMDA receptor-dependent LTP in the epileptic CA1 area via upregulation of NR2B. Neurobiol Dis 54:183-193.

Nabi A, Mirzadeh M, Gibou F, Moehlis J (2013a) Minimum energy desynchronizing control for coupled neurons. J Comput Neurosci 34:259-271.

Nabi A, Stigen T, Moehlis J, Netoff T (2013b) Minimum energy control for in vitro neurons. J Neural Eng 10:036005.

Nagai Y, Goldstein LH, Fenwick PB, Trimble MR (2004) Clinical efficacy of galvanic skin response biofeedback training in reducing seizures in adult epilepsy: a preliminary randomized controlled study. Epilepsy Behav 5:216-223.

Nagaraj V, Lamperski A, Netoff TI (2016) Seizure Control in a Computational Model Using a Reinforcement Learning Stimulation Paradigm. Int J Neural Syst:1750012.

Nagaraj V, Lee ST, Krook-Magnuson E, Soltesz I, Benquet P, Irazoqui PP, Netoff TI (2015) Future of seizure prediction and intervention: closing the loop. J Clin Neurophysiol 32:194-206.

Nelson TS, Suhr CL, Freestone DR, Lai A, Halliday AJ, McLean KJ, Burkitt AN, Cook MJ (2011) Closed-loop seizure control with very high frequency electrical stimulation at seizure onset in the GAERS model of absence epilepsy. Int J Neural Syst 21:163-173.

Netoff TI, Pecora LM, Schiff SJ (2004) Analytical coupling detection in the presence of noise and nonlinearity. Physical reviewE, Statistical, nonlinear, and soft matter physics 69:017201.

Netoff TI, Acker CD, Bettencourt JC, White JA (2005) Beyond two-cell networks: experimental measurement of neuronal responses to multiple synaptic inputs. J Comput Neurosci 18:287-295.

Nijsen TM, Arends JB, Griep PA, Cluitmans PJ (2005) The potential value of three-dimensional accelerometry for detection of motor seizures in severe epilepsy. Epilepsy Behav 7:74-84.

Ogata K (2010) Modern control engineering, 5th Edition. Boston: Prentice-Hall.

Panuccio G, Guez A, Vincent R, Avoli M, Pineau J (2013) Adaptive control of epileptiform excitability in an in vitro model of limbic seizures. Experimental neurology 241:179.

Pasnicu A, Denoyer Y, Haegelen C, Pasqualini E, Biraben A (2013) Modulation of paroxysmal activity in focal cortical dysplasia by centromedian thalamic nucleus stimulation. Epilepsy Res 104:264-268.

Patel KS, Zhao M, Ma H, Schwartz TH (2013) Imaging preictal hemodynamic changes in neocortical epilepsy. Neurosurg Focus 34:E10.

Pineau J, Guez A, Vincent R, Panuccio G, Avoli M (2009) Treating epilepsy via adaptive neurostimulation: a reinforcement learning approach. Int J Neural Syst 19:227-240.

Proix T, Bartolomei F, Chauvel P, Bernard C, Jirsa VK (2014) Permittivity coupling across brain regions determines seizure recruitment in partial epilepsy. J Neurosci 34:15009-15021.

Prokhorov DV, Wunsch DC (1997) Adaptive critic designs. Neural Networks, IEEE Transactions on 8:997-1007.

Rashid S, Pho G, Czigler M, Werz MA, Durand DM (2012) Low frequency stimulation of ventral hippocampal commissures reduces seizures in a rat model of chronic temporal lobe epilepsy. Epilepsia 53:147-156.

Rogowski Z, Gath I, Bental E (1981) On the prediction of epileptic seizures. Biol Cybern 42:9.

Ross DA, Brunberg JA, Drury I, Henry TR (1996) Intracerebral depth electrode monitoring in partial epilepsy: the morbidity and efficacy of placement using magnetic resonance image-guided stereotactic surgery. Neurosurgery 39:327-333; discussion 333-324.

Rothman SM, Smyth MD, Yang XF, Peterson GP (2005) Focal cooling for epilepsy: an alternative therapy that might actually work. Epilepsy Behav 7:214-221.

Rothman SM, Perry G, Yang XF, Hyrc K, Schmidt BF (2007) Optical suppression of seizure-like activity with an LED. Epilepsy research 74:201.

Sato M, Racine RJ, McIntyre DC (1990) Kindling: basic mechanisms and clinical validity. Electroencephalogr Clin Neurophysiol 76:459-472.

Schade CM, Schultz DM, Tamayo N, Iyer S, Panken E (2011) Automatic adaptation of neurostimulation therapy in response to changes in patient position: results of the Posture Responsive Spinal Cord Stimulation (PRS) Research Study. Pain Physician 14:407-417.

Schevon CA, Trevelyan AJ, Schroeder CE, Goodman RR, McKhann G, Emerson RG (2009) Spatial characterization of interictal high frequency oscillations in epileptic neocortex. Brain 132:3047-3059.

Schiller Y, Bankirer Y (2007) Cellular mechanisms underlying antiepileptic effects of low- and high-frequency electrical stimulation in acute epilepsy in neocortical brain slices in vitro. J Neurophysiol 97:1887-1902.

Schwartz TH (2007) Neurovascular coupling and epilepsy: hemodynamic markers for localizing and predicting seizure onset. Epilepsy currents / American Epilepsy Society 7:91-94.

Shanechi MM, Orsborn AL, Carmena JM (2016) Robust Brain-Machine Interface Design Using Optimal Feedback Control Modeling and Adaptive Point Process Filtering. PLoS Comput Biol 12:e1004730.

Shanechi MM, Chemali JJ, Liberman M, Solt K, Brown EN (2013) A brain-machine interface for control of medically-induced coma. PLoS Comput Biol 9:e1003284.

Smyth MD, Rothman SM (2011) Focal cooling devices for the surgical treatment of epilepsy. Neurosurg Clin N Am 22:533-546, vii.

Stead M, Bower M, Brinkmann BH, Lee K, Marsh WR, Meyer FB, Litt B, Van Gompel J, Worrell GA (2010) Microseizures and the spatiotemporal scales of human partial epilepsy. Brain 133:2789-2797.

Sterman MB (2000) Basic concepts and clinical findings in the treatment of seizure disorders with EEG operant conditioning. Clin Electroencephalogr 31:45-55.

Stigen T, Danzl P, Moehlis J, Netoff T (2011) Controlling spike timing and synchrony in oscillatory neurons. J Neurophysiol 105:2074-2082.

Stypulkowski PH, Stanslaski SR, Jensen RM, Denison TJ, Giftakis JE (2014) Brain stimulation for epilepsy--local and remote modulation of network excitability. Brain Stimul 7:350-358.

Sutton R, Barto A (1998) Reinforcement learning: An introduction. Cambridge: MIT press.

Tan G, Thornby J, Hammond DC, Strehl U, Canady B, Arnemann K, Kaiser DA (2009) Meta-analysis of EEG biofeedback in treating epilepsy. Clin EEG Neurosci 40:173-179.

Tang H, Long H, Zeng C, Li Y, Bi F, Wang J, Qian H, Xiao B (2012) Rapamycin suppresses the recurrent excitatory circuits of dentate gyrus in a mouse model of temporal lobe epilepsy. Biochemical and biophysical research communications 420:199-204.

Taylor J (1997) Introduction to error analysis, the study of uncertainties in physical measurements.

Tellez-Zenteno JF, McLachlan RS, Parrent A, Kubu CS, Wiebe S (2006) Hippocampal electrical stimulation in mesial temporal lobe epilepsy. Neurology 66:1490-1494.

Toprani S, Durand DM (2013) Fiber tract stimulation can reduce epileptiform activity in an in-vitro bilateral hippocampal slice preparation. Exp Neurol 240:28-43.

Truccolo W, Donoghue JA, Hochberg LR, Eskandar EN, Madsen JR, Anderson WS, Brown EN, Halgren E, Cash SS (2011) Single-neuron dynamics in human focal epilepsy. Nature neuroscience 14:635.

Van Gompel JJ, Bower MR, Worrell GA, Stead M, Chang SY, Goerss SJ, Kim I, Bennet KE, Meyer FB, Marsh WR, Blaha CD, Lee KH (2014) Increased cortical extracellular adenosine correlates with seizure termination. Epilepsia 55:233-244.

Velasco AL, Velasco F, Velasco M, Jiménez F, Carrillo-Ruiz JD, Castro G (2007) The role of neuromodulation of the hippocampus in the treatment of intractable complex partial seizures of the temporal lobe. Acta Neurochir Suppl 97:329-332.

Viglione S, Walsh G (1975) Proceedings: Epileptic seizure prediction. Electroencephalography and clinical neurophysiology 39:435-436.

Viventi J et al. (2011) Flexible, foldable, actively multiplexed, high-density electrode array for mapping brain activity in vivo. Nat Neurosci 14:1599-1605.

Vonck K, Boon P, Claeys P, Dedeurwaerdere S, Achten R, Van Roost D (2005) Long-term deep brain stimulation for refractory temporal lobe epilepsy. Epilepsia 46 Suppl 5:98-99.

Ward M, Qing K, Otto K, Worth R, John S, Irazoqui P (2014) A Flexible Platform for Biofeedback-driven Control and Personalization of Electrical Nerve Stimulation Therapy. IEEE Trans Neural Syst Rehabil Eng.

Wilson CJ, Beverlin B, Netoff T (2011) Chaotic desynchronization as the therapeutic mechanism of deep brain stimulation. Front Syst Neurosci 5:50.

Wilson D, Holt AB, Netoff TI, Moehlis J (2015) Optimal entrainment of heterogeneous noisy neurons. Front Neurosci 9:192.

Wongsarnpigoon A, Grill WM (2010) Energy-efficient waveform shapes for neural stimulation revealed with a genetic algorithm. J Neural Eng 7:046009.

Yang XF, Rothman SM (2001) Focal cooling rapidly terminates experimental neocortical seizures. Annals of Neurology 49:721.

Yang XF, Chang JH, Rothman SM (2003) Long-lasting anticonvulsant effect of focal cooling on experimental neocortical seizures. Epilepsia 44:1500-1505.

Yang XF, Kennedy BR, Lomber SG, Schmidt RE, Rothman SM (2006) Cooling produces minimal neuropathology in neocortex and hippocampus. Neurobiol Dis 23:637-643.

Yeo SH, Franklin DW, Wolpert DM (2016) When Optimal Feedback Control Is Not Enough: Feedforward Strategies Are Required for Optimal Control with Active Sensing. PLoS Comput Biol 12:e1005190.

Zhao M, Nguyen J, Ma H, Nishimura N, Schaffer CB, Schwartz TH (2011) Preictal and ictal neurovascular and metabolic coupling surrounding a seizure focus. J Neurosci 31:13292-13300.

Zhu-Ge ZB, Zhu YY, Wu DC, Wang S, Liu LY, Hu WW, Chen Z (2007) Unilateral low-frequency stimulation of central piriform cortex inhibits amygdaloid-kindled seizures in Sprague-Dawley rats. Neuroscience 146:901.

Zijlmans M, Flanagan D, Gotman J (2002) Heart rate changes and ECG abnormalities during epileptic seizures: prevalence and definition of an objective clinical sign. Epilepsia 43:847-854.

# Appendix A Epileptor Model

## *A.1 State Equations*

$$\dot{x}_1 = \frac{1}{\tau_1}(y_1 - f_1(x_1, x_2) - z) \quad (1)$$

$$\dot{y}_1 = \frac{1}{\tau_1}(1 - 5x_1^2 - y_1) \quad (2)$$

$$\dot{z} = \frac{1}{\tau_0}(h(x_1) - z - I_1) \quad (3)$$

$$\dot{x}_2 = \frac{1}{\tau_1}\left(-y_2 + x_2 - x_2^3 + I_2 + 1.8u(x_1) - 0.3(z - 3.5)\right) \quad (4)$$

$$\dot{y}_2 = \frac{1}{\tau_2}\left(-y_2 + f_2(x_1, x_2)\right) \quad (5)$$

$$f_1(x_1, x_2) = \begin{cases} x_1^3 - 3x_1^2 & if\ x_1 < 0 \\ (x_2 - 0.6(z-4)^2 x_1) & if\ x_1 \geq 0 \end{cases} \quad (6)$$

$$f_1(x_1, x_2) = \begin{cases} 0 & if\ x_2 < -0.25 \\ 6(x_2 + 0.25)x_1 & if\ x_2 \geq -0.25 \end{cases} \quad (7)$$

$$\dot{u} = -\gamma(u - 0.1x_1) \quad (8)$$

$$h(x_1) = x_0 + \frac{10}{\left(1 + e^{\frac{-x_1 - 0.5}{0.1}}\right)} \quad (9)$$

$I_1 = 3.1, I_2 = 0.45, \tau_0 = 800\ or\ 400,\ \tau_1 = 0.005,\ \tau_2 = 0.01, \gamma = 0.01$

## *A.2 Reduced Model*

We used a two dimensional Epileptor model reduction to analytically compute the minimum energy stimulus parameters to suppress seizures given the time constant of the $z$ variable 18. The $z$ variable equation remained the same (Eq. 3) while the $x_1$ variable

was changed.

$$\dot{x}_1 = -x_1^3 - 2x_1^2 + 1 - z + I_1 \qquad (10)$$

$$\dot{z} = \frac{1}{\tau_0}(h(x_1) - z) \qquad (11)$$

Where $I_1$ is the constant current used in the full Epileptor model.

# Appendix B Reinforcement Learning

## *B.1 TD(0) Algorithm*

The expected reward matrix $Q$ is initialized to a value greater than the average reward $\mu_0$ plus some Gaussian noise disturbance with variance ≈0.001.

$Initialize\ Q(s,a) > \mu_0$ (1)

Select action $a'$

*Begin loop*

Apply action $a' \rightarrow a$

Measure new state $s'$

Measure actual reward $R$

$Calculate\ error$:

$\delta = R - Q(s,a)$ (2)

$Update\ Expected\ Reward$:

$Q(s,a) = Q(s,a) + \alpha * \delta$ (3)

*Select new action*:

$$P(a'|s') = \frac{e^{\frac{Q(s\prime,a\prime)}{\tau_S}}}{\sum_{i=1}^{n} e^{\frac{Q(s\prime,a\prime)}{\tau_S}}} \quad (4)$$

Where,

$$\alpha = (\alpha_0 - a_\infty)\beta^{count} + \alpha_\infty \quad (5)$$

$$\beta \approx e^{-w/ISI} \quad (6)$$

$$count = \begin{cases} count+1, & if\ s\prime = s\ and\ a\prime = a \\ count, & otherwise \end{cases} \quad (7)$$

*End loop*

# Appendix C State space control

## *C.1 Ho-Kalman Algorithm*

The impulse response of the system can be represented using the following equation:

$$y_t = \sum_{k=0}^{\tau} G_k u_{t-k} + v_k \quad (1)$$

Where,

$$G_k = \begin{cases} D, & k = 0 \\ CA^{k-1}B, & k > 0 \end{cases} \quad (2)$$

is the transfer function of the system that maps inputs to outputs.

The impulse response can be used to construct block-Hankel matrices:

$$H = \begin{bmatrix} G_1 & \dots & G_l \\ \vdots & \ddots & \vdots \\ G_r & \dots & G_{r+l-1} \end{bmatrix} \quad \text{and, } \bar{H} = \begin{bmatrix} G_2 & \dots & G_{l+1} \\ \vdots & \ddots & \vdots \\ G_{r+1} & \dots & G_{r+l} \end{bmatrix} \quad (3)$$

Both Observability ($O_r$) and Controllability ($C_l$) matrices can be extracted from the

block Hankel matrices:

$$O_r = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{r-1} \end{bmatrix} \quad \text{and, } C_l = [B \quad AB \quad \ldots \quad AB^{l-1}] \tag{4}$$

We can then rewrite the block-Hankel matrices as:

$$H = O_r C_l \quad \text{and, } \bar{H} = O_r A C_l \tag{5}$$

Since we are dealing with noisy data, we are unable to extract the exact Observability and Controllability matrices from the neural evoked response (impulse response). We can instead use Singular Value Decomposition (SVD) to estimate the Observability and Controllability matrices from the block-Hankel matrix:

$$\hat{H} = U_n \Sigma_n V_n^T \tag{6}$$

$$\hat{O}_r = U_n \Sigma_n^{\frac{1}{2}} \text{ and } \hat{C}_l = \Sigma_n^{\frac{1}{2}} V_n^T \tag{7}$$

The state space matrices are then calculated as follows:

$$\hat{A} = \hat{O}_r^{+} \bar{H} \hat{C}_l^{+} \tag{8}$$

$$\hat{B} = \hat{C}_{l(:,1)}$$

$$\hat{C} = \hat{O}_{r(1:n,:)}$$

$$\hat{D} = G_0$$

## *C.2 Model accuracy across prediction horizons*

In order to determine model prediction accuracy into the future at different time horizons we used a popular method described in (Ljung, 1999).

Given a dynamical system with the form:

$$x(t+1) = Ax(t) + Bu(t) \quad (1)$$

$$y(t) = Cx(t) + Du(t)$$

where $x \in R^{Nx1}$ is a vector containing the unobservable states of the system, and $A$, $B$, $C$ are matrices that contain the coefficients for the model. A constant gain value $D$ is included.

We can design a Kalman Filter around our dynamical systems model to forecast future states at different time horizons:

$$X = A'X + B'u(t) + L'y(t) =$$

$$\begin{bmatrix} x(t+1|t) \\ x(t+1|t-1) \\ \vdots \\ x(t+1|t-\tau+1) \\ x(t+1|t-\tau) \end{bmatrix}$$

$$= \begin{bmatrix} A+LC & 0 & \cdots & 0 & 0 \\ A & 0 & \cdots & 0 & 0 \\ 0 & A & \cdots & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & \cdots & A & 0 \end{bmatrix} \begin{bmatrix} x(t|t-1) \\ x(t|t-2) \\ \vdots \\ x(t|t-\tau) \\ x(t|t-\tau-1) \end{bmatrix} + \begin{bmatrix} B \\ \vdots \\ \vdots \\ \vdots \\ B \end{bmatrix} u(t) + \begin{bmatrix} -L \\ 0 \\ \vdots \\ \vdots \\ 0 \end{bmatrix} y(t)$$

$$\begin{bmatrix} y(t|t-1) \\ \vdots \\ y(t|t-\tau) \end{bmatrix} = C \otimes \begin{bmatrix} x(t|t-1) \\ \vdots \\ x(t|t-\tau) \end{bmatrix} \quad (2)$$

Where $X \in R^{\tau xN}$ is the state prediction horizon matrix. Using this framework $\tau$ is the largest prediction horizon that needs to be forecast. The Kalman Gain is the $L \in R^{Nx1}$ vector and is only used in the calculation of the one time-step prediction horizon. The LFP prediction is determined by multiplying the $C \in R^{1xN}$ vector with the state prediction matrix $X$ using the Kronecker product ($\otimes$) .

We calculated model accuracy using the Coefficient of Determination ($R^2$) metric which is formulated as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y(i) - LFP_{noisy}(i)\right)^2}{\sum_{i=1}^{n}(y(i) - \bar{y})^2} \quad (3)$$

where $y$ is the model prediction values of the cross validation data set, $\bar{y}$ is the mean of the model prediction values, and $LFP_{noisy}$ is the raw LFP data.