Chronic Monitoring and Longitudinal Efficacy of Neuromodulation for Neurological Movement Disorders

Tomasz Fraczek

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2023

Reading Committee:
Jeffrey Herron, Chair
Howard Chizeck, Chair
Eric Shea Brown

Program Authorized to Offer Degree:
Neuroscience

University of Washington

Abstract

Chronic Monitoring and Longitudinal Efficacy of Neuromodulation for Neurological Movement Disorders

Tomasz Fraczek

Co-Chairs of the Supervisory Committee:
Jeffrey Herron
Department of Neurological Surgery

Howard Chizeck
Department of Electrical and Computer Engineering

Modern chronic stimulation therapies with sensing capability offer a unique opportunity to effectively manage the symptoms of a disease while also providing an understanding of the progression of neurological disorders. I designed and implemented a framework for monitoring the long-term effectiveness of deep brain stimulation (DBS). Using this framework, I compared several adaptive DBS algorithms for managing essential tremor symptoms to determine how to design more effective next generation adaptive DBS. I also show that the beta band biomarker used to decode movement is sufficiently stable over time for the adaptive paradigm to remain effective for many months. In a second implementation focused on live distributed data collection, I demonstrate our framework for optimizing adaptive deep brain stimulation remotely in the patient's home. I also develop a method using this framework to quantify the long-term accuracy of neural based linear discriminant decoders. I evaluate its accuracy in the case of a particular patient-specific adaptive DBS algorithm for Parkinson's disease. This work serves as a first step towards the data collection and longitudinal analysis capabilities needed to make adaptive DBS a therapy feasible in the chronic at home environment, and facilitate a deeper understanding of the longitudinal effects of neurological movement disorders.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

This dissertation work is the compendium of many years of hard work, fraught with many dead ends, unexpected challenges, and sudden moments of euphoric clarity. None of it would have been possible without the help I received along the way.

I would like to thank my family Witold Fraczek, Joanna Werner-Fraczek, and Zofia Fraczek, who always had my back. Their resilience and dedication through out their life has been an absolute inspiration. Their faith in me has been a constant pillar of support that I have found indispensable.

I would like to thank my wonderful girlfriend Anais who has stood by my side and put up with all my PhD related stress. She offered the daily gentle support and encouragement that enabled be to keep pushing forward with the eork.

I would like to deeply thank my advisors, Jeffery Herron and Howard Chizeck, who have been instrumental to my success as a graduate student. The hardest questions they asked me always led to the biggest breakthroughs. Without their constant support and insight I would not have been able to complete this thesis.

I would like to thank the previous students in the lab, Ben Ferleger, Margaret Thompson, Andrew Haddock, Brady Houston, and Timothy Brown. They are the giants upon whos shoulders I stood to acomplish everything that is detailed here. Special thanks to Ben Ferleger who welcomed me into the lab when I first joined and mentored me as I figured out my focus for the remainder of my PhD.

I would like to thank my entire network of friends I formed here in Seattle who have made these last few years one constant adventure.

I would also like to thank my funding sources, the University of Washington, the

Neuroscience program, and especially the Weill Neurohub Foundation. Thank you to Medtronic who provided the devices implanted in our patients and Rune Labs for assisting in the data collection processes.

DEDICATION

to my family and friends,

whose support let me achieve my goals

Chapter 1

INTRODUCTION

Recent advancements in medical technologies are leading to an increased number of implanted devices for the treatment of chronic diseases [93]. These devices are a promising alternative to conventional treatments. Instead of taking medication on a daily basis that has wide-ranging side effects, a single surgery can implant a therapeutic device that can remain effective for years. Many devices are also able to adjust the therapy delivered to the patient either automatically based on onboard sensors, or by direct input through a connected patient controller. However, the necessity of the implantation surgery makes it particularly important that these devices remain effective for long periods of time [100]. The ability of many of these new devices to send diagnostic data along with other sensors to monitor the patient's wellbeing are extremely useful to solving this issue. However, coordinating data ingestion from multiple independent sources is a significant challenge, especially in the research and development environment [72]. If correctly processed, this rich longitudinal data can help us design devices and algorithms that are robust to both short term and long term changes in the patient's therapeutic needs. Successful optimization of such devices will allow patients to manage their symptoms with minimal complications and live a normal life.

In this work, I focus on deep brain stimulation (DBS), a specific kind of neuromodulation therapy approved by the FDA for treatment of Parkinsons, Essential Tremor, epilepsy, and under investigation for the treatment of many more conditions [100, 93]. DBS has been demonstrated both in clinic and in practice to be highly effective at managing the symptoms of these diseases. However, constant natural

changes in patient state and symptom intensity have promoted the development of adaptive DBS (aDBS) paradigms that automatically adjust stimulation to minimize disease symptoms while avoiding side effects [7, 98]. Ideally, these algorithms strive to manage both rapid changes in the patients stimulation needs as well as long term neural and disease changes [100, 44]. These algorithms should also be deployed on a fully implanted device, to allow the patient to go about their daily life without requiring the patient to remain tethered to external hardware [44, 100]. This means that a significant portion of the algorithm must be implemented within the restricted capabilities of the implanted hardware. Such an optimized implementation requires significant amounts of multi-modal data to design, deploy, and test. Collecting this dataset in an entirely remote fashion from the patient's home environment allows us to have a unique perspective on how these algorithms can function in the patient's daily life.

Achieving the goal of fully deployable aDBS requires overcoming a number of challenges outlined in this thesis. In Chapter 1, I develop a framework for ingesting and combining multi-modal data to completely capture a patient's state, and track changes over the course of many months. This allows us to have a unique perspective on the long term efficacy of these aDBS therapies and to better understand how to optimize them for long-term stability. In Chapter 2, we apply this framework to a previously collected clinical experiment dataset of aDBS development in essential tremor patients to validate the long term stability of movement biomarkers. In Chapter 3, we use an expanded version of the framework to collect live data from the patient's home environment during their daily routine. We then analyze almost two years' worth of data from this patient to understand long term changes in neural biomarkers and how they affect the decoding ability of the device. This work opens the door to novel aDBS algorithms that account for long term changes in neural biomarkers and thereby maintain clinical effectiveness for long periods of time.

Chapter 2

THE FRAMEWORK

2.1 Introduction

In today's modern increasingly connected world the daily practice of medical research is changing at an increasing pace. Small internet connected devices are becoming commonplace, forming the so-called internet of things (IoT). Although the nature and development of the IoT is outside the scope of this work, it's potential for application to neuroscientific, and biomedical research in general is extraordinary [140, 82]. Remote monitoring also can provide many practical benefits in the clinical setting, such as reduced costs, better patient outcomes, and reduced hospitalizations due to early interventions [42]. However, a majority of the remote monitoring literature focuses on the clinical impact of remote monitoring on patient outcomes, and relatively little research is devoted to the benefit of chronic monitoring to research and a more profound understanding of the disease [82, 36].

Remote at-home monitoring of patients is an essential component of understanding chronic diseases. Clinical care of chronic diseases been hampered by the limited contact between the patient and the clinician responsible for providing care []. Therefore, the clinician only gets small snapshots of the patient's disease state, providing a potentially skewed view of the disease progression and the effectiveness of the therapy [42]. This can be especially poignant in the case of diseases with a circadian component like Parkinson's disease [145, 166]. Other disease, may be affected by dietary intake, or symptoms may change in the often stressful nature of the clinical setting. In all of these cases, the IoT can offer a huge advantage.

As medical device technology develops, more and more patients are implanted

with therapeutic devices to manage disease symptoms [72, 140, 146]. For some of these devices, like insulin pumps for diabetes, a precise and reliable biomarker (blood glucose level) is available that precisely specifies the needed therapy level (amount of insulin introduced [165]. For therapies where an undisputable biomarker of symptoms is not feasibly available, many devices assume that the patient's symptoms are more or less constant and therefore deliver therapy at roughly a constant rate [93, 101]. The placement of small monitoring devices in the patient's home can allow for a more holistic view of the disease and treatment [94, 49]. These can include cameras, as well as body-worn sensors such as accelerometers or neural monitors. Chronic monitoring of this form will allow us to understand both the short term variation in symptoms as well as look at long term effects, both of which will allow us to develop improved treatments.

In the field of adaptive deep brain stimulation (aDBS) research one of the great outstanding questions is the numerous overlapping timescales at which the adaptive algorithms need to be updated [101, 7]. These are the quick changes, related to the patient's movement and sudden changes of symptom intensity [49], medium timescale changes over the course of the day due to the circadian rhythm and medication [50, 148, 105], and slow variation over the timescale of weeks and months due to the progression of the disease [101]. Changes on all of these timescales are closely entangled, making if nearly impossible to separate out these effects with in clinical experiments. As a result, understanding all of these timescales is essential to be able to deliver useful aDBS treatments that are effective on short time scales and continue to deliver clinical benefit for long periods of time.

Collecting data that could answer these questions is a huge challenge. The at home environment, by its very nature, is much more dynamic and less controlled than a clinical lab experiment. As a result, the data collection philosophy is different than that commonly used during tightly specified experiments. Huge volumes of metadata must be collected for the data acquired from the at home environment to

be amenable to analysis [83]. Managing many remotely deployed devices, each with uncertain connections and outside of the researchers direct control also brings many software and robustness challenges [76, 80, 160]. These are all challenges that must be considered carefully, but can be overcome with the right software and approach.

In this chapter, we review the requirements such a data collection system would need to fulfill, along with several software solutions that try to fill this gap. We then cover how we worked to build a software package capable of transforming a retrospective dataset into one usable for longitudinal analysis and what we learned from this process. Based on this experience, we developed a software ecosystem, DataNet, that allows both remote data collection and packages it into a longitudinal analysis dataset. Lastly, we review holistic learning as well as ethical considerations stemming from this sort of work.

2.1.1 Guiding Design Principles

The advantages of a distributed chronic monitoring framework come with their own challenges which mandate consideration. Specifically, the data from multiple independent sources needs to be processed, synchronized and integrated together to be useful. The data has to be de-identified and transmitted securely to protect patient privacy. The accumulation of data has to be managed and indexed in an automated way to prevent the formation of an unmanageable dataset. Additionally, the system must be robust enough to survive the constantly changing research environment, ensuring that any piece of data is traceable back through the entire acquisition pipeline. This capability is enabled thanks to massive collection of detailed metadata. Lastly, all the data must be compiled into a widely used format to enable collaboration with researchers around the world.

2.1.1.1 Live Distributed framework

The first challenge is the very practical problem of bringing together data from multiple independent data sources into a single cohesive dataset [146, 76, 128, 160, 150, 132]. This process can in general be considered to occur in two separate steps. The first, consists of collecting all the data for each modality from the device they are recorded on. This can be done in several ways, but generally will consist of automatically having each device upload to a centrally accessible repository independently [128]. Care must be taken at this stage to ensure that the data is collected reliably. Since these data become the source material, and cannot be reconstructed, it is generally beneficial to minimize the amount of pre-processing done at this stage [80]. This allows the automation scripts to be as simple and robust as possible. It is also important that the ingestion scripts for multiple data modalities are as independent as possible [160, 80]. Once the data is accessible on the cloud, the distributed nature of cloud infrastructure allow us to consider the data as 'safe': it is unlikely to be lost unless it is intentionally deleted. Here, it can be stored and processed into the final ready format. Once all the raw data has been safely moved to a centralized endpoint, it is ready for the more complex secondary processing steps needed to bring it into a usable format. This includes time alignment, noise/artifact removal, as well as more practical considerations such as metadata parsing and the calculation of derivative features. These steps in the pipeline are secondary to the first round of automation as they can be repeated. The cleaned versions of the data can either be dumped back into the data lake, or packaged into a more refined format.

2.1.1.2 Patient Privacy

Storing medical patient data requires a great deal of consideration. If the data processed by the pipeline is considered identifiable patient health information (PHI), then the governing body (NHI, HIPPAA) will specify stringent regulations on how

the data must safely be stored [108, 109, 111]. These regulations differ by country, and administrative district: but they largely all have the same goal: ensure that the patient personal and potentially damaging data is never made available to the public. Important to this legislation is the balance between ensuring the privacy of patient information while still providing a framework where this data can be used and shared for the purpose of research [111, 164, 73]. Data leakage and privacy violations are unfortunately not an uncommon phenomenon and can take many forms from malicious to accidental which are reviewed extensively elsewhere [2, 109, 73, 1]. Here, we review strategies that can be taken to minimize leakage of private data via both malicious intent, as well as ensure standards to minimize accidental leakage of data.

The best, yet underappreciated, approach to ensuring the patient's data is not leaked is simply not having the data. Although this extreme would prevent any research or data analysis from being done, it does point us in the direction of an important conclusion: the more data we gather, the better we will be able to understand the needs of the patient but the greater of a risk we expose both the patient and ourselves to [111, 108, 164, 37]. If we only collect anonymous accelerometer data in the clinic, then the range of questions we can answer is limited, but the patient's privacy is more or less ensured, even if all the data were to be released. However, if we collect and store video data from everywhere in the patient's home, 24/7, then we can potentially answer a huge host of questions about the patient's disease states over time. The patient's privacy is greatly at risk. Even a partial data leak could be catastrophic to the patients wellbeing. As a result of this tradeoff we must be very conscious of the data we collect, and only collect private information when it is absolutely necessary.

Once we have private information, there are several ways that we can ensure this data remains private. A great amount of research has been done into designing hardware and software architectures to ensure this data security [33, 3, 149, 147, 134, 60]. Fundamentally, many of these methods revolve around obfuscating and de-

identifying the data so that it cannot be linked back to the patient [164, 37]. Locally and immediately processing data into non-identifiable features is a great approach for mature projects [37, 60]. In this scenario, the data is processed in a fully automated fashion directly on the device where it is recorded, and only secondary derivative data, which can more easily be anonymized, is saved or transmitted. The original most sensitive data can then be quickly deleted ensuring data privacy. This is a wonderful approach in principle, but it is unfortunately not always suitable to a research environment. Since we as researchers and engineers unfortunately almost never know the answers to the questions we ask, it can be difficult to precisely predict the data needed to answer these questions. Moreover, as our research pipelines change, and mature, the features used originally often lose value. All data collection and all experiments need to be performed again. Since we are dealing with human subjects, this is often simply not possible and can lead to a loss of projects or treatment for the patient. As a result, to be able to maximize the learnings and positive outcomes to our patients, we must store the raw, sensitive forms of data.

A common solution to this problem is the usage of encryption to be able to store and transfer data while ensuring that third parties are not able to read the data [164, 23, 60]. Storing encrypted data ensures that no privacy leak can occur even if the device containing the data is stolen, and a VPN can ensure that all network traffic between computers on the network is encrypted and is therefore not decipherable by a potentially malicious interceptor. Particularly relevant to our work is the fact that a VPN such as wireguard can allow us to remotely access machines and transfer data with all the security of a local connection [31]. If implemented perfectly, modern encryption standards can ensure that the encrypted data cannot be decrypted without prior knowledge of the key in any reasonable timeframe, given currently available technology. However, these methods are only as secure as the encryption key management of the team using them is [149, 135, 164]. If the passwords securing the encryption are compromised, then the entire encryption architecture is compromised

[135]. As a result, encryption cannot be considered the cure-all for data security, and care has to be taken at every step of building the data collection architecture that encryption is maintained and keys are retained.

2.1.1.3 Nuances of the Research Environment

An important distinction designing data infrastructure for the research environment and for a more traditional production environment is the fact that the research environment is much more dynamic. In a production environment, the requirements and expectations are relatively well specified ahead of time, and are generally not expected to significantly change over time. Most of the research on the development of data collection architectures is focused mature production environments such as hospitals or specific clinical studies with tightly limited goals [146, 83, 128, 160]. On the contrary, the research environment is necessarily expected to involve much more variation. This is particularly true in and research engineering context, which is constantly experimenting with new technologies and approaches to best solve the problem at hand. As a result, the data collection system needs to be prepared to handle this variability by minimizing assumptions.

The research environment brings the need to handle a significantly increased frequency of faults in the pipeline compared to most production systems [81, 80]. There are a few different ways in which this can manifest itself. First, a data collection environment that is in active development will likely involve changes of the recording hardware. Early experiments might be carried out with one recording tool that then proves insufficient and needs to be replaced or upgraded. The new device may record different modalities at different frequencies and output them in a completely different format. In this situation, the system needs to be able adapt to the new recording hardware with minimal software changes. Similarly, in a development environment that is being constantly refined it is especially likely that parts of the system are unstable, leading to common intermittent and permanent faults of all kinds, where

the data collection pipeline can either temporarily break of fail completely [80]. This makes software isolation a particularly important component of the system design so that the rest of the data collection infrastructure can still continue to function, salvaging as much data as is possible [161, 81]. Moreover, the pipeline should minimize assumptions of what the data should look like. Any assumptions that do need to be made should be externalized to the peripheral parts of the pipeline. This is a known design pattern that is particularly important in this context, and managing data in this way results in the formation of a data lake, which brings it's own management challenges [110, 41].

The final component of managing the research environment is to be able to keep track of where data has come from. Especially since the code is being changed during the development process it is important to be able to trace the data back through the pipeline. This allows the team to go back through the pipeline to detect bugs, and hopefully be able to identify and disinfect the affected data. Moreover, in an engineering environment it is not always feasible to explicitly track every single change that has been made to the system. If all the data includes metadata that specifies the exact code that was used at every step of the data collection and processing process, then this is taken into account.

2.1.1.4 Managing the data lake

As the size of the dataset increases it is ever easier for the data lake to outgrow processing capacity. In this scenario, more and more data is collected faster than it can be processed and utilized in analysis. It is a self-propagating problem, where every next bit of data that comes in makes the situation worse. As partial analysis is done on little segments, the number of formats, parsing scripts and discordant data storages grows exponentially. Very quickly, the data lake becomes so convoluted that no single script is able to parse through all of it. This makes it nearly impossible to effectively search for the relevant data needed for any specific analysis without a

huge amount of manual data parsing. Analyses slow down, but deadlines continue to pressure the team to continue analysis without taking the time to fix the data store, making it ever more fragmented and unmanageable.

There are several key tools and methods that need to be implemented to effectively manage this problem and ensure the dataset remain continuously useful, including a logical shallow organization of the data store that houses the lake and an accessible metadata catalog [133]. These techniques need to be followed quite intentionally thought the entire data collection process, as retro-active attempts to correct early errors are very costly.

Shallow data organization is an essential component of data lakes and ensures that the lake remains parse-able and usable as new data is added [41, 133]. The naïve approach in the research community is to keep data by excessive use of folders, where days of experiments are grouped by experiment type, and then data within those days is further organized by some metadata values such as recording type or participant. Although this approach seems convenient when manually looking for data, it does not scale well to larger and more complicated datasets. This sort of organization requires the management software to know all possible organizations of the data, and the meanings of all possible folder groupings. This task quickly becomes near impossible to automate in an effective way. Instead, it is best to organize both the raw data and the processed data into a reasonable shallow directory tree and instead have all the relevant metadata stored in a file format alongside the data. This could be something as simple as a metadata .json file accompanying the raw data .csv files in the data directory, or something more sophisticated like the attribute and dataset structure enforced in Neurodata without Borders (NWB) archives [150]. Limiting the level of nesting, and thereby the depth of the data lake, allows the data management code to remain more flexible and able to effectively search through the data store for the relevant data. For example, if the data is organized by <date> -> <patient> -> <modality> -> <recording session> -> <experiment name> -> <prompt type>,

then selecting all data for a specific patient and prompt type not only requires the software know the entire directory specification, but also all recording sessions and experiment names. If one particular experiment then also had sub-experiments, the whole infrastructure quickly breaks down. On the contrary, having all the data in a single folder with hashed filenames may be very easy for a computer to search, but becomes nigh impossible to oversee or understand manually. The more idealistic reader may suggest that humans shouldn't be interfering with the raw data store anyways, it should be entirely sanitized and run by automated processes. Although this could be a valid point, it is simply not an attainable ideal in a real research setting. Despite the greatest intentions, at some point manual or semi-manual intervention and verification will be needed. A shallow data lake, with data organized on only a few levels, promises to balance these needs [110]. For example, if the data was instead organized by <patient> -> <modality> -> <session-id> then the iteration process for searching through all the data becomes much simple and more robust. All the metadata that was before stored across the directory organization can now be stored in metadata files. This allows a human to effectively manually locate any specific piece of data if needed, but also allows automated systems to be simple and effective. Moreover, if done well, this structure does not give undue primacy to any particular bit of metadata. During analysis, the data can be just as easily be divided by experiment as it can be by prompt type. Overall, a well-designed shallow data organization schema allows efficient simple code, efficient parsing of data and flexible metadata based selection while still enabling humans to still be able to manually navigate the raw data store.

One drawback of this shallow data organization is that manually finding data in the resulting data warehouse can be challenging [110]. This problem can be solved by usage of an indexing table that makes important metadata easily visible and searchable [41, 133]. The exact format of this indexing database does not much matter: it can be anything from a spreadsheet to a fully-fledged SQL database. What is

critically important is that the database is kept up to date [110]. A database that is out of date almost instantly becomes useless. This means that database updates must either be automatic, or there must be a specific person assigned to ensuring the data base remains updated. Secondly the type of metadata exposed in the database must be metadata that is relevant for finding data for analysis. Using a comprehensive database of metadata alongside a flat data storage can simultaneously ensure that data is both easy to parse and extend programmatically while remaining easy to search.

2.1.1.5 Data Sharing

In today's modern research environment collaboration and data sharing is essential, as demonstrated by the existence of the BRAIN initiative [71]. A leading way that researchers in neuroscience can share data is using the NWB data format [150, 131]. This data format, further reviewed below, is designed to provide data along with meaningful metadata to allow for collaborative data analysis.

2.1.1.6 Metadata Collection

Each of the above requirements have either directly or implicitly required that metadata be collected alongside the raw data [133, 110, 151]. This requirement has been identified in the field many times, and was even one of the driving factors behind the creation of the NWB data format [150, 131]. One of the most challenging components of metadata collection is understanding the depth of metadata that should be collected. Even my fellow researchers have said that a particular form of metadata need not be collected, only to months later realize that the experimental protocol had changed without a log of the specific change.

2.1.2 Existing Data Solutions

There has been extensive research on defining the requirements for distributed data collection systems and theoretical developments on the architectures possible [146, 83]. There are many potential solutions for enterprise scale clinical remote monitoring for entire hospitals or hospital systems [128, 161]. There are far fewer solutions that address the needs of a smaller research team. However, are projects that have attempted to solve parts of the problem and address parts of the needs expanded above.

2.1.2.1 NWB and NWBQuery

By far the most significant player in the space of neuroscientific data management solutions is NWB. As described above, NWB consists primarily of an open-source data specification standard, supported by a suite of community data interfacing and analysis packages [151, 150, 131, 132]. In the past years, NWB has become an important vehicle for fostering neuroscientific collaborations and data sharing as visible through and increasing number of publications citing NWB [132, 58]. In addition to the core format, NWB is supported by an ever-expanding ecosystem of user-built packages for processing, managing, analyzing and visualizing NWB datasets [131, 132].

One package in the NWB ecosystem that is particularly relevant to the challenges we aim to solve is NWBQuery [70]. This set of three related packages is designed to help in searching through large NWB datasets, either by directly parsing the metadata in the datasets, or by constructing a SQL database to enable quicker future searches. These work for any NWB datasets, but as the name implies, require a complete NWB dataset to work well. This makes them very powerful tools for working with the data once it has been shared, but limiting their utility in the live data collection setting. Once aspect of the NWB environment that is missing is the initial data acquisition capabilities. As a whole, the NWB ecosystem is designed to facilitate the sharing of complete, prepared datasets between research groups, and the collection of that data

is left outside of the specification. In fact, there are several components of the NWB approach that do not make it suitable as a data storage format in a live collection setting. Most important among these is that the NWB datasets are static, and once created cannot be edited. This is important for reliable data sharing, but causes problems in data collection. As a result, there is a need for a framework that can robustly collect and analyze neuroscientific data to be later exported to NWB for sharing with other groups.

2.1.2.2 RUNE

Some of the missing distributed data collection capabilities are provided by the RUNE platform [25]. Developed by rune labs, this platform focuses on remote at-home neurological data collection. IN addition to a focus on facilitating patient-clinician interactions through the StrivePD app, Rune Labs also provides a platform to assist research teams such as ourselves with data collection explicitly for the purpose of better understanding chronic diseases. This includes the machinery to ingest many different modalities of data, visualize data availability, and make the data available for download. Thanks to a partnership with Rune Labs, we were able to make use of some of their capabilities in our framework, specifically for the collection of Apple watch data. Despite its capabilities as a data collection platform, RUNE has a few limitations. Most importantly, it a closed source enterprise service. This makes it difficult to adapt to the specific needs on any research project. If a new data modality needs to be collected, there is no good way to include it in the RUNE API. Future versions of the RUNE API may improve on these limitations, but the RUNE software will likely remain a closed-source enterprise solution, resulting in the need for an open-source alternative.

Figure 2.1: A summary of all the available data collected during the Activa PC+S tremor experiments developing aDBS for essential tremor. Four patient participated in the study, which lasted for almost 5 years. Patients would visit the clinic once every few months to participate in aDBS experiments. Note that different kinds of experiments were carried out may times, and the aDBS paradigm used was constantly updated.

2.2 V1: Retrospective Activa Dataset

The majority of the data for this dataset was collected before I joined the lab, as part of a 5 year project to demonstrate the capabilities of aDBS in patients with essential tremor (ET). As a result, the data was not intended for centralized parsing: the data for each set of experiments was specifically customized to answer just the specifics of the question being investigated. The only commonalities were those left over from the software originally written to capture data from the Activa PC+S device. However, this software had been modified by the successive students. These conditions led to a specifically challenging design process for longitudinal analysis of this dataset.

2.2.1 Existing Data

Data was collected from four ET patients at the University of Washington Medical Center, by previous graduate students during the years (2015-2020). All patients received unilateral implantation of the Medtronic investigational Activa PC+S system, under an FDA Investigational Device Exemption (IDE), as treatment for pharmacologically refractive ET. All data was collected as part of IRB approved experiments focused on the development and optimization of aDBS algorithms. The data were stored and organized as individual, largely unrelated, experiments. Each experiment was composed of data specifically required by the particular question being investigated. Data modalities include neural data from the Activa PC+S implant, inertial measurements from a wrist-worn smartwatch, and logging of both experiment-specific and patient state.

2.2.1.1 Neural Data

Neural data is collected via the Medtronic Activa PC+S system in the form of LFPs from the stimulation site and from the surface of the cortex. In the case of our patients, stimulation was applied in the ventral intermediate nucleus of the thalamus (VIM). Cortical recordings were collected from a strip of four electrodes laid across the central sulcus, in the region associated with motion of the hand. Electrode placements were verified by CT scans after implantation.

All neural recordings included here are differential recordings between electrodes. Data can be collected at various sampling rates depending on the recording mode. During onboard recordings, data can be collected from as single source at up to 800 Hertz. These data are saved in a Medtronic-defined XML format and can later be transferred for storage and analysis via USB. When recording via the Nexus D or Nexus E systems, data are streamed to the computer in real time. This allows data to be used in experimental aDBS applications, but limits the sampling rate to 422

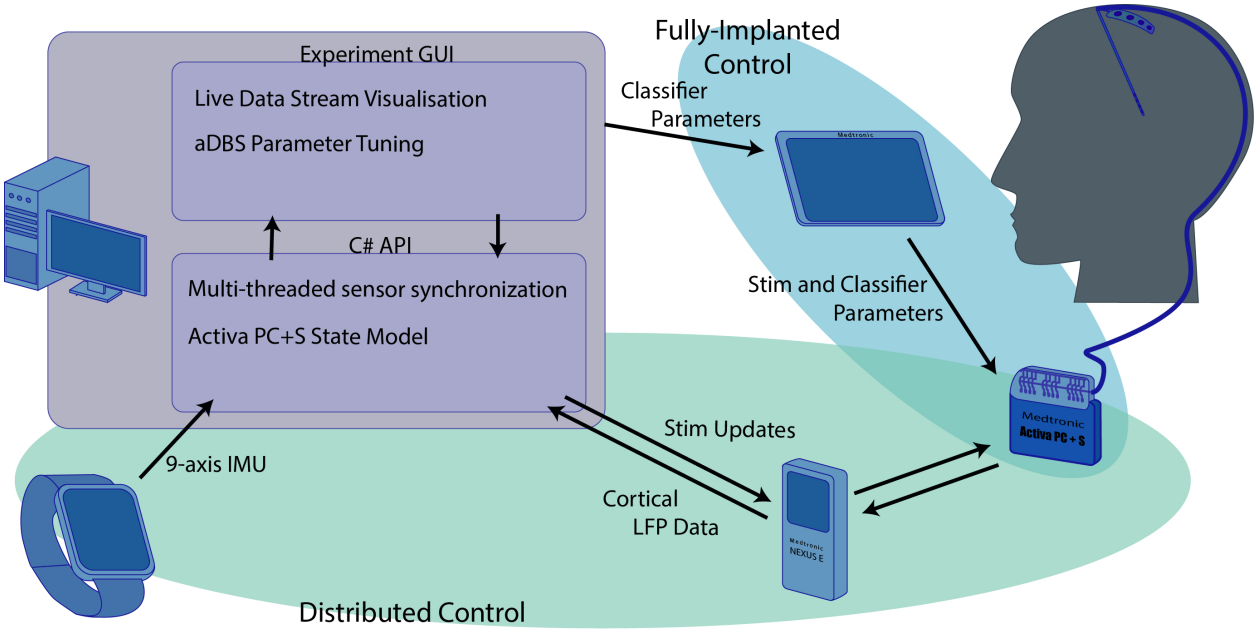

Figure 2.2: A schematic of the data collection setup previously developed by the lab to collect data during aDBS experiments for ET. Data could either be collected through the distributed pipeline (green) which included streamed cortical and VIM neural data during aDBS experiments controlled by a desktop computer, or via the fully-implanted pipeline, which included logged neural data for experiments where aDBS was controlled by a pre-trained classifier inside the IPG. Both kinds of data include smartwatch IMU, movement prompts, and a log of stimulation where available.

Hertz if recording from a single source, or 200 Hertz simultaneously recording from two sources. These data are stored in a site-specific plain-text format, designed for easy use with our aDBS applications. This difference in sampling rates and storage formats must be accounted for when developing a standardized analysis framework.

2.2.1.2 Inertial Measurements

Inertial data is collected from a wrist-worn smartwatch, referred to as the inertial motion unit (IMU). Gyroscope and accelerometer data is streamed in real time via Bluetooth at a sampling rate of 100 Hertz. The resulting time-series is saved in a site-specific format parallel to the other data modalities. Most of our analysis used only the gyroscope component of the data, and focused on the magnitude of the three component (x, y, z) vector.

2.2.1.3 Experiment State

Experimental changes to DBS during aDBS experiments were limited to stimulation amplitude changes, and remained well within the clinical constraints set by the neurologist. Changes to DBS frequency were clinically motivated and were made by the patient's neurologist, independently of aDBS experiments. As a result, stimulation frequency is not consistent across all data and must be tracked. Maximal stimulation amplitude was also varied in this same way, independently of aDBS experiments. Some metadata (patient ID, session ID, date, experiment type) can be deduced from a standardized (BIDS inspired) directory naming structure. Any other external experimental parameters must be either interpreted from the recorded data, or be manually extracted from session documents.

2.2.2 Implementation

The framework I built is designed to discover the maximal amount of compatible data with minimal requirements on manual preparation of the data storage. Data from various measurement sources is integrated and made available in a standardized way. The design choices described here ensure that there is no innate hierarchy in the loaded data, and large-scale analysis can be performed across any variable with equal ease. The code for this implementation of the longitudinal analysis framework

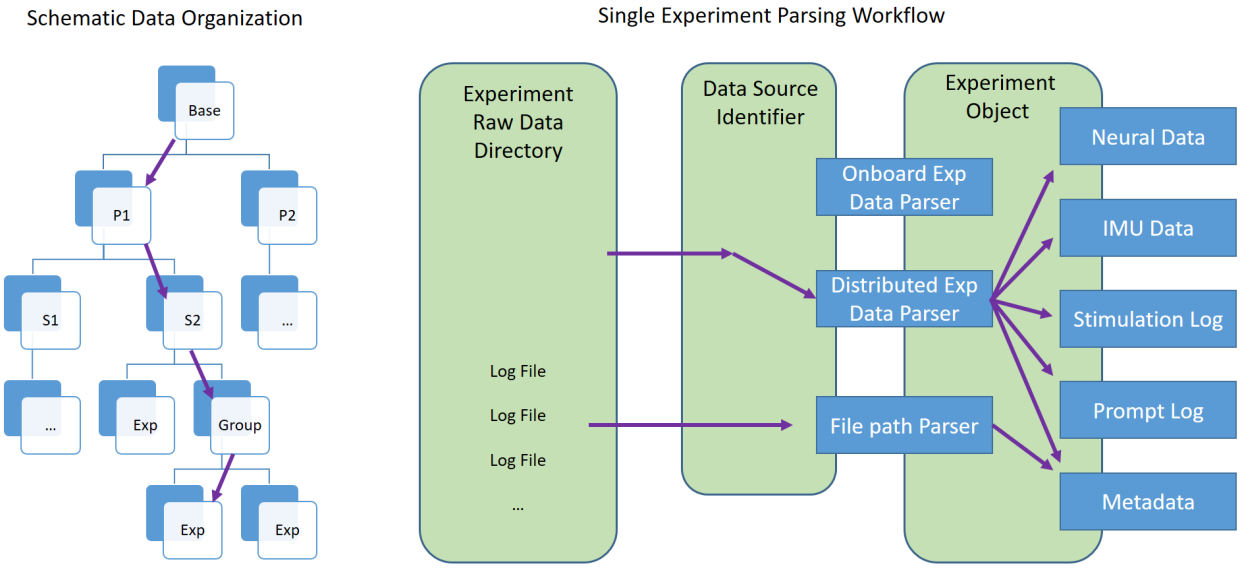

Figure 2.3: Schematic overview of the parsing logic of the retrospective dataset into searchable and standardized Experiment objects that provide reliable access to the data. The directory tree is searched recursively for directories that contain valid experimental data (left). Whenever a directory is found, all the log files are passed into the parsing workflow shown on the right. A source identifier determines the format in which the data was saved determining which parser needs to be used. Each of these parsers is then responsible for decoding the data in the log files into standardized DataFrames. A file path parser collects additional metadata that is not stored in the raw files. All of this data is made available in a uniform way through the Experiment object which enables search and programmatic analysis of the dataset.

is available on Github at: `https://github.com/uw-herron-lab/activa_common`.

2.2.2.1 Data Search and Selection

The core of our implementation is focused on discovery of raw data, parsing into a standardized flat data structure that is amenable to longitudinal analysis and search.

Data is discovered using a recursive search through the directory tree that compares the contents of any directory for all potential end directories that could contain experimental data. This allows the framework to handle any data organization, and effectively extracts arbitrarily nested data into a 'flat' data format. Each of the final directories is passed to the Experiment parsing logic. This begins by comparing the contents of the given directory to a collection of pre-defined sets of required files. Each of these sets corresponds to a particular format in which the data was originally collected or stored. In our dataset, there are two dominant formats, one collected through the distributed pipeline, and one through the fully-implanted pipeline. Although the data available in both of these types of experiments was similar, the format in which the data is saved is entirely different, requiring separate and independent parsing for both types. Each of the two parsers is responsible for ingesting the available data and then populating the fields of an 'Experiment' class, which standardizes access to the data across all storage types. There is also a parallel pile path parsing process which attempts to interpret metadata about the particular experiment from the file path. Information such as the IDs of the patient, session, and experiment name are only available in this way. This allows the data parsing capabilities to be easily extended without requiring any changes to the downstream processing software.

The Experiment class is responsible for facilitating all data search and analysis. When first initialized, the parsers for each modality attached to the particular Experiment instance load only minimal metadata. This allows the user to quickly select the experiments that are relevant to the particular analysis being performed without having to wait for all the data to be loaded from disk. Once further analysis needs to be performed, each requested modality is loaded from disk into reliably structured DataFrames. Since the format of these DataFrames is the same across all Experiments, longitudinal analysis can be easily performed. In addition to the raw data, the Experiment object is responsible for identifying patient and experiment states over time.

2.2.2.2 State Flagging

Any analysis of the long term effects of DBS must take into account the acute effects of DBS and patient states. Therefore, the Experiment object provides several functions for converting logged data into epochs of when the patient is in a particular state. Moreover, there are functions that, given a particular timestamp, will return a description of the patient state. These are all implemented independently for each dimension of patient state, as described below.

Stimulation Epochs DBS amplitude is the only stimulation parameter that is varied in the present data set within the course of a single experiment. Therefore, it is chosen as the primary method of identifying DBS state. Stimulation amplitude is extracted from a log of amplitude changes, and converted to an independent time series. For most experiments in the present data-set, stimulation was binary: either off or on at the maximum value, with ramp-up and ramp-down periods. These ramping periods are essential to prevent discomfort to the patient. For ease of use and presentation stimulation was classified as 'on' whenever the stimulation amplitude was above the threshold of 100 millivolts, and off otherwise. This patient state is available as either a binary description, a list of stimulation 'on' epochs, or a full vector of stimulation amplitude over time.

Movement Epochs Some of the experiments included prompted patient movement. During such experiments, patients were shown a screen with one of two instructions: "Rest" or "Hand". Patients were instructed to sit, relaxed, in their chair when "Rest" was displayed on screen, and to hold their arm out in front of themselves, at around shoulder height, when "Hand" was displayed. This action was effective at eliciting tremor in most of the patients, and was therefore used as a test condition for aDBS algorithms. All prompt epochs were saved in parallel with the neural and IMU data. These are used as ground truth for determining patient movement. The

first two seconds of each prompt interval, when the patient is actively reacting to the prompt and extending their arm, are labeled as "Starting". The two seconds after each prompted movement, when the patient is returning to rest, are labeled as "Ending". The remaining time periods are labeled as either "At Rest" or "Moving". This state is made available primarily as a list of movement or rest epochs depending on the function arguments.

2.2.3 Lessons Learned

One of the main challenges we faced when working with the system is performing longitudinal comparisons using data that was not necessarily intended for this purpose. Over the course of the data collection period, the patients we worked with participated in a variety of different experiments, outlined in section 2. Several of these experiments were repeated at very different time points, which in theory could allow for longitudinal comparisons across these timescales. However, as the development of aDBS progressed, the details of the experiments were changed. For example, the prompted movement evolved from a nose to target and back movement task to a task where the patient simple held their arm outstretched. Both these tasks were very successful in causing tremor in the patient and very useful for evaluating the aDBS algorithm in testing at the time, but are difficult to directly compare to each other. As a result, we had to be very selective and cautious about the data we included in our comparisons and were very limited by the kinds of metrics we could employ.

This challenge was all the more exacerbated by the incomplete metadata collected with these experiments. The prompted movement task is again a prime example of this challenge. Experiments of this sort were always labeled as 'prompted movement' but the exact kind of movement the patient was prompted to perform was not explicitly logged anywhere in the dataset. This had to be interpreted from conversations with the previous students, interpretations of the papers written, and by scanning through old lab notes. Not only was this process time consuming and ineffective, but it

yielded an imperfect reconstruction of the task used for each experiment. A similar example occurred with the aDBS software algorithms used. During each experiment, the student running the experiment was careful to log which aDBS algorithm and approach they were testing at the time. However, these were only logged as algorithm names. And although the broad strokes of each of the algorithms was consistent over time, there were subtle changes between each experiment. Decoding functions changed form, parameters were adjusted, and outputs were tweaked. In some cases, these changes were significant enough that the behavior of the algorithm changed drastically. Without a way to compare the algorithms between these iterations, it was difficult to compare different algorithms in a longitudinal analysis of aDBS.

2.3 V2: Prospective Weil Dataset

Based on the learnings from the retrospective dataset we set out build a more robust dataset in our new project. This dataset would be more challenging as it involved distributed data sources, outside of the clinic, all happening without direct researcher oversight. However, it also presented the opportunity to test our design principles to be able to effectively collect and use a uniquely rich and powerful dataset.

2.3.1 Data-Net Ecosystem

The DataNet Ecosystem is a collection if interacting software packages designed to facilitate remote chronic collection of multimodal data for neuro-modulation research.

2.3.1.1 Particular Requirements

In addition to fulfilling the basic requirements of data collection and management outlined in the introduction, our software for this application would face an additional challenge: all the data collection would be taking place remotely in the patient's home without direct researcher oversight. This particular stipulation meant that all of the

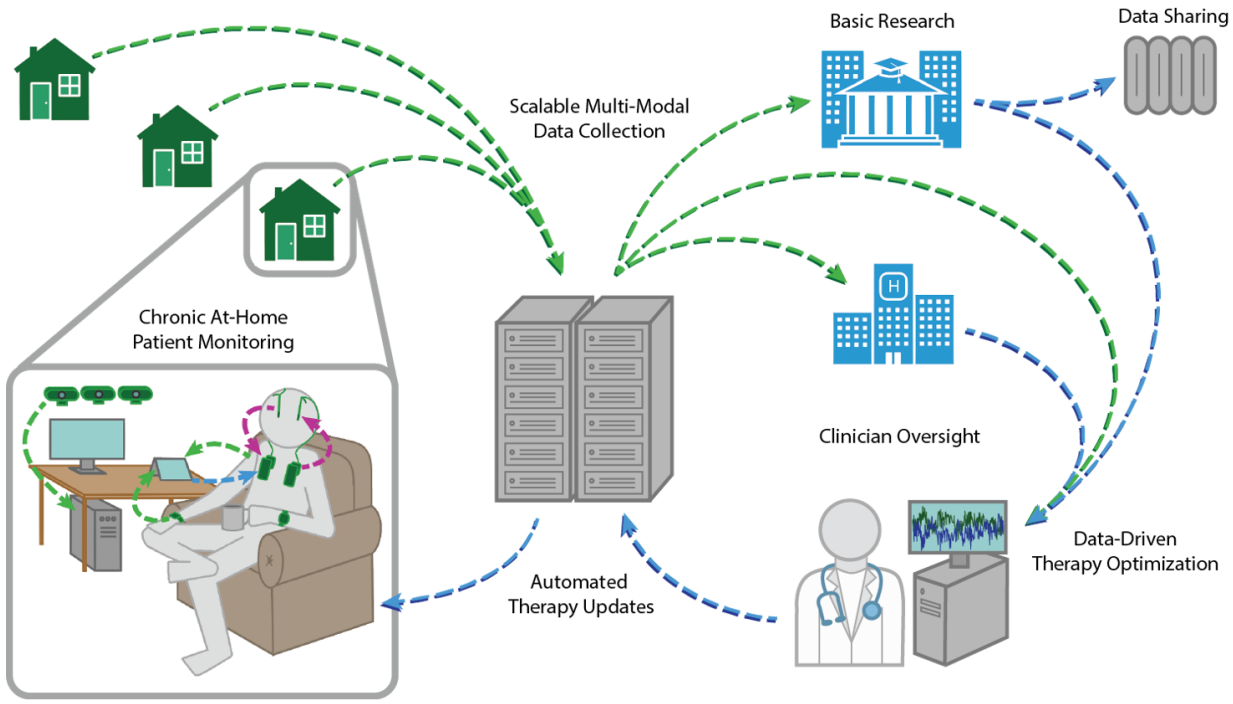

Figure 2.4: Overview of the DataNet remote data collection platform. This platform was designed to be scalable and provide chronic monitoring of patient data while robustly handling the practical considerations of a research and engineering environment. Data is collected from the patient in their home environment and collected on a central cloud storage location. From there it can be used by both basic researchers and clinicians to provide optimized automatic therapy updates back to the patient.

existing requirements would be more difficult to meet. It also implied two additional requirements: complete remote control and simplicity of usage.

Deployment in the patient's home means that limiting impact on the patient's life is important. After the initial deployment, our system can operate without requireing regular access to the hardware. All of our deployed computers are remotely accessible at all points in time. To accomplish this, all computers have either an ssh server or some other remote access software such as NoMachine running at all times. Addi-

tionally, in the event of an unexpected shutdown, the system can recover and restart with minimal input from the patient. Everything is automated to set itself up on reboot, so in the event of a catastrophic shutdown, we only need the patient to press the power button and the system can restore function.

All our systems are easy enough to use that elderly patients without technical skills can still incorporate them into their daily routine. This is closely related to the low impact challenge. We asked our patient to regularly participate in recordings. Any additional technical challenge in doing this would make any patient hesitant to participate in the study. Therefore, regardless of how complicated the background processing is, the interaction of the patient with our systems to start and end recordings is made to be as clear and simple as possible. That means that all the devices should connect automatically, and that recordings can be started with few button presses. The patient also receives clear information when recordings are in progress and is given an easy way to terminate recordings, especially video.

2.3.1.2 Software Sections

To ensure a balance between code flexibility, code reusability, and efficient deployment of the different services, the code is split up amongst several related code repositories. These are all managed with git, using the git-flow repository management scheme [144]. All the repositories are hosted on github, and all data is generated with versions of the code that directly correspond to commits available on GitHub. This allows us to track the exact version of the code used to generate any piece of data by simply logging the repository URL and the git commit hash.

Data-net-common This repository is designed to be a way the different repositories and related services to share commonly reused functions. These are generally utility functions, and this repository is a dependency for all the other repositories in the data-net-ecosystem. This package is available on github at: https://github.com/Weill-

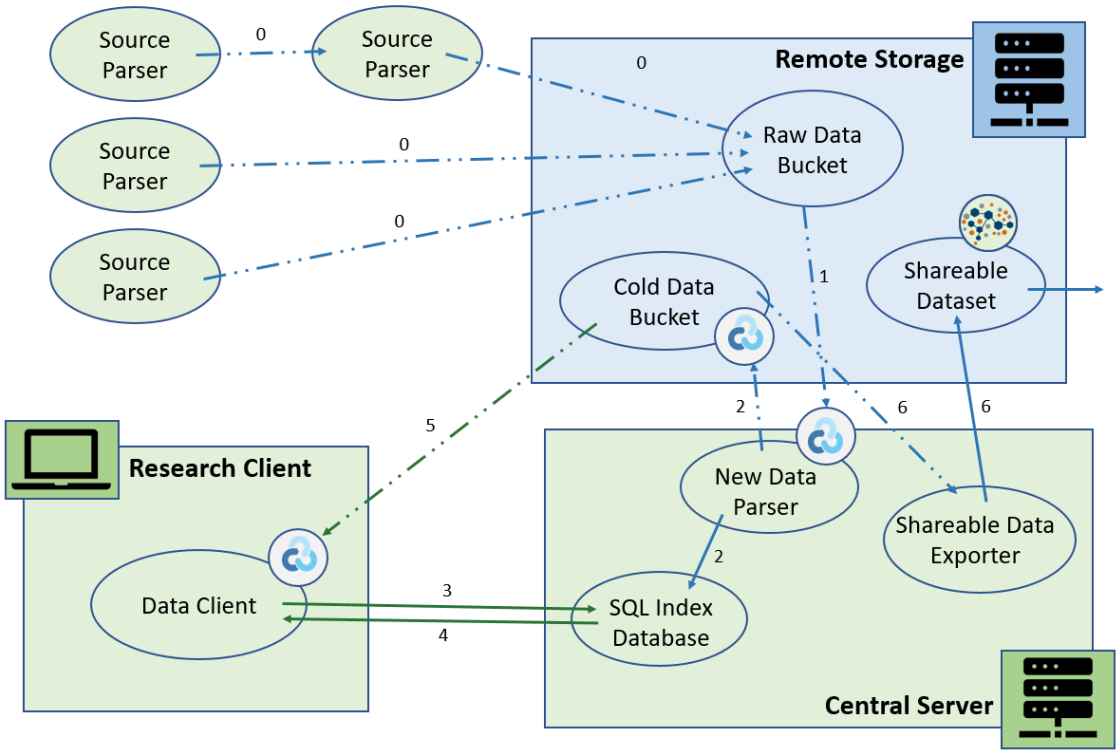

Figure 2.5: Software overview figure, which shows how the different components of the Data-Net framework interact. Green boxes show software running within our VPN, where the blue boxes show the cloud outside of our VPN. Dashed lines show explicitly encrypted data transfer. Data for each modality is independently collected by the source parsers, and is transferred to the raw data storage (0). This bucket is regularly queried by the new data parser (1). When new data is found, the new data parser processes it into a time-aligned and more usable format, as well as indexing it in the SQL database (2), where it can be accessed by researchers. When a researcher would like to perform analysis, they can query the database with a set of conditions (3), and receive back a list of all matching time ranges in the data (4). They can then easily load only the portion of the data relevant to their analysis (5). Whenever data needs to be shared, a de-identified NWB copy of the data can be generated (6).

Neurohub-OPTiMaL/data-net-common

Data-net-source This repository is a source for a common framework of functions to collect data from a remote device and move it to a the centrally coordinated location where it can be integrated with other data streams. Each instance of one of these 'SourceParser' classes is responsible for monitoring and collecting data from a single device. For example, one of our source parsers run on the video collection machine and is responsible for regularly uploading all recently recorded videos to the remote storage endpoint. Different instances can use different transfer protocols but most of the parsers in our implementation use reliable and secure data transfer. This package is available on github at: https://github.com/Weill-Neurohub-OPTiMaL/data-net-subject

Data-net-server This repository is the main workhorse of the live data packaging and data analysis aspects of the data-net ecosystem. It is responsible for polling the outputs of the data-net-source repositories, collecting and aligning data from all modalities into a live compacted dataset, provided a data lookup index, and exporting the data into a final shareable output format. These three primary functions are broken into the three primary modules: storage, index, and export. This package is available on github at: https://github.com/Weill-Neurohub-OPTiMaL/data-net-server

Storage The storage code is organized first by source group and then by specific modality. Each modality has a dedicated Reader and Writer class which inherits from a base DataReader and DataWriter. This inheritance architecture is used both to help with code reusability, as well as to enforce a strict architecture for the data. Both the Reader and Writer for each modality also inherit from a common Spec class, which contains common setting specific to that modality.

The modality specific Readers as generally pretty simple child classes of the base DataReader. The base DataReader is responsible for providing easy access to the most common and universally required metadata, such as start time. It is also responsible for managing loading of all the HDF5 datasets from the data store. This consists both of functionality for preparing slices of the dataset both along the time and data dimensions, along with the ability to read the data that corresponds to those slices into python memory. The reader can then return either the raw data, or package the data into a neatly labeled DataFrames, along with all the metadata attributes of the modality packaged into a python dictionary.

The modality specific Writers are much more complicated by comparison. The base DataWriter they all inherit from implements the basic code that all the Writers use to write the prepared and collected data to the HDF5 data file. It also defines the external facing functions and the two main function hooks that the modality specific Writers must implement to prepare the raw data into a format ready for writing.

These two are the prep_data and the prep_attributes functions. The data for these two functions is drawn from the folders associated with the particular modality on the particular day. The exact preparation process depends on the modality but does follow an set of common requirements that allow for increased reusability of the code.

Index The index module of data-net-server is responsible for creating and maintaining a SQL indexing database of all the available data. This database is intended to facilitate lookup of the data by providing a minute-by-minute table of what data modalities are available when, along with any other custom-defined columns. To facilitate this process, all the stored data is saved in minute-long chunks. This enables the indexing process to open the data for each minute individually, without the need to load all the data for that modality into memory. This is accomplished through the RowReader objects, which are wrappers around the regular CombinedReader,

specifically intended to load the raw data from HDF5 for a narrow time slice.

Data-net-client The data-net-client is the primary component of the data-net ecosystem that is intended to be installed on the researcher's computer. It is designed to interface with the services running on the server and facilitate the lookup of data for the researcher. This package is available on github at: https://github.com/Weill-Neurohub-OPTiMaL/data-net-client. The intended workflow is that the researcher formulates a SQL query that describes the data that they need for their analysis. The data-net client then queries the index to find all matching rows of data. The lookup data from these rows can then be either returned directly, or can be converted into a contiguous interval format. This function looks for all time-adjacent sets of rows returned and converts them into a lookup format that retains the address to the data file along with a time range that fulfills the original query conditions. Last, the user can call the fetch function with these lookup ranges to get a list of pretty-formatted DataFrame for their desired modality for all ranges.

2.3.1.3 Hardware Components and Networking

Here we review the hardware which was used in the deployment of the software described above.

Central Storage Essential to the creation of this inter-connected data infrastructure is a central remote storage where all the raw data recorded by each of the devices can be independently be uploaded. We chose to use Wasabi, and s3-bucket based remote storage service. Using a service like this has several advantages in regards to data safety. We use wasabi for several purposes. Each parser has a dedicated directory where it uploads the results of its ingestion processing. The central processing and packaging services can then monitor these directories to source the raw data to be packaged into a ready-to-analyze format.

Summit RC+S The Medtronic Summit RC+S is a research platform developed by Medtronic and available through IDEs. (Stanslaski et al., 2018). The Summit RC+S system (often referred to as just 'RC+S') can record from 4 channels at 1 kHz, along with onboard digital spectral power estimations. Channels can be configured to perform differential recordings between any pair of electrodes, with 4 electrodes on the stimulation lead and 4 electrodes on the attached cortical strip. The implanted pulse generator (IPG) of the summit also includes an accelerometer which streams data out along with the neural data. The Summit also supports 9 state state table for switching stimulation parameters in a fully embedded fashion based on two independent LD classifiers on the power band estimates. Data from each RC+S is streamed via medical RF to a relay device worn in a specially designed vest, which then relays the data via Bluetooth to a tablet that can finally relay the data to a central server for further processing [49].

Our patient had two RCS+S devices implanted, one on each side of the body. The depth electrodes were placed so they could apply stimulation to the sub-thalamic nucleus (STN) and the electrocorticography strip was placed over the motor and sensory cortices, roughly overlapping the hand/arm motor areas. The stimulation parameters were set by a neurologist according to the clinical standard of care. In addition, the patient often used their patient controller to adjust their stimulation over the course of the day to manage the effects of medication. The patient participated in several other studies through the time period they were participating in our study.

Data from the Summit RC+S is streamed to the attached Medtronic research tablet. This tablet runs medical-grade software written by the team at UCSF that is overseen by the IRB. This software allows the patient to easily swap stimulation settings, and stream all data out to a specific remote endpoint, in this case a UCSF HIPPA Dropbox. This data is streamed out directly in the form of the raw .json files as specified by the Medtronic research API [49, 144].We then have a source parser regularly polling this Dropbox endpoint looking for ready data files. The source parser

is responsible for downloading these raw file, processing them through the OpenMind RC+S pipeline to turn it from json data packets into easy-to-analyze .csv files. In addition, this software also handles some of the time alignment offsets that occur between packets. These files along with the remaining json metadata is uploaded to our central storage endpoint.

The data from the RC+S device included several modalities of data all synced to the same onboard clocks of the RC+S device. This includes 4 channels of neural data, which can be from wither the depth or the surface electrodes. The electrodes used along with the onboard analog filter prams used in the recordings are stored as settings files. In addition, the RC+S can record on-board power-band estimates, and the linear discriminant classifier states used in the control matrix. These are also saved as .csv's. Stimulation settings, either as constant open-loop stimulation settings or constantly fluctuation adaptive settings are also saved. In addition to neural and onboard data, each RC+S device has an 3-axis accelerometer included in the chest-IPG. This data is streamed out at <rate> and is also saved.

Apple Watches Our patient wears two apple watches which enable recording of motion and symptom data. These watches are incorporated with the Apple Health-Kit systems. Through a partnership with RUNE-Labs and the StrivePD app, the data from both apple watches is uploaded to the RUNE-labs remote platform via the patient's smartphone. Here both the raw data and processed .csv's are available for download, wither through the web interface or through the API. Our source parser periodically queries the rune-labs API and uploads prepared csv files to the central wasabi server.

The Apple Smartwatches provide several modalities of data. Primarily, they have full 9-axis inertial motion units (IMUs) with provide acceleration data. There are 3 axes of raw accelerometer data (which include the gravity vector), 3 axes of gyroscope data which provides angular velocity information, and a computed 3 axes

of accelerometer data with the gravity vector removed. Based on this, the RUNE software calculates estimates of tremor and dyskinesia probability and severity. In addition, the apple watch contains a LED against the skin which measures the patient's heart rate every few seconds. All this data is queried and saved as .csv's for further processing.

Cameras We Have 3 cameras deployed to the patient's office to provide ground truth monitoring of the patient's movement. These are Logitech Brio cameras, connected to the recording computer with USBs. To ensure that the cameras are able of providing a robust stereoscopic view of the patient they are positioned at least several meters apart, and are connected to the computer with 3m long USB cables. Full details of the camera setup are available in [144] Recordings are managed though a custom written application. This application serves several purposes: it manages the recordings, and informs the patient of the recording status. Recordings can be scheduled ahead of time using .json config files that are automatically polled from the application Git repository. In addition, recordings can be started manually remotely. The recording status is displayed to the patient on a touchscreen. The screen also contains several large buttons that allow the patient to pause or completely stop the recording at any time. In addition, the most computationally intensive process was the processing of video into Pose. To handle this workload, we found we needed a dedicated computing machine with powerful GPU. This machine was responsible for taking pre-recorded video from the remote data storage and converting it into pose data. This process was significantly slower than data collection requiring asynchronous processing.

Deployed Computer To oversee video recordings and enable manual control of the recording system, we deployed a desktop computer to the patient's home. This computer was made remotely available at all times. More details of the computer

	Central Server	**Deployed Computer**	**Pose Processor**	**RCS Input Processor**	**Research Computer**
Processor	Inter i7-3770 @ 3.40 GHz	Intel i7-10700 @ 2.90 GHz	Intel i7-4790K @ 4 GHz	Intel i7-10700 @ 2.90 GHz	AMD Ryzen 9 3900X
RAM	8 GB	32 GB	32 GB	32 GB	64 GB
Storage	180 GB SSD	512 GB NVMe	500 GB	512 GB NVMe	250 GB SSD
	1 TB HDD	2 TB HDD	1TB HDD	2 TB HDD	2 TB HDD
GPU	—	GeForce GTX 1650 Super	GeForce RTX 3080 12GB	—	NVIDIA GeForce GT 710

Table 2.1: The specification of all the machines that significantly contributed to the deployment of our version of the DataNet ecosystem. Specific hardware requirements are dependent on the data collected.

setup are available in [144].

Central Server This central server is the cornerstone of all active components of the data management environment. It manages the VPN that all research machines are connected to, runs the software combines all the data into a ready dataset, and hosts the data-index database that research computers can query. In our case, the central server was a desktop computer that we dedicated to the task.

2.3.2 Practicalities of Implementation

A sample of the data collected during an example week of recordings is shown in Fig X. Multiple data streams were collected in parallel from all the source devices. A variety of experiments were performed during this time. In this section, we review the

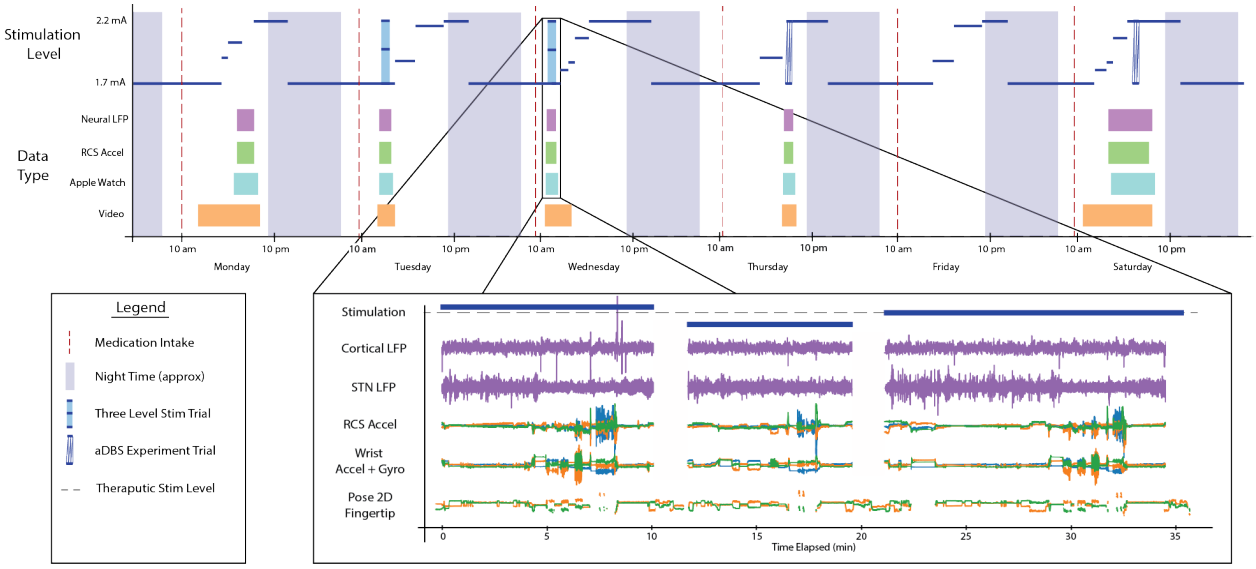

Figure 2.6: Summary of the simultaneous data collection of the system over one example week of recordings. The overlap in data collection for the different modalities is shown in the schematic on top. An expanded view, for a sample of the multi-stim dataset and a subset of the available channels, is shown in the bottom panel.

total data that was collected and review the benefits to analysis the DataNet system permitted.

2.3.2.1 Data Collected

Data was collected from the patient's home during a variety of different experiments. A sample schematic of a week of data collection showing all the modalities collected simultaneously is shown in Fig 2.6. These are described below:

Multi-Stim This experiment series was designed to investigate the effect of the three different stimulation levels. It also served as a way to calibrate future aDBS paradigm development and as a way to disentangle the effect of DBS stimulation from other effects. In this paradigm, complete multi-modal data was recorded when

the patient performed a series of prescribed tasks at a variety of different stimulation levels. The stimulation was applied in an open-loop, or continuous mode, at a low, normal, and high stim level. Stimulation was set to 50%, 100% and 110% of the patients normal preferred stimulation amplitude. The order of these stimulation levels was different every day, and all possible orders were iterated through. During the experiment, the patient turned all their data collection devices, and perform the prescribed series of tasks. After completing the task set of stimulation the patient would take a break for a minute or two and then press a button to advance to the next stimulation setting. The tasks were inspired from the standard clinical tasks used as standard of care in clinical practice, combined with tasks that could more easily generalize to the patient's daily life.

Move aDBS This was designed to provide a proof of concept aDBS system and demonstrate the effectiveness of aDBS in the chronic at-home environment. Details of this paradigm are available in full in [30] but are reviewed here for completeness as well. This data collection paradigm was very similar to the Multi-Stim experiments, except for the stimulation applied. Stimulation was applied as cDBS at the normal amplitude, as an aDBS algorithm that stimulated during movement, and as an 'inverted' aDBS algorithm, that stimulated during rest. Movement was decoded by an onboard neural classifier specifically trained for this purpose. A battery of clinically-inspired tasks was performed by a patient in each of the three stimulation conditions.

Naturalistic Behavior This data collection protocol was designed for us to be able to generalize our results towards normal daily life. During this protocol, the patient was asked to proceed with their regular daily routine. Data from all modalities was streamed, and yielded valuable data any time the patient was in the instrumented room, which in our deployment was the patient's work-from-home office. Data collection was considered to be opportunistic, and no strict regime on the start or ends of

recordings was imposed on the patient.

2.3.2.2 Data Inventory

The multimodal data collection of this project allowed us to compile a one-of-a-kind dataset of the effects of aDBS on the patient in their home environment. A total of 110.2 hours of neural data on each of 4 recording channels were collected over 83 days. Concurrently, we collected 187 hours of smartwatch data, 90.4 hours of 2D pose data from each of 3 video cameras. To our knowledge, this is one of the largest at-home datasets collected for a single patient, with some of the best temporal distribution. The limitation of this dataset however, is that it only contains data from one patient, so conclusions drawn from it will not necessarily generalize to the rest of the population.

2.3.2.3 Additional Challenges Faced

Time Alignment Time alignment is a common concern with data collection that includes many different simultaneous data modalities. As a result, we were very aware that this could be a concern. However, all our devices were internet connected and capable of remaining synced to GMT. Moreover, since the time accuracy needed for our analysis is on the order of several tens of milliseconds, we believed that internet time synchronization should be sufficient for our purposes. However, early on in our investigation, we noticed that some of our time streams, particularly the apple watch data, could sometimes be offset by as much as 20s. The exact cause of this offset is not entirely clear due to the proprietary-closed source nature of the apple watch and RUNE software, but we tracked it down to intermittent Bluetooth connectivity between the apple watches and the patient's smartphone. Additionally, due to scheduling and startup delays on the video recording computer, the video data would often have time offsets on the order of 1 second. Clearly, we would need some sort of time-alignment step in the post-processing after all.

In theory, both data streams should be grossly correlated to each other, since large movements should cause a signal in both. This correlation can be dramatically improved by having the patient perform a task that is intended to be clearly visible on both time streams. In this case, we asked the patient to firmly and a-periodically tap their chest with both closed fists several times. This creates several spikes that were visible in all the accelerometer time streams as well as being visible in the pose data.

In keeping with the automation motivation of the system, we developed a cross-correlation based method of verifying the time alignment. The two accelerometer data streams being aligned are normalized across all three axes, using the L2 norm. This removes the directionality information, but preserves data about the intensity of the movement, making the comparison between the two time streams more reliable. We then calculate the normalized cross correlation between the two resulting time series. This is a metric closely related to the standard cross correlation, but it is normalized by the mean and standard deviations of both time series. This ensures that the resulting values are bounded by +- 1, where value of 1 indicates that the two timeseries are identical, whereas a value of -1 indicates that the two time series are exactly the negative of each other. A value of 0 indicates no correlation. Since we do not care about the relative phase information between the two signals we can take the absolute value of the cross correlation value. We then compute the XCorr value for a variety of different relative time lags. Now, the time lag with the largest XCorr value should correspond to the time offset between the two time series.

The method for ZNCC works well in some cases but it has a few critical vulnerabilities. For example, for some movements, the two accelerometer signals can be phase shifted with respect to one another. If a phase aligned and phase shifted movement are both included in the analyzed epochs, then the ZNCC can have either multiple or even no clear peak. The normalization of ZNCC allows us to automatically identify and discard these alignments. This method works best if both signals are relatively

noise free and windowed to an epoch with large, synchronized effects in both traces. These problems were significantly improved if only the chest tapping interval was included in the time alignment process. However, this required that the time alignment interval was known ahead of time, which could be achieved in one of two ways. Manual inspection of the data post hoc works very well, but at that point we might as well just use the time alignment GUI to do the alignment directly. The other option is to somehow mark the time near the time-alignment taps during the data recording. For the context of this study, this could be accomplished by having the patient press a button right before they perform the chest taps. This method would be recommended for prospective datasets as it would allow near-complete automation. For our purpose however, and given the amount of data that had already been collected, it was easiest to simply us the time alignment GUI to perform the time alignment. This method was fast and flexible, and could handle all the unexpected one-off scenarios we encountered.

To facilitate manual inspection and time alignment of the raw data streams we developed a manual time alignment GUI. The appearance and function of this GUI is summarized in Fig 2.7.The main component of the GUI is a large plot of all the time series being currently aligned plotted on top of each other, with time on the x-axis, and z-scored data amplitude on the y-axis. One timeseries is considered to be 'true' time, and the other data series are aligned relative to 'true' time. The x-axis can be zoomed in and out by pressing the 'x', and 'z' keys, and the entire view window can be shifted left and right with the corresponding arrow keys. The aligning data series can be shifted left and right by holding 'Shift' while pressing the arrow keys. The size of both shift kinds if shift depend on the x-axis scale: each shift is 20% of the horizontal window. If finer shifting is necessary, you can either zoom in, or hold 'Ctrl' to enter fine adjust mode, and shift by a single sample interval on the true time data. The up and down arrow keys will scale the aligning time series relative to the true time series up and down. 'Q' and 'E' will scale the y-axis without adjusting the scaling of the

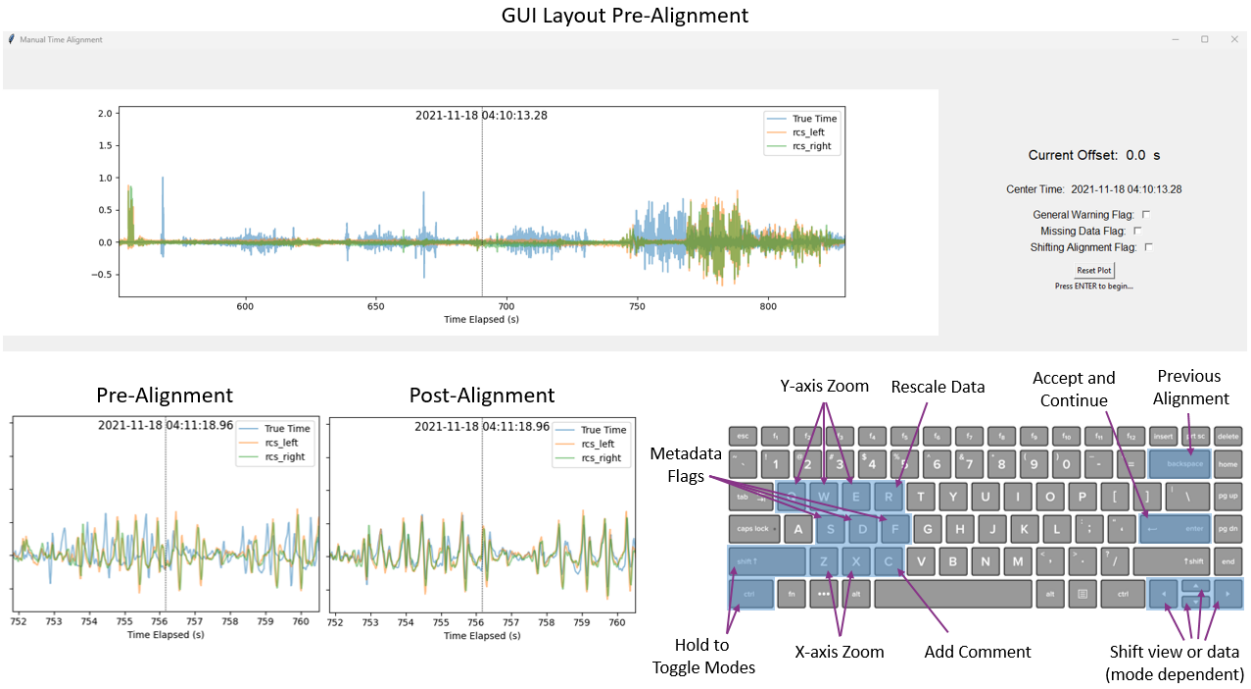

Figure 2.7: Overview of the GUI used to perform time alignment between the different modalities. The top row shows a screenshot of the GUI at the start of aligning the two RCS accelerometer timeseries (orange and green) to the left apple watch accelerometer data, (blue) which is considered to be 'true' time. The plots in the bottom left show a zoomed in example of the timeseries before and after alignment. The keyboard layout, in the bottom right, highlights that the GUI is designed to be used entirely with the keyboard, allowing for rapid navigation and alignment.

data. Automatic z-score based rescaling of the currently visible data and y-axes bac ne achieved by pressing 'R' and 'W' respectively. The plot can be reset to its initial state by pressing 'Delete'. Several metadata flags related to the data and alignment quality can be toggled by pressing 'S', 'D', and 'F', and a custom comment can be entered by pressing 'C'. Once alignment is satisfactory, pressing 'enter' to save the time alignment and move to the next time series to align. The time alignments are returned

from the GUI as a dictionary which can either be further processed or saved to json. Each time-series being aligned must be a one dimensional vector over time. Therefore, multidimensional data, like the accelerometer data we use, must be somehow collapsed into a 1-d vector. In our case, we use an L2 norm across all the components, and align based on the amplitude. Complete usage instructions and the code is available on github (https://github.com/Weill-Neurohub-OPTiMaL/ManualTimeAlignGUI).

2.3.2.4 Integration with existing NWB software

Data-Net exists in parallel to NWB. It is a package focused on real-time distributed data collection in a med-device engineering environment. NWB, by contrast, is a fantastic data format for sharing neuroscientific data in comprehensive and understandable way. As a result, it is important that our framework can export the data collected into the NWB format. In addition to facilitating data sharing, this would allow the usage of all the tools that have been developed for analyzing NWB datasets.

2.3.3 Evaluation

We believe that we have successfully, designed, developed and deployed a robust remote data collection platform. We have met the main component of each of the design requirements we specified above, with some minor caveats, described below:

Live Capabilities We never quite reached the fully automated and live goal of the system. Although data was collected reliably, several steps of the data collection process remained partly manual by the time data collection stopped. The biggest of these was the data aggregation step: collecting data which was automatically uploaded to the remote repository and collecting it into a ready dataset. The main reason for this was the difficulty we faced with time alignment: this process remained largely manual, making automation of the surrounding steps largely pointless.

Privacy All the computers in our research network were connected using the WireGuard VPN ensuring that all communications were encrypted. All hard drives where data was stored were encrypted and all computers were password protected. All data stored in remote cloud endpoints was encrypted. As a result, we believe we did a good job ensuring patient privacy as far as system security was concerned. We did however face several potential complications with privacy. In initial discussions with the patient, they proposed placing the recording system in the bedroom. We rejected this proposal due to privacy concerns, and elected the patient's office instead. Even then, despite the fact the cameras were oriented to capture just the patient, we did capture other members of the household on camera. In future systems this could pose concerns, as all members of the household had to consent to appearing on camera, not just the patient.

Robustness The system was able to collect data very robustly. As the patient came in and out of the instrumented room, recordings would drop out and resume smoothly. Data was collected successfully for all the different data collection protocols. As engineering changes to the device and stimulation settings were made, there was minimal disruption to the data collection capabilities. The data organization in the initial remote repository was designed well enough that hardware and protocol changes did not impact the data transfer pipeline. As new data modalities were added, the system was able to be smoothly expanded to accommodate.

Data Lake Management As data collection proceeded and multiple data collection protocols were running simultaneously, it was very important that we kept track of what recordings corresponded to what data. This was done by several parallel systems. We maintained a GitHub wiki where the details of all experimental protocols were available as a reference to all members of the team. We also maintained a google spreadsheet of all the recording IDs along with easy-to-interpret metadata. However,

the most comprehensive source of metadata to be able to manage the dataset was the data-net-index. The biggest advantage of this system was that it was mostly automatic: whenever new data was added to the dataset then the index was expanded to include the new data.

2.4 Discussion

We believe that the data architecture discussed here is capable of facilitating remote chronic data collection. As more sensors move into the patient's home, this sort of framework will be ever more important to be able to collect meaningful conclusions from these distributed systems. In the next two chapters, we review our usage of the two frameworks described above to answer neuroscientific questions.

2.4.1 Ethics Considerations

As we already alluded to, building comprehensive at-home data collection systems such as the one described here comes with severe and wide-ranging ethical consequences. These considerations come primarily in two broad families: those related to patient's privacy, the unintended consequences adaptive stimulation can have on the patient's actions, and the effect of chronic neural implants on the patient's perception of self. A complete analysis of the ethical considerations is out of the scope of this work, but we will also review them here, as they should be an important consideration to anyone developing these systems.

Privacy As has already been mentioned, monitoring systems such as the one described here, are necessarily an invasion of the patient's privacy [3, 149, 146]. Chronic video recordings can capture private behavior of not only the patient, but also of other persons who come into the camera's field of view. As our ability to decode neural data increases, it becomes conceivable that even neural data, which is generally not considered identifiable data, could be linked directly back to individual patients

[**?**]. This data could potentially be then be used to decode many aspects that the patient would rather remain private [16, 9]. Many of these privacy considerations are similar to privacy considerations that arise with the deployment of IoT devices automatic, only more severe due to the more personal nature of the data involved [3, 15]. Research on privacy in existing IoT devices highlights the widespread nature and challenge of solving these problems [149, 60].These risks need to be taken into account whenever these systems are designed and deployed.

These privacy risks only increase as the scale of the deployed system increases and parties with potentially conflicting interests become involved. Privacy policies used by many of the current corporate leaders in IoT devices are not necessarily made to be easily understandable to the user, and privacy policy change notifications lack an opt-out option, forcing the user to either accept or discontinue their usage of the service [122]. This approach to customer's data privacy is already concerning with IoT devices, but could potentially be deadly if continued with medical devices. For example, a patient could be forced to choose between allowing a company to use their private data in an un-acceptable way, and discontinuing the use of a potentially life-saving medical implant. Cases brought before the FTC in recent years have highlighted that companies offering IoT services may have a conflict of interest in ensuring best practices in user data privacy [122, 1]. Research has proposed several ways in which these privacy concerns could be managed at scale with clever system architecture [149, 60], however, it remains to be seen if these recommendations will be followed in commercially deployed systems despite potential conflicts of interest.

Security Security concerns with these sorts of systems follow a similar pattern to the privacy concerns detailed above, but are even more difficult to solve due to the existence of an intentionally malicious third party. However, it seems that this space receives more direct attention from technological developments than privacy. Security violations do not only endanger public release of the patients private data,

but also could allow the malicious actor to directly target the patients well-being [33]. As implantable devices become more profoundly intertwined with the patient these security concerns only grow.

Consequences of aDBS A number of potential ethical and legal concerns arise errors in the function of the implantable device. These concerns are particularly prevalent with aDBS stimulation for motor diseases. This is due to the fact that inappropriate stimulation can have huge effects on the patients motor capabilities with little to no warning. These errors could occur due to internal errors, or even due to unpredictable external influence. Consider, for example, a PD patient with an aDBS device implanted driving a car when a device error occurs. If device suddenly shut off or otherwise stopped delivering clinically effective levels of stimulation, the patient could suddenly begin experiencing severe tremors or experiencing bradykinesia. A difficulty with initiating movement like this could make the patient unable to make a turn in the road and crashing. On the other hand, if stimulation is suddenly increased, the patient could experience dyskinesia. These sudden jerking movements could cause the patient to uncontrollably swerve the car. In both of these cases, the patient could experience significant harm both to themselves and their personal property (ie their car). Moreover, the patient could cause serious harm to other persons on the road or nearby. This example illustrates the severity of the safety problem that needs to be solved as these sorts of devices are widely deployed. One possible solution is a 'freeze mode' that the patient can choose to enable whenever they are performing an activity where adverse motor symptoms could be dangerous. In this freeze mode, the device would disconnect from all non-mandatory networks, and pause the adaptive algorithm. This would mean that stimulation would be delivered in an open-loop configuration, temporarily reducing the clinical effectiveness of stimulation, but greatly decreasing the risk of a malfunction.

Another important consideration of device malfunction are the legal and liability

consequences of such an event. In some ways, these considerations mirror the legal battles around liability in self-driving cars. One can imagine that if the crash was caused entirely by a failure of the device causing the patient to be unable to steer, then the liability would lie with the manufacturer. However, in this case the legal troubles might be even more significant, since it would be very difficult to establish with certainty whether the movement that caused the crash was a failure of the patient to pay attention, or a failure of the device. It could be possible to implement some sort of black-box logging system on the device, but then we face increased privacy and security challenges. A proposal to resolve these issues is far outside the scope of this work, but they are an important consideration to remember when thinking about the future of these systems.

Neural implants and the self The last important consideration of DBS and aDBS that we will review here is the impact of brain implants and stimulation on one's perception of the self [69]. On a societal level, we often identify ourselves and our personality with our mind, and by extension the brain. Therefore, any interference in the way the brain functions, such as a lobotomy or even just DBS stimulation, poses a significant threat to the patient's perception of self. There have been several reported cases where patients receiving aDBS treatment had either significant personality changes, or adverse psychological effects on either mood or self-control [125]. This view however, has been increasingly challenged in recent years [48]. Some patients also experience severe psychologic effects that can be explained by a sudden unpleasant perspective shift from a dualist view where they consider their self to be a disembodied "mind" to a more neuro-centric and materialistic view of their selves[99]. Continued ethical and philosophical discourse in this field is necessary so that we can develop potential life-saving neurological technologies in a safe and ethical way.

Chapter 3

APPLICATIONS TO ESSENTIAL TREMOR

3.1 Introduction

3.1.1 What is Essential Tremor?

Essential Tremor (ET) is the most common neurological movement disorders [91]. By some estimates, it affects as much as 1% of the world's adult population and up to 4.5% of the senior population to some extent [57, 29]. ET manifests itself primarily as a 2–8 Hz tremor during active motion or holding of posture [168, 57]. Classically, the strongest tremor is apparent in the extremities, especially the hands, but will often also be accompanied by trunk tremor [57]. Since many other movement disorders share tremor as a symptom and the progress of the disease is slow, it is generally wildly under-diagnosed [91]. As it is a neurodegenerative disorder, the symptoms tend to increase in severity over time. This usually means that the amplitude of the tremor increases and that new regions of the body can be affected by tremor.

Despite its prevalence, the a lingering misconception that essential tremor is largely benign means that it has been understudied for many years and the pathophysiology is still poorly understood [139, 11, 90, 168, 118]. Studies generally agree that ET is related to pathological tremor frequency oscillations in the 'tremor-network', consisting of the cerebellum, inferior olive, thalamus, and motor cortex [126]. Recent research has also suggested that ET is not in fact a single disease, but may be a family of related conditions with similar symptoms [11, 64, 168, 28]. Of these, the two most common explanations are the cerebellar origin and lewy body origin[152, 10]. The Lewy-body model of ET proposes that the pathology of ET is driven by the accumulation of Lewy bodies in the brainstem, but especially the locus ceruleus [92]. In the

cerebellum-driven model of ET, which has been gaining popularity in recent years, disruptions in the function of cerebellar circuits leads to a resonance that propagates thought the tremor network, all the way to the affected extremities[168, 119]. The exact form of these disruptions is still under investigation but several promising hypotheses have been proposed. These include neurodegenerative reduction of purkinjie cells [28, 10, 118], excessive pruning of synapses between purkinjie and climbing fibers mediated by a geneticly-driven GluRδ2 protein deficiency[119], and dysfunction of the GABAergic system in the cerebellum [51]. Importantly, some have argued that these distinct hypotheses are merely different compensatory behaviors for the primary pathology [28]. Clearly, more research in this field is needed to conclusively determine the pathology of ET and hopefully lead to more effective treatments.

Once patients are diagnosed, initial treatment is usually pharmacological, but for severe, pharmacologically refractive cases DBS is a promising option [29, 91, 93]. Deep brain stimulation (DBS) is a common therapy used to treat neurological disorders. It has been approved by the FDA to treat ET, Parkinson's disease (PD), dystonia, and epilepsy; and is under investigation for treatment of depression, addiction, Tourette syndrome, and many others [107, 93]. For ET, the DBS lead is usually implanted into the ventral intermediate nucleus of the thalamus (VIM) [93]. The mechanism of action of aDBS is not fully understood, and is likely to be slightly different for different diseases and different implant locations [107, 159]. However, for ET specifically, it is well accepted that DBS disrupts synchronous oscillations propagating and amplifying thought the tremor network [29, 126, 93]. Despite the fact that the mechanism of action is still under debate, DBS demonstrate very high levels of effectiveness in suppressing symptoms [29, 107, 93].

In current clinical practice, conventional or continuous DBS (cDBS) is used in an open-loop fashion [91, 93]. Stimulation is configured manually by a clinician and the applied stimulation pattern is fixed . Parameter tuning is a lengthy process that, even with the expertise of a neurologist, may require several visits before a satisfactory

setting is found [55, 100]. Optimal stimulation settings are those that significantly suppress tremor, without causing intolerable side effects. The patient is provided a "patient programmer" that they can use to turn the stimulator on or off, but this control is rather coarse at best and used primarily to conserve battery at night while patients sleep. As a result, stimulation is often delivered even when it is not necessary, which may unnecessarily increase exposure to side effects of stimulation[100].

3.1.2 Adaptive Stimulation for ET

Adaptive DBS (aDBS) offers to solve many of the limitations of cDBS systems. In this approach, stimulation is delivered in a closed-loop format that allows the system to adapt to the patient's state. Stimulation can be applied only when necessary, thereby reducing side effects while maintaining clinical efficacy. Since the stimulation could adapt to the severity of symptoms, stimulation could theoretically always be delivered at the optimal level. Moreover, recent evidence suggests that intermittent stimulation may be more effective at suppressing symptoms than cDBS and prevent accommodation effects [87, 85, 44, 170]. ET is a particularly attractive application for this approach since the primary symptom, tremor, manifests itself almost exclusively during movement [57, 168]. This clearly defines the periods when stimulation would be the most beneficial, greatly reducing the complexity of the control problem to be solved. It is worth noting that naming several conventions exist, with adaptive, closed-loop, and responsive DBS having overlapping definitions. In this work we use adaptive DBS as an umbrella term to describe the various ways in which we have automatically adjusted stimulation based on biomarkers of the patient's state.

3.1.2.1 Initial evaluations of aDBS for ET: The Activa PC+S Study

When the aDBS research for ET was started at UW, aDBS had been demonstrated successfully in patients with PD. However, there was only one known attempt at developing aDBS for ET. In that study, the authors used motion detected through

an EMG system on the patient's arm as a control variable for turning DBS on and off [170]. In this section, I review the development process of aDBS in ET patients carried out at the University of Washington. I begin with an overview of both the hardware and software in the research platform we developed around the Medtronic Activa PC+S. I then outline how the UW team developed an EMG-driven aDBS demonstration to a proof-of-concept BCI aDBS system using cortical LFPs as a control signal. By the end of the study, a clinically translatable, fully implantable aDBS paradigm, accompanied by a largely automated programming process had been developed. Consequently, aDBS seems to have now reached the threshold where it could be evaluated as a clinical therapy to improve patients' lives. Throughout this initial process, interviews with the patients were conducted to assess their level of comfort and gain a greater understanding of the patient experience. I addition to providing an overview of engineering and research process involved in developing aDBS for ET, this section gives a detailed description of the experiment that provided the data used in my analysis later in this chapter.

Hardware Used The central component of the research system was the investigational-use Medtronic Activa PC+S, used with FDA permission under an investigational device exemption (IDE, clinical trial number NCT02443181) [141]. This device consists of a pulse generator implanted (IPG) in the chest which controls both stimulation and sensing capability. The IPG was connected via a subcutaneous clinical lead extension to the stimulation and sensing electrodes. For this protocol, stimulation was delivered using a clinical standard four-electrode DBS stimulation lead, the Medtronic Model 3387, implanted into the ventral intermediate nucleus of the thalamus (VIM). Sensing was performed utilizing a Medtronic Resume-II four contact strip electrode placed on the surface of the cortex, spanning the central sulcus, roughly over the hand motor area. This configuration allowed for standard tremor-mitigating stimulation to be delivered to the VIM while also allowing the sensing of cortical local field potentials

(LFPs) related to hand motor activity. The IPG supports both cDBS and aDBS. cDBS can be configured using the Medtronic 8,840 clinical programmer. aDBS can be performed either in a distributed fashion with control decisions made outside the device, or in a fully implanted fashion with stimulation decisions made onboard after configuration using the Activa PC+S Sensing Tablet.

For distributed control of stimulation, the IPG can be paired via a short-range inductive connection with the Medtronic Nexus-D or Nexus-E telemetry bridges. The choice of sampling frequency for the neural data was largely driven by the hardware specifications. This setup can stream raw LFP data to a desktop computer via a USB connection with a sampling rate of up to 422Hz if streaming from one electrode, or 200 Hz if streaming from two simultaneously. The work presented used the 422 Hz LFP data streams. The benefit to aDBS control of the addition of VIM data was found to not be worth the loss in sampling rate, since LFP data from the VIM was heavily contaminated by stimulation artifacts. However, data was transmitted in discreet packets every 400ms. In practice, the half-duplex inductive link's bandwidth limitations resulted in a small window for stimulation updates to be transmitted to the PC+S without resulting in streamed neural data loss, so stimulation updates in a distributed algorithm needed to be performed with 400 ms resolution.

All neural recordings consisted of differential voltage recordings between pairs of electrodes. In this use-case recordings were selected to utilize pairs of electrodes that lay on opposite sides of the central sulcus. These were identified as the pair of electrodes with the highest beta-band power while the patient was at rest, determined through a standardized montage sweep provided by the Activa PC+S instruments. The Activa PC+S can also stream analog estimates of power bands at a sampling frequency of 5Hz, with frequency ranges configured on the proprietary Medtronic Sensing Tablet. When not streaming, the sampling rate of the IPG can be increased up to 800 Hz for raw LFP data, and the resulting LFP or power band data can be downloaded to a USB drive via the Medtronic tablet. In the fully implanted aDBS configuration the

IPG uses the analog power band estimates from the attached electrodes in combination with a simple linear classifier to control stimulation. Both the power bands and the weights of the classifier are also configured via the Medtronic Sensing Tablet

EMG-Driven aDBS Initial feasibility studies began with the implantation of the first patient in 2015. The goal of this initial work was to demonstrate that system integration was successful and could be used to prototype aDBS paradigms, which had never been done before [61, 62]. These initial experiments would validate the system and allow correlation of the various modalities of data to the patient's state. Using this data, development and verification of neural biomarkers of movement and tremor required for neural-driven aDBS began [62, 63]. An IMU-based measurement of tremor that could track fast changes in tremor severity was developed to objectively evaluate the effectiveness of any aDBS paradigm. This would then be a test bed that would enable the identification of biomarkers that would not only reliably distinguish times when stimulation was needed but also be robust to changes in stimulation.
Taking inspiration from prior literature described in Yamamoto et al. (2013) where EMG was used to drive a DBS system using a re-engineered patient programmer, we first implemented prototype aDBS paradigms driven by EMG and IMU signals. Each of these devices was used to monitor the patient and detect whether they were actively moving. When movement was detected, stimulation was rapidly ramped up to the therapeutic threshold and maintained until the patient returned to rest. A representative trial for the EMG system is shown in Figure 3.1. These trials consisted of a comparison of the relevant data during repeated rest, movement, and imagined movement trials. During rest, the patient was asked to simply sit in the chair, while all sensors recorded data to use as a baseline. During movement, they were prompted to raise their hand (at the time of the green vertical lines) and hold it out in front of themselves, until prompted to return to rest (at the time of the red vertical lines). This movement was found to reliably elicit tremor for this patient. Prompt intervals

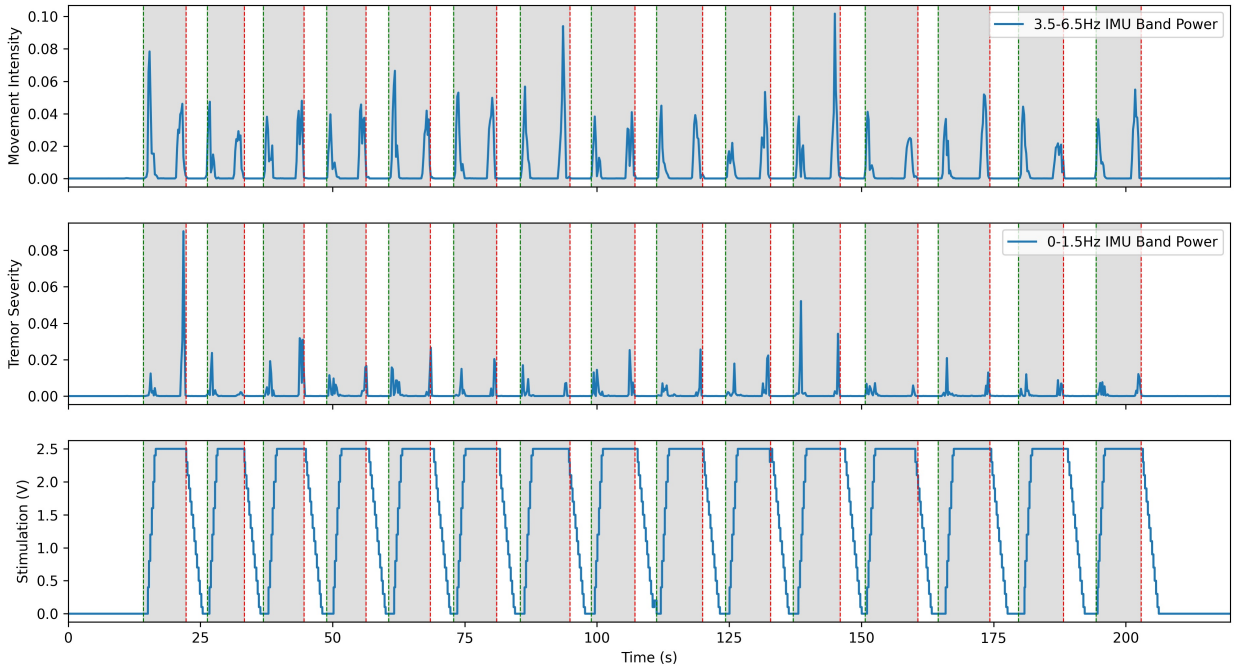

Figure 3.1: Representative example of a EMG-driven aDBS trial performed with one of the UW patients. Movement was detected directly from an EMG armband and used to stimulate only when the patient was moving. This system was very effective at stimulating only during movement and demonstrated the potential of aDBS for ET. However, it required a significant amount of additional tethered hardware that made it impractical outside of the clinical environment

of various lengths, were interleaved so that multiple comparisons could be collected quickly. A similar approach was used for imagined movements, but instead of moving, the patient was asked to instead imagine performing the same movement. This prompt paradigm was used as a template for many of the later experiments through the study. Analysis of the IMU data allowed development of a metric for tremor severity, described below and shown in the second plot below, which could be computed in near real time. This metric correlated to the tremor observed in the patient, based on the

FTM scale, while reducing the movement onset and offset artifacts [63]. By comparing the neural data obtained during these trials we tested whether our system was able to detect beta band desynchronization both during overt and imagined movement. Moreover, these changes were apparent even during stimulation, despite the dramatic changes in the frequency spectrum observed during DBS.

This system fulfilled its primary goal as a technology demonstrator. aDBS triggered by movement, particularly in the case of EMG, was successful in suppressing tremor while delivering less total stimulation. aDBS driven directly by tremor severity interpreted from IMU data was less successful, due to feedback causing the stimulation to fluctuate wildly. Beta band desynchronization was shown to be reliably identifiable with the hardware available, and therefore a potential control variable for future aDBS systems triggering off of movement-related biomarkers.

Although this initial work showed the potential of our system as an investigational device and the promise of aDBS for ET, it also highlighted many of the challenges that would need to be resolved over the rest of the study. Conversations with the patients exposed their reticence to undergo battery replacement surgery, thereby highlighting the importance of conserving battery power. Since streaming neural data used approximately 10 times as much battery as normal operation, experimental sessions were kept succinct and avoided unnecessarily draining the patient's battery. Moreover, using aDBS to minimize the energy usage of stimulation would be an important consideration throughout the rest of the project. The delays inherent in the distributed Activa PC+S system made the control scheme for distributed aDBS difficult to implement. We often observed transient periods of significant tremor at the onset of movement, before the aDBS control caught up and turned-on stimulation. Solving this issue would be one of the main targets we would pursue.

Distributed BCI Control A neural driven, BCI aDBS system was built informed by the above results [61, 63, 65, 22]. A system like this, driven by the well-documented

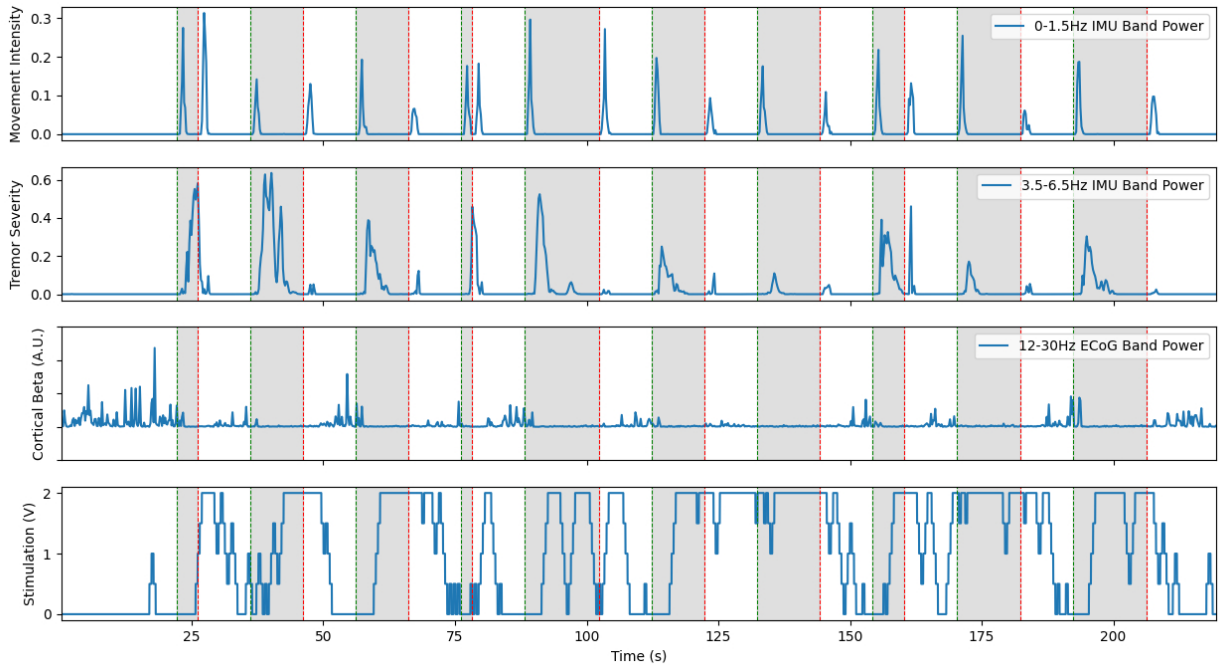

Figure 3.2: An example of a BCI-driven aDBS trial conducted with one of the UW patients. Note that the decoding algorithm reliably detects the onset of stimulation, but usually with a significant delay, especially before stimulation is fully ramped up. The means that the patient often experienced a significant amount of tremor at the onset of movement, which was successfully suppressed a short few seconds after movement onset.

beta band desynchronization phenomenon, would allow stimulation to only be delivered when necessary during movement [124, 156, 162]. Since the Activa PC+S is capable of using cortical LFP power bands to control stimulation, validating beta desynchronization driven aDBS in a distributed fashion, would pave the way for fully implanted aDBS systems.

The chosen to detect movement was the well documented phenomenon of event-related-desynchronization in the beta band (12–30 Hz). This approach consisted of training linear discriminant classifiers to detect significant drops in power in the

beta band that corresponded to movement for each patient. Since the application of high-frequency DBS significantly altered the power spectra visible on cortical recordings, two classifiers were trained in parallel, one with DBS off and one with DBS on. Classifier training used the prompted movement task described above. Average power spectra were computed in each of the four states for that patient on that day. A weight was assigned to each frequency bin, based on how much the power in that bin changed between the rest and prompted movement states for each stimulation state. Once the classifiers were trained, the adaptive DBS algorithm proceeded as follows. Starting in the stim off state, the off classifier listened to the neural data stream, calculated power spectra using Welch's method with a Hann window and normalized by the average and standard deviation of the classifier training data. The dot product of this power spectrum with the classifier weights was fed it into a logistic regression function. When this result crossed a pre-set threshold, indicating the onset of volitional movement, stimulation was ramped up to its maximum clinically permitted value over the course of a few seconds, and the system switched to using the stim-on classifier. Since the ramping of stimulation is known to lead to the greatest number of side effects, the ramp rate was carefully tuned to be the fastest pre-set possible ramp rate that was tolerable for the patient. When the stim-on classifier detected that the beta band power had risen back up to levels indicating rest, stim was ramped back down and the system switched back to the stim-off classifier. The progression of cortical beta is shown in the third row of Figure 5. The thresholds were tuned for stimulation sensitivity, as reliably delivering stimulation during movement was considered more important than reliably turning stimulation off when at rest.

A representative trial of this aDBS paradigm is shown in Figure 3.2. Although the system was always able to identify movement periods in this excerpt, the identification was often delayed and noisy. This was in large part due to the comparatively low spatial and temporal resolution of the cortical strips. Only four electrodes were available, and only a single pair could be used at a time to provide the differential recordings

required for the system. The strip was placed during the implantation surgery and could not be adjusted afterwards. This meant that any imperfections in the initial placement and shifts over time could leave the electrode not in position to optimally observe beta band desynchronization. This effect would have to be independently taken into account when designing a fully deployable aDBS system.

Fully Implanted aDBS The fully implanted system was the culmination of all the work performed and the first potentially clinically translatable system [44]. As identified in previous work, the time delays inherent in the distributed architecture were a major source of aDBS paradigm design difficulties. Since the fully-implanted aDBS algorithm would function entirely within the IPG, these communication delays would be massively reduced. Additionally, there would be no constraints placed on the patient in terms of additional wearable hardware. As a result, IMU data from the android smartwatch and the computational power of a desktop PC could only be used in the training process. The real time updating of stimulation would have to rely entirely on the capabilities of the IPG. By automating the programming methods used to develop the initial aDBS system, patients could benefit from aDBS without requiring the prolonged manual tuning of the classifier necessary in previous versions. This meant that the resultant system should be easily adaptable to new patients to reduce the large time commitment from both patient and clinician required to tune stimulation parameters.

Classifier training proceeded in a semi-automated fashion leveraging the convenience of a desktop PC to determine a classifier that could be used entirely within the capabilities of the implanted device. Due to hardware limitations, there was no explicit way to pass stimulation state as a parameter to the onboard mechanism. To overcome this, classifiers were trained for each of the four possible states the patient and stimulation system could find themselves in: Stim off, rest; stim on, rest; stim off, movement; and stim on, movement. Thirty seconds of data were collected for each

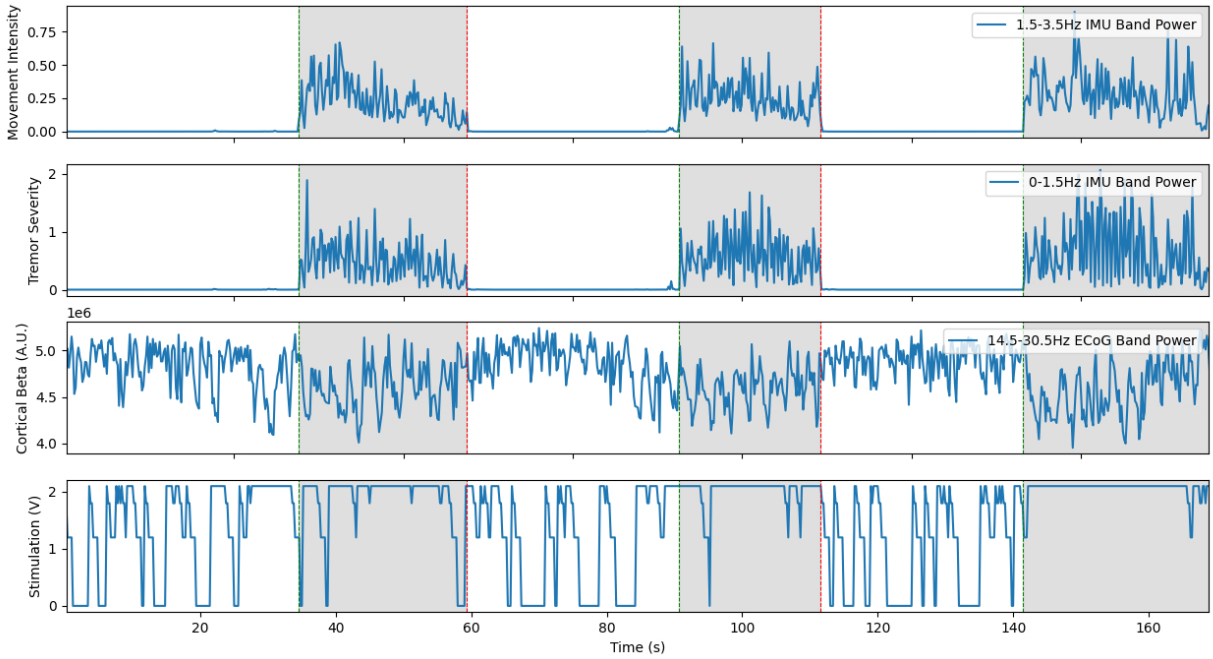

Figure 3.3: An example trial of a fully implanted, cortical driven aDBS paradigm in one of the UW ET patients. Since the classifier is strongly biased towards having stimulation on, the stimulator turns on even during rest. Unlike the previous examples, the cortical beta shown here is the power band estimate computed online by the Activa PC+S during the experiment. Although this is noisier than the offline spectra, the system was able to respond very quickly to the onset of movement and turn on stimulation.

state, with optional repetition of selected tasks. For each of the patients, data from these individual classifiers was combined, which resulted in a classifier that used the power in the band near the stimulation frequency to determine the stimulation state, and the power in the beta band (usually 12.0–28) to distinguish whether the patient was moving or not. Again, since unnecessary stimulation was considered less of an issue that missed stimulation, the classifier was strongly biased towards favoring the stimulated state. This classifier was then uploaded to the IPG and used to switch stimulation on and off. Total time to train this system was under 20 min for each

patient.

An example trial of this system is shown in Fig 3.3. Overall, the dramatic reduction in control delays seem to outweigh limitation on the complexity of the classifier. Additionally, the ability of this system to be deployed to patients without the need for a tether to external hardware cannot be overstated. This is the first reported, full-translatable, aDBS system for ET. There is still much that future studies could do to improve performance with newer hardware. However, this system showed better performance in TEED and greater tremor suppression than cDBS. Although more study is necessary to understand the source of this improved performance, it is an encouraging sign that aDBS systems could improve the lives of patients even more than cDBS systems when deployed in clinical practice.

3.1.2.2 Work from Other Groups

In related work published by our colleagues at the University of Florida, the value of embedded closed-loop for the treatment of ET has been further demonstrated upon in an expanded group of patients over a time period of several months [115]. This study used similar cortical biomarkers, but achieved much higher specificity, potentially due to a more sophisticated paradigm for training their neural classifiers. Our approach for the classifier was based on a large block structure. We collected data during movement and rest, with stimulation both on and off. The classifier was then trained by comparing the power spectra between these four conditions. Conversely, the UF team used data from a prompted movement task to train their classifiers. It seems that these repeated small samples were better able to generalize to the tested behavior, potentially by capturing the transitions between states.

Recent work has shown evidence that there may be alternative biomarkers of movement visible in VIMLFPs [116]. The potential of this VIM approach was further demonstrated by the researchers at Oxford, who developed a VIM based aDBS paradigm using extremely detailed neural recordings [59, 127]. This work was performed using

per-cutaneous leads in the operating room, allowing the team to use the full capabilities of in-clinic hardware, especially the high recording frequency and artifact filtering capabilities. However, the features used for online decoding could be implemented in the fully implanted context with some further advancements in the implantable hardware.

3.1.3 Importance of Longitudinal Changes

As we've already hinted in previous sections, for adaptive algorithms to be successful, a robust biomarker must be found that can control stimulation. Here we evaluate the stability of the well documented phenomenon of beta band (12-30Hz) event related de-synchronization (ERD) in the motor cortex during movement [156]. For these sort of systems to be feasible in real world applications these biomarkers must be robust; there can be no significant short term change in the characteristics of cortical rhythms as a result of DBS that could obscure beta ERD signals. Additionally, the control signal must not be obscured by long-term decay of the signal-to-noise ratio (SNR) of the electrodes or changed by drift in the underlying neural dynamics. The slower the long-term changes, the less often the algorithm will have to be adjusted while remaining effective. Fully implanted systems have many advantages that would facilitate adoption, primarily focused on the complete lack of external hardware. Many experimental systems rely on distributed setups, built around a connection between the implanted device and a computer, which make them impractical in real world situations. Wearable technologies such as smartwatches are more practical, but come with their own problems. They introduce an additional external point of failure, and can significantly increase latencies in the control system. With all system components implanted, patients would need not carry around additional hardware or undergo extra surgeries. Moreover, cortical local field potentials are known to present many characteristics that can be correlated to patient activity and have already been demonstrated as control of aDBS. Fully implanted systems based on ECoG strips, are therefore a

good solution. This leaves concerns about the long term stability of these signals. Previous work has shown potential interactions between stimulation sites in the deep brain and cortical regions that could lead to change over time. There have also been promising results from investigation into the long-term decay of SNR as a result of scarring around the electrodes. To the best of our knowledge, the stability of beta ERD as an aDBS control signal has not been investigated. In this study, we take advantage of an existing dataset to determine beta ERD is robust against both short term interferences from aDBS and long-term drift as a result of chronic stimulation.

3.1.4 Work in this Chapter

In this chapter I cover how we brought our longitudinal analysis methodology to better understand the timescales involved in aDBS for ET. To better understand the requirements on such an aDBS algorithm, I review the nature of ET as a disease. I then cover the work that has been done to develop aDBS for ET. The six-year project to develop a movement-driven aDBS algorithm for ET at UW provided a unique opportunity to perform a longitudinal analysis. This analysis allows me to compare the different aDBS algorithms against each other from the perspective of time. By evaluating the long-terms stability of the control signals used, we can provide an estimate of the timescale on which aDBS of the sort developed at UW could remain effective. Lastly, we review lessons and implications on the future of aDBS for ET.

3.2 Methods

All data collected with our patients during the entire aDBS for ET project was considered candidate data for this analysis. A complete description of all the potential data and its limitations is contained in Chapter 1 Section 2.

3.2.1 Cross-algorithm Comparisons

3.2.1.1 Data Selection

The most important constraints for this analysis were the changing experimental protocol, especially the specific task performed by the patients, and a limited amount of metadata collection, that made tracking changes difficult. Thankfully, the prompted movement experiment paradigm was used during trials of each of the three different paradigms used. The exact type of movement the patient was prompted to perform differed between patients, but all of these were selected to be the style of movement that maximally elicited tremor in the particular patient. As a result, we only use the data from prompted movement aDBS trial experiments, that featured simultaneous smartwatch IMU, time-series neural data, and explicitly logged stimulation amplitude. This left us with 18 experiments distributed across all 4 patients. There were 2 EMG-aDBS trials, 12 BCI-aDBS trials, and 4 fully-implanted aDBS trials.

3.2.1.2 Metrics

To compare between the different algorithms used a number of separate metrics are employed to capture the different aspects of aDBS paradigm performance.

Latency : These metrics are intended to capture the time delay between when a patient begins movement and the time that the system is able to deliver stimulation. We focus on two metrics in particular. Time to Stim Onset is the time measured in seconds between the moment the patient was prompted to begin movement and the time that stimulation first began to increase. Although this low level of stimulation is unlikely to provide measurable benefits in terms of symptom suppression, this metric captures the minimum time delay involved in a single loop of the aDBS system. This includes the time needed to collect data about the patient's movement state, decode this data into a stimulation update decision, and for this decision to be conveyed back

to the stimulator where stimulation amplitude can be changed. We also use a related metric: time to max stim, which additionally includes the time it takes stimulation to ramp up to its clinically target amplitude. This metric measures the amount of time that the patient could expect to have to wait between the time they initiate a movement and their tremor symptoms would be suppressed.

Stim Decision Correctness : Evaluating the correctness of the stimulation decision mandates that we take into account the full 2x2 error matrix. Stimulation on during movement is considered a true positive while lack of stimulation during rest is considered a true negative. Stimulation turning on during rest is a false positive and stimulation turning off during rest is a false negative. These can be evaluated as simple percentages by comparing the stimulation decisions to the movement prompts. In this analysis, however, we instead use the time-averaged amplitude of stimulation delivered during both rest and movement. Both of these normalized by the target amplitude on that trial to give a ration of stimulation in each condition between 0 and 1. This metric has the advantage of taking into account the effects of stimulation ramp up and ramp down, which is more directly relevant to the experience of the patient.

Electricity/Battery Usage : To assess the total amount of electricity injected into the brain and the potential drain on the battery, we use a metric adapted from [77, 44]. This metric, based on the total electrical energy delivered depends on the all the stimulation parameters as well as the amount of time spent at each stimulation amplitude. To make comparisons easier here, we instead present relative values for TEED. *TEED vs cDBS* gives the proportion of total electrical energy saved during the trial by our aDBS paradigm compared to constant stimulation delivered with the same parameters. *TEED vs optimal aDBS* instead gives the proportion of total electrical energy delivered by our aDBS algorithm vs a theoretical optimal aDBS algorithm

that stimulates only when the patient is moving. Together, these two metrics allow us to compare the battery savings of this our aDBS paradigm to both current clinical standard and the target solution.

Symptom Management : To evaluate how well the different aDBS paradigms were able to treat the patient's symptoms, we use two metrics. The first, tremor severity, originally developed by [63], measures the proportion of tremor band power due to tremor:

$$TremorSeverity = \frac{(\text{IMU Tremor Band Power})^2}{\text{IMU Total Power}} \tag{3.1}$$

To take into account metrics other that movement, we also use the Fahn-Telosa-Marin (FTM) tremor rating scale [40]. This is a series of motor tasks carried out by the patient which is then rated by a trained clinician who assigns a numeric symptom severity score. For our patients, their performance on the FTM tasks was videotaped during no DBS stimulation, cDBS, and the particular aDBS paradigm under test at the time. Several blinded clinicians then reviewed these videos and assigned the patient a score in each condition. Results from all clinicians were averaged together to obtain the final estimate of symptom suppression in each condition.

3.2.2 Longitudinal Analysis

Data Selection This study takes advantage of longitudinal data collected at the University of Washington over the course of the development of aDBS algorithms for ET using ECoG strips. As a result, it provides a unique opportunity to evaluate the long term effectiveness of this approach to DBS. The data available for our analysis consists of a cumulative 32 hours collected from 3 patients during a series of full-day visits spread out over many months. Of this data, only the experiments that contained a recorded ground truth for patient movement were amenable to all of the analysis

described below. As a result, special care is taken to avoid errors due to comparisons across inconsistent sample sizes.

Frequency Band Selection Selection of the frequency band used in this analysis is a critical step. Based on the literature, beta ERD usually occurs in the 12-30Hz range. This band is therefore of great interest as a potential control signal. The aDBS algorithms used during the generation of the data were sensitive primarily to changes in the 4- 30Hz band. The short term analysis to detect potential entrainment as a result of aDBS focuses on the 4-30Hz band and the long term analysis to detect slow drift as a result of stimulation itself focuses on the 12-30Hz (beta) band. The converse is also checked. Since the original data was acquired at either 200 Hz or 422Hz, we are well below the Nyquist frequency.

Interval Selection Although our patients were engaged and participated actively in experiments, movement prompt display times are not as close to the true movement times as might be desired. Since patients demonstrated a reaction delay, beta ERD would likely not be observed in a comparison directly around the prompt. To calculate beta ERD, we use Welch's method as implemented by the scipy package in python to calculate the power spectrum for an interval of length ta just before the onset of the movement prompt. Power in the desired band is then calculated by integration with Simpson's rule using the scipy package. A time period equivalent to the estimated max reaction time t_r is then discarded, and a second sample of beta band power is calculated from another interval of length ta. Beta ERD magnitude (ΔPower) is then calculated as the difference in band power between the pre- and post-movement interval. Optimal duration of both t_a and t_r are estimated to maximize the number of beta ERD events observed.

Statistical tools Due to the low number of patients in this and much other experimental human work, we must be careful to avoid introducing error through our statistical analysis. In particular, we focus on the Kruskal-Wallis test, the non-parametric equivalent of one-way ANOVA. Since it is non-parametric, it operates on the median rather than the mean, and can safely be applied to the non-normal distributions observed in our data set. Additionally, this test is robust against differences in the number of samples in the groups being compared. Showing these distributions in a meaningful way is also difficult for the same reason. Therefore, when possible we elect to use the boxen plot: a more detailed variant of the well-known box-and-whisker plot. The central bar represents the median, and each narrowing represents a 10% quantile away from that median. The triangular flags at either end of the boxen plot represent the minimum and maximum values.

3.3 Results

Over the course of this study, we made several important advancements. We demonstrated the feasibility of aDBS for ET using cortical LFPs as a control signal. We then developed preliminary machine-learning driven BCI and volitional control systems. These elements were then put together to create a fully implantable, clinically translatable aDBS system for ET. Over the course of this process, we came to many conclusions that we hope will be helpful to future generations of aDBS for ET. Here we outline some of the most important conclusions that we came to over the course of this study, both practical and theoretical. We also provide our outlook on what the key hardware and ethical challenges that need to be solved in aDBS for ET.

3.3.1 Retrospective Evaluation of UW Work

In previous sections in this chapter, we reviewed the results for each family of aDBS algorithms tested. In this section, we focus on direct comparisons of the different algorithms and their strengths and weaknesses both relative to each other, relative to

the current cDBS standard of clinical care, and relative to the theoretical 'optimal' ET aDBS algorithm.

3.3.1.1 Conceptual comparison

In a theoretically ideal aDBS system for ET, the patient would have elevated stimulation amplitude whenever they are moving, and reduced stimulation amplitude when they are at rest. There is no guarantee that this algorithm would necessarily deliver the best symptom management, however it is the target algorithm as defined for our research.

Patient comfort and freedom of movement are important factors to consider in a deployable system. In this respect, the fully implanted aDBS system has a great advantage over the two distributed systems, as it can function while the patient is free to walk around. The two distributed systems in contrast require that the patient remains tethered to a desktop PC, limiting their practical applications. The EMG driven system also requires the patient to wear an EMG armband, more hardware than any other system. Since patients prefer simpler systems with less external hardware, this was a significant disadvantage of the EMG system.

3.3.1.2 Latency Differences

One of the most prevalent issues that we ran into was the control delay inherent in the distributed system. The differences in these delays for all the algorithms tested are shown in Fig 3.4 (C) and (F). A large part of this delay can be attributed to the hardware limitations of the Activa PC+S. However, any distributed system will necessarily have larger delays due to the increased communication distance over a fully implanted one. Conversely, a distributed system will have much higher computational power and flexibility than a fully implanted one. This can be clearly seen in the difference in stimulation delay between both families of distributed algorithms and the fully implanted system. The median time delay to the onset of stimulation was 1.48s

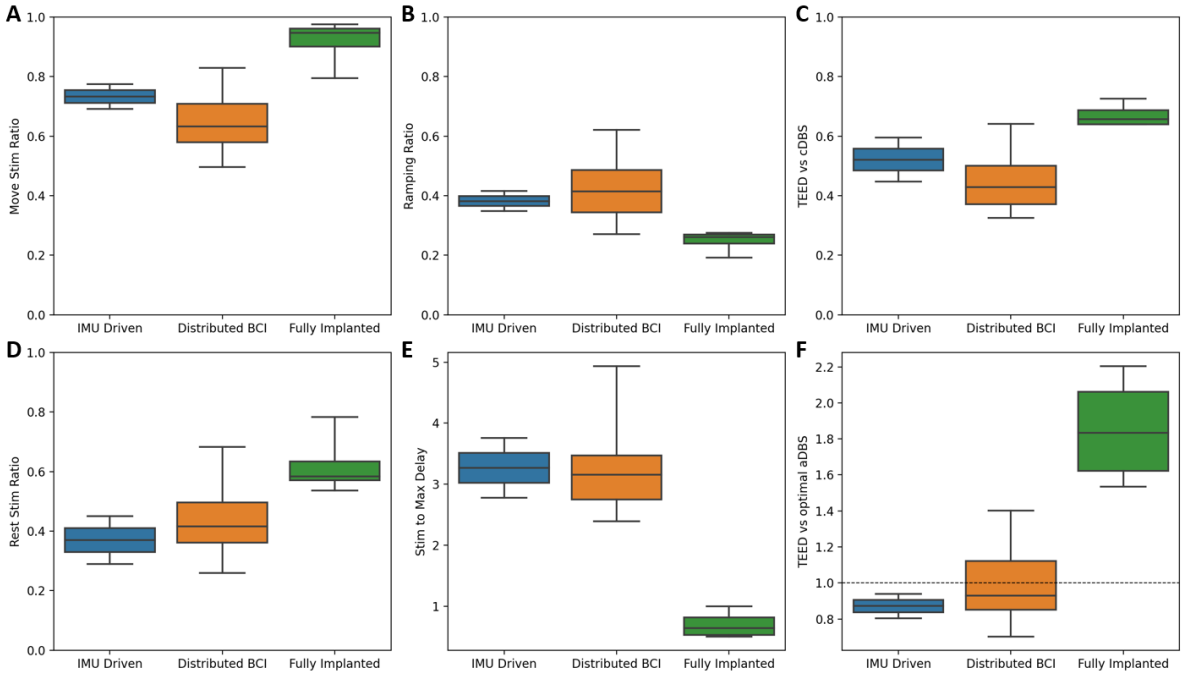

Figure 3.4: Comparison of the different families of aDBS algorithms trialed at UW. The proportion of time during which stimulation is on during prompted movement, which an ideal aDBS algorithm would keep close to 1.0 is shown in (A). The ratio during rest, which should be lower than during movement, is shown in (D). The ratio of time throughout the experiment during which stimulation was wither ramping up or down is shown in (B). The ratio of time at which stimulation was at its 'high' value is shown in (E). The time delay from the movement prompt to the onset of delay or stim reaching in maximum amplitude are shown in (C) and (D) respectively.

and 0.79s for the two distributed systems. This was due to the transmission delays inherent in the system architecture, the extra time required to compute power spectra, and the limitations of the ramp rate. These confounding effects can be clearly seen in Figure 5 as stimulation starts well into the gray prompted movement periods (bottom row), leading to a large burst of tremor before stimulation becomes effective (second

row). This is due to the wide variety of differently tuned and optimized algorithms tested in this family. The fully implanted system by contrast had a median delay of 0.25s, which was a dramatic improvement in latency. As a result, this method did not experience the same bursts of tremor at the onset of movement that hampered the distributed approaches.

3.3.1.3 Stimulation Accuracy and Precision

The stimulation ratios for all the different families of algorithms during both movement and rest are shown in Fig 3.4 (A) and (D).

The ability of these algorithms to generalize to other kinds of movement is seen in the fact that the neural BCI approach resulted in a system that could control the delivery of stimulation with a sensitivity of 90 and 100% for the prompted movement and FTM drawing tasks, respectively [65].

The fully implanted algorithm had an average stimulation amplitude weighted ratio of 91.5% during movement and 62.1% during rest, which corresponds to a sensitivity of 91.8%, and a false positive rate of 28.7% [44].

3.3.1.4 Electricity and Battery Usage

The amount of electricity savings during any given period of time that could potentially be achieved by and aDBS algorithm depends on the amount of movement during the period under investigation. If the patient is moving the whole time and symptoms need to be managed, then even the 'optimal' aDBS algorithm would offer no reduction in total stimulation delivered compared to cDBS. Conversely, if the patient is at rest the whole time, then the possible stimulation savings are 100%. The median percentage of time during all of our prompted movement experiments that the patient was actually moving was 45.3%. Therefore, we will use this movement ratio for all our comparisons in this section.

A comparison of the TEED of the aDBS algorithms in the dataset to the TEED delivered by cDBS is shown in Fig 3.4. The fully implanted algorithms provide the least energy savings of only 34.2%. This is expected due to the intentional bias of the classifier towards stimulation. The distributed algorithms provide much greater energy savings, with median savings of 47.9% and 57.1% for the EMG and BCI driven aDBS algorithms respectively. Importantly, all of these algorithms provide dramatic energy savings compared to cDBS.

A comparison of the TEED of each algorithm to the optimal aDBS algorithm reveals that further work could continue to improve the energy savings. The fully implanted algorithm delivered a median of 183.3% of the total energy an optimal algorithm would have delivered. The median TEED of the two distributed classifiers is actually below that of the optimal aDBS algorithm. Although this does mean that these algorithms used less energy than the optimal classifier, these cam at the cost of significant delays to the onset of stimulation, and occasional instances of the stimulation switching off even as the patient was moving, as can be seen in the example stimulation trace in Fig 3.2.

3.3.1.5 Symptom Management

At the end of the day, the most important metric for improving the patient's quality of life is the algorithms ability to manage the symptoms of ET. In this respect, all of the algorithms tested performed at least as well as cDBS. For BCI-Driven aDBS, over all patients and all sessions, we found a 46.0% average improvement in clinical FTM scores over the no stimulation condition, compared to a 42% average improvement during cDBS. Although the difference between each of these and the stimulation off state was statistically significant, the difference between the two stimulation paradigms was not.

The fully implanted algorithms even seemed to deliver better management of tremor than cDBS.

As has been well documented; side effects of stimulation, especially paresthesia, are often exacerbated while stimulation is being ramped up. The total ratio of time through the aDBS trial that stimulation was either ramping up or down is shown in Fig 3.4 (B). This is the proportion of time that the patient could potentially experience paresthesia if ramping rates are too fast. To ensure the comfort of our patients, we always set a maximum ramp rate for each patient that did not induce intolerable side effects. Several studies have noted the nuances of stimulation ramp rates but to our knowledge no conclusive best practices have been established (Petrucci et al., 2021). In our patients, we noticed that the maximum tolerable ramp rate differed drastically. Further studies will be required to better understand this phenomenon, however, reducing the ratio of time during which the stimulation is ramping by reducing the number of unnecessary state switches is desirable. With a better understanding of the nuances of ramp rates, stimulation could be applied in a way to circumvent paresthesia while still allowing for fast control of DBS.

Due to the communication delays discussed above, we often observed large bursts of transient tremor at the onset of movement. Limitations on the ramp rate necessary for patient comfort mean that the delay between the need for stimulation and when stimulation reached clinically effective levels was even longer. Future studies should therefore investigate aDBS systems with different levels of minimal and maximal stimulation. Instead of switching stimulation between the on and off states, we would instead switch between high and low amplitude stimulation states. The low voltage state would be set low enough to not be noticeable for the patient, and the high level would be at the clinically effective threshold. Since the difference in voltage between states would be smaller, this would enable the system to respond faster without inducing ramping side effects.

3.3.1.6 Comparison Summary

It is clear that there are significant tradeoffs between the three algorithms analyzed here. The EMG-driven aDBS algorithm provides the highest accuracy in stimulation decisions and greatest energy savings. However, it requires the patient to wear additional hardware in the form of an EMG arm-band and requires the patient to be tethered to external hardware to function. The BCI aDBS reduces these requirements on external hardware, but in turn results in decreased accuracy of movement detection which leads to less effective tremor suppression. The fully implanted aDBS algorithm balances the complete lack of external hardware during normal operation which leads to further decreased classifier performance by having the smallest energy savings compared to cDBS.

As we demonstrated in our embedded aDBS experiments, the tradeoffs between these two types of systems can be minimized by utilizing a hybrid approach. By training the system in a distributed fashion, we can utilize the full computational power that larger hardware offers to adapt stimulation in a high dimensional parameter space. By using this system to set the parameters on a simpler, fully-embedded classifier, we can retain the fast response times of a fully embedded system. Future studies will endeavor to further tighten and automate this two-stage control loop. It is likely that the classifier used in the implanted system will not remain effective for long periods of time as the patient's state and medications change. Multi-modal monitoring of the patient's symptoms, with simultaneous streaming of neural, IMU, and even video data, would detect when these changes occur. The system would then either prompt the patient to re-train the classifier or substitute in a previously trained classifier that would better suit the patient's current state.

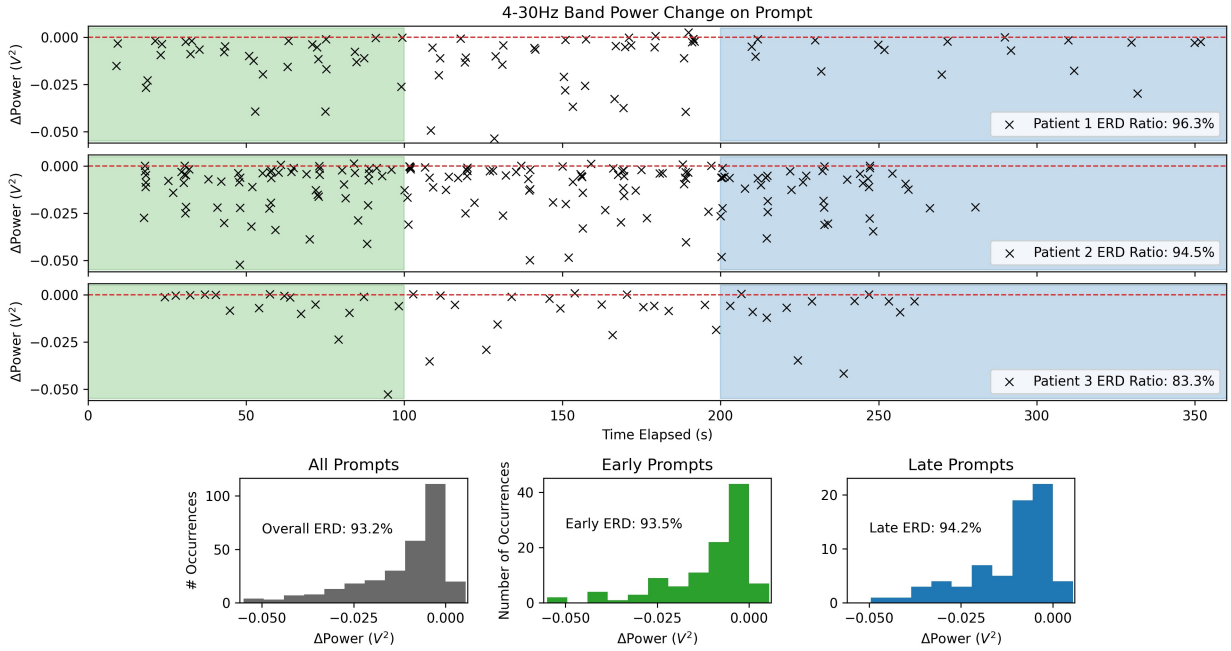

Figure 3.5: Demonstration of the short-terms stability of the beta-band controls signals used in the UW aDBS algorithms during active aDBS. Beta desynchronization during all movement onset events from aDBS trials is aggregated, per-patient, in the top three rows. As can be seen in the bottom histograms, there is no significant change between the distribution of ERD early and late in each trial.

3.3.2 Longitudinal effects of chronic DBS

293 movement prompts from three patients were included in the analysis. A grid search was performed to find the pair of t_a and t_r that gave the maximum proportion of observed ERD events in the 12-30Hz band, to correct for reaction delays. Options between 1 and 10 seconds were considered for ta and between 0 and 10 seconds for t_r . Of these combinations, two pairs gave similarly good results. A ta = 4s with t_r=3s performed the best, with 95.0% of events showing beta ERD. This was closely followed by t_a=2s with t_r=3s at 94.5% of events showing beta ERD. This difference is small enough to likely be insignificant (0.04 times the standard deviation). Since any

resulting control algorithms benefit from shorter response times, we chose to carry out our analysis using t_a=2s with t_r=3s. Results for the short term analysis are shown in Fig 3.5. The top three plots show 4-30Hz band ERD for each of the three patients with relevant data. All recording sessions from all experimental visits were overlaid and aligned at the start time of each individual experiment. The t_a=2s and t_r=3s interval configuration was used to calculate the power change, plotted at the time the movement prompt was shown. The red dashed line is a visual reference for the level where no ERD would have been observed. Visual inspection reveals no obvious trend over time, and this is verified by statistical comparison of prompts that occurred early and prompts that occurred late in individual experiments. The distribution of powers for prompts in the first 100 seconds of each experiment (green region and histogram) and after the 200 second point (blue) show qualitatively similar distributions. This is supported by a Kruskal-Wallis test between these two, which fails to reject the null hypothesis that the early and late samples came from a distribution with the same median ($p = 0.493$). Repeating this process with the 12-30Hz band yields similar results ($p = 0.267$). These results would indicate that the beta ERD are robust to interference from DBS on time scales of a few hundred seconds.

Longitudinal analysis of the same data over the course of months since implantation is shown in Fig. 3.6. Data from all movement prompts in a single session is gathered together into a single boxen plot. All sessions from all patients are shown in the plot. Since sessions were limited by patient availability, intervals between visits are not necessarily uniform. Visual inspection shows that there may be a slight increase in beta band ERD, meaning that the band power change becomes more negative, over the course of 6-10 months. This is supported by a Kruskal-Wallis test comparing the first and last sessions for patients 1 and 2 ($p = 0.001$ and $p < 0.001$ respectively). Patient 3 does not show any significant change ($p = 0.143$), though this may be due to the fact that there are only two sessions, separated by only two months. This result strongly suggests that ERD in the beta band is a viable control signal and robust

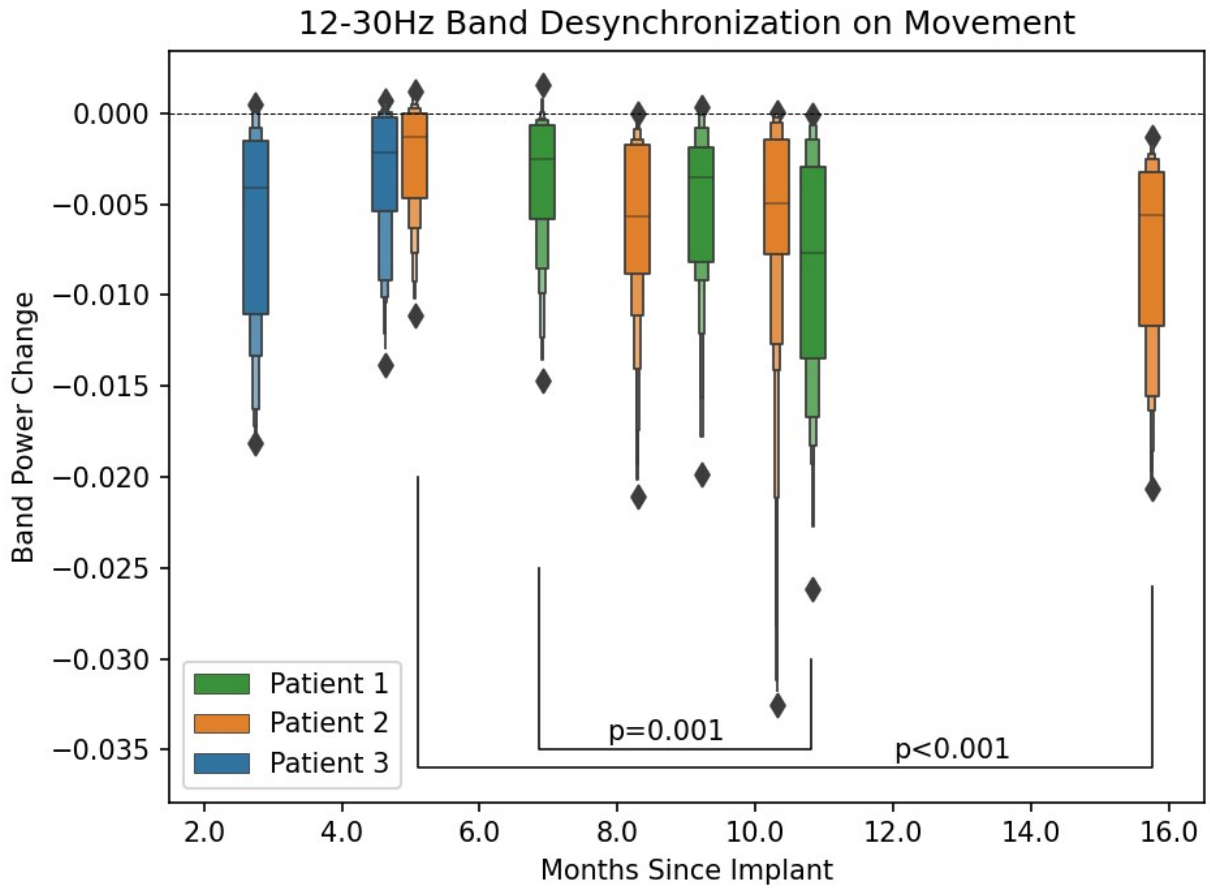

Figure 3.6: Summary of the longitudinal effects observed on the control signals in our aDBS ET patients. Each boxen corresponds to all movement-labeled data collected for one patient during one in-clinic experimental session, central line corresponds to the median. The horizontal dotted line shows the level of no beta band desynchronization, and observations below this line show increased amplitude of beta-band desynchronization during movement onset. Note that as time since implant continues to pass beta band signals remain well-distinguishable for all of our patients. Significant increases in the amplitude of beta band desynchronization were observed between the first and last recording session for both patients 1 and 2.

over long periods of time. Additionally, it shows that cortical surface electrodes are capable of providing informative signals as much as 16 months after implantation.

3.4 Discussion

3.4.0.1 Note on Limitations

The analyses described in the section above have several limitations. The secondary analysis nature of this study and imperfect metadata collection means that not all possible confounding factors could be clearly controlled for. Moreover, only a small number of subjects with a relatively limited number of session for each was available. All these factors may limit how well these results generalize. Future studies are needed to validate these results and investigate the precise nature of the long term increase in beta ERD.

3.4.1 Designing aDBS Systems

Benefit of Nested-Loop approach In the fully implanted aDBS system we describe above, training of the implanted algorithm proceeds in a distributed fashion. In our experiments, we train and test the classifiers within a reasonably short time scale of less than a few hours. However, this begs the question of how often these sorts of classifiers will need to be updated. If training need only be repeated once every few months, then a system like this could be tuned during routine clinical visits. Thanks to the automated nature of the system, this could be carried out by a trained technician instead of a neurologist. However, if the algorithm requires updating on a near daily basis, then any clinically translatable approach would need to be deployable to the patient's home. Both the algorithm update itself and detection of the ineffectiveness of the current algorithm would have to be fully automated. It is also not clear how well the control variables used in our experiments would translate to daily activities. In either case, more work will be required to evaluate the robustness

of aDBS paradigms developed in this way.

A detailed understanding of this trend would allow aDBS algorithms to slowly adapt to any changes and remain in an optimally functioning state. Since this would decrease the frequency at which the aDBS stimulator would need to be manually updated, it would reduce the number of expensive clinical visits. Taken together, these results emphasize ECoG strips as a part of future aDBS systems. The availability of on-label cortical strips is therefore a crucial step towards widespread adoption of such systems. Once this occurs, patients with severe ET could begin to benefit from the advantages of fully-implanted aDBS systems.

The results of this paper are in clear support of ECoG strips as a practical element of future developments of aDBS algorithms. We have shown that beta ERD is robust over short time scales of several hundred seconds despite the potential interference of aDBS triggered off this same band. We have also shown that beta ERD sensed from cortical strips do not decay over the time course of multiple months. These results are in agreement with similar research previously conducted with PD patients, and long term evaluation of ECoG as an option for long-term BCI systems. Cortical strips are therefore shown to be a reliable source of bio-markers for a fully implanted aDBS systems. Implanted cortical strips also seem to be a safe extension to the regular DBS implantation process. Despite the additional surgical complexity involved with the implantation of cortical strips, which carries inherent risks, the clinical team involved with these patients has not observed any adverse events directly attributable to the implantation or presence of cortical sensing strips implanted for the PC+S project.

Another aspect of aDBS that has been shown to be important is the exact methodology of training the neural data classifier. This is highlighted by the comparison between our results and those recently published from the University of Florida [115]. It is clear that the training data and paradigm must be designed with the ability to translate into a more naturalistic context in mind. Their study achieved higher sensitivity, specificity, and overall accuracy. When classifiers such as these are designed,

we believe it is important to take the patient experience and clinical practice into account. False negatives, lack of stimulation when it is needed, are more detrimental to the patient than false positives, stimulation even when it is not needed. For this reason, it is important that classifiers are optimized primarily for sensitivity instead of just overall accuracy. These considerations will help ensure that aDBS is best optimized for the needs of the patient.

Hardware Matters New implantable DBS systems promise to solve some of these problems by being rechargeable and sampling at much higher rates. This will allow long-term chronic recordings and analysis of a wider range of frequency bands.

One of the big limitations of the data collected and analyzed here were due to the intrinsic hardware limitations of the Activa PC+S system. This system was a first-of-its kind system that enabled groundbreaking research. However, from the perspective of time the limitations become apparent. The sampling rate of 422Hz is quite low compared to more common neural recordings. Moreover, the system is pretty noisy meaning that band power data beyond about 75Hz is below the noise floor. The on-board power band estimates therefore are very sensitive to artifacts and are very noisy. Embedded algorithm parameters can only be uploaded and data can only be streamed through separate USB connections. This means that the patient must be physically tethered to a computer during data collection, and limits the practical deploy ability of the aDBS paradigms.

Next Generation aDBS systems such as the Medtronic Summit RC+S improved significantly on these limitations [141]. This research platform had a higher sampling rate, better artifact reduction, and far improved onboard processing and data streaming capabilities. It is only thanks to the advanced capabilities of this platform that we were able to perform the work in the next chapter. The upcoming Medtronic Percept device, unfortunately removes many of the capabilities of the RC+S, providing only much slower streaming of pre-selected power band data. While these capabilities may

be sufficient for some simple clinical applications of aDBS, they provide only a limited view into the inner workings of the disease and patients state. It is our hope that new aDBS devices capable of performing groundbreaking research will become available.

3.4.1.1 Alternative Approaches

A potential direction of research that we did not fully explore in our study is the potential to smoothly adapt the levels of stimulation to the severity of symptoms. Systems like this could follow a similar approach to the beta thermostat approach demonstrated in the PD literature [163], adjusting stimulation to even out therapy when provided in conjunction with medication. We performed proof of- concept distributed studies using bollinger bands to drive 'graded-DBS' [22]. This work showed promise in handling the non-stationary dynamics and adapting the stimulation algorithms in a patient-specific way. However, hardware limitations prevented us from implementing this approach in a fully implanted context. This approach would allow the aDBS system to adapt stimulation levels more precisely to the changes in patient. As a further extension to this cortically driven aDBS, we also endeavored to build a volitional aDBS system. This approach offered to make aDBS more flexible and applicable to more diseases by handing control of the stimulation directly back to the patient. The idea was that although controlling stimulation would take conscious effort initially, repeated training and daily use would allow the patient to develop automatic, almost subconscious control of the stimulation. As indicated by our conversations with the patients, this also had the potential to greatly improve the patient experience by increasing the sensation of agency and strengthening the identification of the device as a part of themselves [19, 63].

3.4.2 Outstanding Challenges in aDBS for ET

Unclear Nature of ET Future studies of aDBS with larger numbers of patients will also be capable of investigating the variations observed between patients. Recent

work is increasingly suggesting that ET is not a single disorder, but rather a family of related disorders that need to be treated slightly differently [139]. This is also supported by the wide variation in effective stimulation settings observed even in our relatively small cohort of patients. For most ET patients, the recommended stimulation frequency is close to 140 Hz, but for one of our patients we found the most effective stimulation occurred near frequencies of about 90 Hz. A survey of optimal stimulation parameters determined in an automated way, matched with neural recordings on and off stimulation would be a promising avenue to investigate these differences. In this context, aDBS is firmly in the regime of personalized medicine. Future aDBS applications should retain the focus on tuning stimulation individually to the needs of each specific patient. Broad generalization is useful only in so far as it simplifies the training process of each patient individualized aDBS paradigm and highlights the nuances of each patient's needs.

One of the great remaining hurdles in the development of aDBS for ET and treatments for ET in general is the limited understanding of the pathology of the disease. Although progress has been made recently suggesting the involvement of multiple network components and the central role of the cerebellum many unanswered questions remain [126, 45, 67, 119, 118]. With this growing evidence for the role of the cerebellum, it will be important to identify whether the cerebellum is the sole generator of pathological oscillations, or do further changes need to happen for resonant frequencies to arise. Moreover, it is unclear whether pathological changes occur in other brain region that facilitate the propagation of tremor oscillations. An improved understanding of the pathology of ET could lead to new stimulation targets and stimulation paradigms that better counteract the symptoms of ET, or even pharmacological treatments that directly target the underlying pathological changes. Due to the acute control possible in aDBS systems, this research will be able to help answer many of these important questions.

No Explanatory Model aDBS for ET lacks the sort of robust neural biomarkers that are directly correlated to symptoms. In PD, for example, abnormally long bursts of STN beta activity has been shown to be directly correlated to symptoms, while a sharp peak in the gamma band has been shown to be directly correlated to stimulation-induced dyskinesia, resulting in a perfect control signal for ramping DBS amplitude up and down [169, 87, 105]. The ET DBS systems described here must rely on proxy biomarkers such as beta band desynchronization as a measure of movement, which is correlated with symptoms. Since this biomarker imperfectly follows symptoms, there is a hard cap on how close to optimal performance our systems can come. Thankfully future aDBS research is well positioned to begin unraveling these questions. We look forward to seeing how this work develops the field, as VIM aDBS would remove the need to implant the additional cortical strip required for our approach [115]. If a direct biomarker of tremor severity, identifiable both during and off stimulation, could be found, the aDBS could be driven exactly as needed. Further verification and development of both these and cortical biomarkers of tremor is essential for robust aDBS systems for ET. The presence of multiple simultaneously computable biomarkers would allow for cross-validation and increased robustness.

For these reasons, one of the developments that could most dramatically push aDBS for ET forward is a well-verified, explanatory model of ET. Specifically, such a model should explain at a high level the interactions between brain areas that give rise to pathological ET tremor oscillations. Though one such promising model has recently been proposed, more work is needed to verify this model and determine how it can be fit to patient data [171, 35]. This modeling effort should proceed in conjunction with the imaging-based modeling efforts [27, 4, 103]. As these models move towards more predictive power, they will would allow for more insight during the process of selecting the implant site and tuning aDBS parameters for new patients [104]. Neural and biophysical recordings of patient state could be used to cluster the patient with other patients that display similar symptoms, with the expectation that similar stimulation

would be similarly effective for patients within a cluster. This could dramatically reduce the number of parameters sets that need to be tested to find an effective stimulation paradigm. When repeated with multiple patients this would result in a map of ET disease states and related diseases. Even once DBS parameters are set, it is likely that over the course of the patient's daily life, optimal stimulation parameters will change. As the patient takes medicine, for example, the stimulation amplitude and frequency change. A reliable model could then be used to help inform automatic switching of stimulation parameters as the patient's state changes. Fundamentally, a good enough model could provide insight into the pathology of the disease and aid the search for ET treatments that target the underlying cause rather than just treat symptoms.

Challenge of Clinical Translatability As has been noted in several reviews in this field, there is a large gap between the experimental demonstration of aDBS and a clinically translatable treatment [7, 101]. We took part in collaborations to address some of these challenges with our development of automated tools for optimization of DBS and aDBS paradigms [55]. However, several challenges remain.

One of the largest challenges to the clinical translatability of the neural-driven aDBS systems we developed over the course of this study is the availability of on-label cortical strips. These are a required component for the implementation of aDBS systems with that use movement intention sensed from cortex as a control parameter. As implied by our BCI control work, these could also be used in a volitional fashion to seamlessly offer the patient a greater degree of control over stimulation. Recent work along with our own analysis have shown that cortical electrodes such as these retain a high signal to noise ratio for years after implantation [113, 46]. This has been further confirmed by the group at UF, who have shown that aDBS driven by cortical strips is robust over several months [115]. Moreover, in the context of our study we observed no adverse effects as a result of the implantation of the cortical strips, which

may motivate work pursuing their clinical validation as a safe extension to existing DBS systems.

Ethical Considerations Finally, there are a number of ethical considerations that arise as aDBS systems are translated into common clinical practice. We have discussed a number of these elsewhere, but we provide a short review here [19]. Concerns have been raised about the potential for stimulation to cause shifts in the user's perception of selfhood and agency e.g., in cases where stimulation causes behavioral changes as a side effect [75]. aDBS might one day mitigate and manage these side effects. Questions remain, however, about how users will interact with more robust aDBS systems; how those interactions impact clinical outcomes and quality of life [18]. It is not clear, for example, how much control users will want over stimulation parameters, or how involved they want to be in aDBS algorithm training, or how different algorithms will impact user experience. To investigate these questions, a neuroethicist on our team (TB) lead a series of longitudinal, semi-structured phenomenological interviews with each patient the goal of which were to give patients the opportunity to describe using the experimental aDBS platform [19]. A more complete discussion on these considerations is available in [20].

Chapter 4

APPLICATIONS TO PARKINSON'S

4.1 Introduction

Parkinson's disease is one of the most commonly known neurological diseases. It is estimated that there are almost a million patients over 45 with Parkinson's disease in the US alone [95]. The disease usually first becomes visible in a patient due to three common factors: stiffness and difficulty initiating movement, known as bradykinesia; tremor during rest, and uncontrolled jerky movements, known as dyskinesia's. One of the possible treatments for these motor symptoms, once medication is reaching the limit of its effectiveness, is DBS.

4.1.1 What is Parkinson's disease

To understand the pathology of Parkinson's disease, it is essential to first understand the structure of the motor circuitry and the basal ganglia. The basal ganglia (BG) consists of a series of interconnected brain regions located in the deep brain, below the cortex. It has been studied extensively, as it is a central component of several essential brain circuits, especially movement and sensation [102]. Due to the neurological motor disease focus of this work, I will focus on the motor circuitry in the BG.

There are several main parts of the basal ganglia , illustrated in Fig. 4.1, based on [79, 54, 102, 5]. The striatum is the largest single structure, and is composed of several nuclei which wrap around the majority of the rest of the basal ganglia. The striatum serves as the main incoming connection to the basal ganglia from the cortex and other regions [79, 102]. The globus pallidus is technically a part of the striatum though it is generally considered as a separate stricture. It is comprised of

two regions, the external capsule of the globus pallidus (GPe) and the internal capsule (GPi) [68]. Together, these two regions are critical in regulating the voluntary control of movement, the GPi in particular acts as the central output nucleus for the BG [68]. Further medial and inferior of the GP is the sub thalamic nucleus (STN). It is physically the smallest of the BG nuclei. While it's exact role and organization is still a matter of debate, it is known to be essential for the regulation of the other BG nuclei [8, 6]. The substantia nigra (SN) so called for its distinctively darker coloration. It is itself comprised of two sections, the substantia nigra pars compacta (SNc), which contains primarily dopaminergic neurons that supply dopamine to the striatum and other BG regions, and the the substatia nigra pars reticulate (SNr) serves mainly to relay signals back through the BG [138]. At the very center of the deep brain lies the thalamus, which is also composed of many sub-nuclei [158]. As a whole the thalamus is involved in many roles but is also the primary relay center between the brain and the rest of the nervous system [79, 158]. These nuclei are all tightly interconnected and work together.

The connections between the BG nuclei that compose the primary motor circuitry have been thoroughly studied. There are a few theories for how the BG nuclei work together to control motor activity, the most common of these is the direct, indirect and hyper direct pathway model, illustrated in Fig 4.1, and originally proposed by [5, 102]. Both the direct and indirect pathways begin with projections from the motor cortex to two groups of medium spiny neurons in the striatum. These two groups are primarily differentiated by their afferent connections and the type of dopamine receptors they primarily express: either the D1 or D2 receptors [138, 47, 54]. Those in the direct pathway, which receive excitatory dopamine connections from the SNc through their D1 receptors, have inhibitory projections to the GPi. The GPi has a direct inhibitory connection to the thalamus, which is then responsible for relaying the motor commands to the muscles, and sending somatosensory responses back to the cortex. Neurons in the indirect pathway in the striatum, identified by their D2

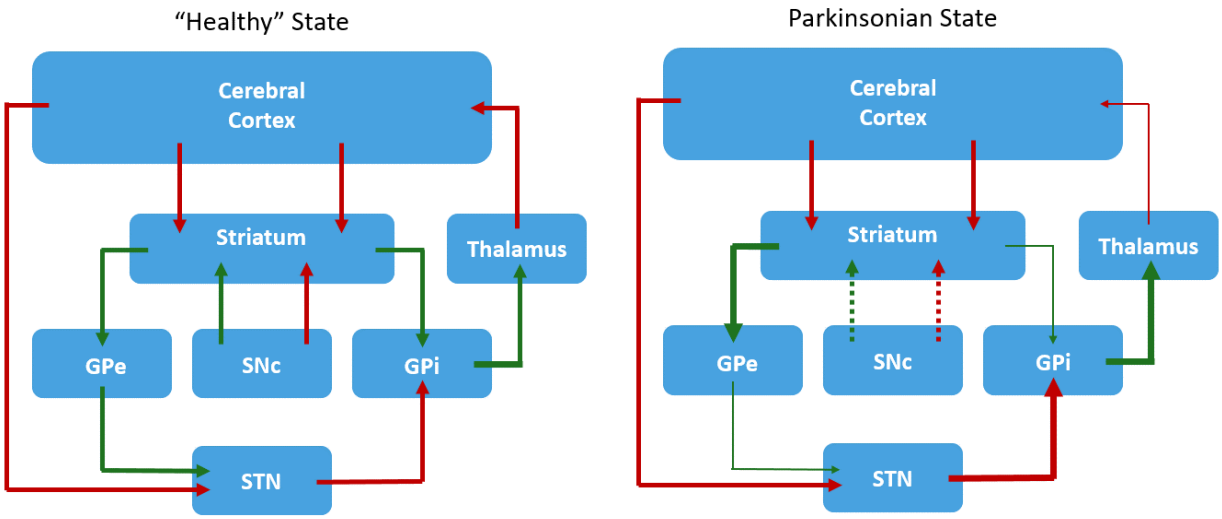

Figure 4.1: Traditional model of the basal ganglia connections showing the effect of Parkinson's on the main motor connections in the basal ganglia. These are connections between the dominant nuclei, which are the Striatum, the internal and external capsules of the globus pallidus (GPi and GPe), the subthalamic nucleus (STN), the substiantia nigra pars compacta (SNc) and the thalamus. Excitatory connections are shown in red, inhibitory connections in green. In the canonical model, these are thought of as the direct pathway (cortex -> striatum -> GPi -> Thalamus), the indirect pathway (cortex -> striatum -> GPe -> STN -> GPi) and the hyper direct pathway (cortex -> STN -> GPi -> thalamus). In Parkinson's, the dopaminergic connections to between the SNc and the striatum are reduced, leading to different effects on the direct and indirect pathways. Influences that are strengthened or weakened by this change are shown as changes in the thickness of the corresponding arrow.

receptors, have inhibitory projections to the GPe, which then also has an inhibitory projection to the STN. The indirect pathway ends with an excitatory projection from the STN to the GPe. The last main pathway in the BG is the more recently discovered hyper direct pathway [79, 54, 117]. This pathway consists of inhibitory projections from the cortex directly to the STN, which then links in to the inhibitory connection

between the STN and GPe, part of the indirect pathway [117]. Together, these opposed pathways allow the circuitry to optimally select from a collection of possible alternative actions [13, 54].

This model of basal ganglia function is effective at explaining the observed changes in Parkinson's disease. One of the hallmarks of Parkinson's disease is the dying off of dopamine producing neurons in the SNc [5, 97]. In some patients, dopaminergic neuron quantities in the SNc may have reduced by up to 60% by time the patient is diagnosed [32]. In the classical model of the BG, this results in a dramatic decrease in the modulatory dopaminergic connections from the SNc to the striatum. These connections are primarily excitatory to the D2 neurons and inhibitory to the D2 neurons. This results in a suppression of the direct pathway connection to the GPi and a strengthening of the indirect pathway connection to the GPi. Both of these changes strengthen the inhibitory connection out to the thalamus, explaining the classic parkinsonian symptoms of bradykinesia, or a difficulty initiating movement. Conversely, when treated with an excessive amount of levodopa, the opposite effect takes place, and the patient will experience dyskinesias [38, 97].

As successful and widely accepted as this model of BG function is, it does of course have its limitations. The first and foremost of these is the fact that it takes into account only a small portion of the known connections in the BG [79, 102]. A more complete illustration of all the known connections is shown in [79]. Proponents of the classical BG model contest that these other connections are responsible for the other functions of the BG and do not significantly alter the simplified picture of the model as presented above [102]. Recent research has also shown that the clear distinction between the D1 and D2 MSN in the striatum is not nearly as clear as is implied by the classical model [12]. This would imply that the direct and indirect pathways are more closely intertwined and not as easily distinguished. Moreover, this model of the basal ganglia focuses on the motor effects of parkinsons disease. Other effects, such as the decline in memory associated with neurodegereation of the hippocampus, need

to be modeled and taken into account separately [26, 66]. It is my hope that future developments in whole-brain conectome projects can lead to models that can take the motor and non-motor symptoms of PD into account simultaneously.

4.1.2 DBS – Current State of the Art

Later in the progression of PD patients may be prescribed and elect to undergo a DBS implantation procedure. This is usually applicable primarily to patients once the medication has reached towards the limits of its effectiveness. In this surgery, a set of electrodes is implanted into the deep brain structures. The exact location varies, but is usually either the sub-thalamic nucleus (STN), the external globus pallidus (GPe) or the internal globus pallidus (GPi). The stimulator system consists of an electrode lead implanted in the target area, an implantable pulse generator (IPG) implanted subdurally over the chest, and a wire under the skin that connects the two. The IPG the delivers continuous rectangular stimulation pulses, based on a series of stimulation parameters. For most devices for patients with PD, these consist of stimulation amplitude (usually between 1 and 5 µA), stimulation frequency (usually close to 130Hz, and pulse width (60-90us), along with appropriate selection of stimulation contacts. A trained neurologist adjusts these stimulation parameters to best treat the patient's symptoms. Due to the huge size of the parameter space, this process often takes several weeks to months. This process is made all the more complicated by the nonlinear interactions between the parameters, which make simple heuristic approaches to adjusting parameters difficult. The patient also usually has a patient controller, which allows them to select from several pre-set stimulation settings as well as subtly adjust the stimulation amplitude within safe limits set by the clinician. This method of stimulation is known as conventional or continuous DBS (cDBS) and is the current standard of clinical care.

The exact mechanism of action of DBS is still a matter of debate. All three of the common target regions (STN, GPe/i) are part of the cortico-thalamic-basal ganglia

motor pathways, as discussed earlier. [98, 107]. It is generally accepted that the periodic pulses of DBS work to modulate the oscillatory activity in the basal ganglia motor circuitry in some way. This modulation [97, 126]. There is evidence that this modulation is a result of both inhibitory and excititory effects [107, 159, 98, 142]. It is therefore likely that future research will reveal that the exact mechanism of action can be location specific and a result of several interacting effects.

4.1.3 aDBS – Next Generation therapy

As effective as this therapy is, is still has several limitations. Most relevant to the work here is the fact that symptoms, and therefore the optimal level of stimulation, varies over time. For example, most patients take some form of levodopa at the same time as DBS. This means that if the stimulation amplitude is set so that symptoms are well managed before the patient takes their medication, then once the medication washes in the stimulation may be significantly to high and the patient may experience dyskinesia [38]. Similarly, the patient's symptom severity depends on movement [123, 86]. In all of these cases, automatically adjusting stimulation amplitude up or down to maintain symptom suppression without causing side effects is desired [87, 163]. This is exactly the purpose of aDBS, which aims to adjust stimulation automatically and in real time. To be able to evaluate the effectiveness of any aDBS therapy, it is valuable to have a way to measure the severity of the symptoms the patient is experiencing. The most common metric used in clinical practice to summarize the overall level of symptoms is the UDPRS scale [39, 52]. However, this scale requires that a trained clinician to evaluate the patients symptoms in a series of tasks and give a rating between I and V in each symptom category, for a final combined score. As has been noted in recent literature this means that the scale can be subjective, and requires the input of a trained clinician [114, 120]. As a results, there has been a recent effort to develop more quantitative measures of PD motor symptoms that can be used as a direct symptom measurement in automated systems [143, 74, 121, 112, 136]. Most of

these focus on the use of accelerometers and video to track and quantify the patients motions. Using these automated metrics, we can begin to optimize aDBS paradigms to minimize symptoms.

4.1.3.1 Compared to ET aDBS

In general, aDBS for PD is a more complicated problem than aDBS for ET. In essential tremor, patients experience essentially just one significant symptom: tremor [57]. This tremor can occur in various body parts, but occurs most commonly during movement [57]. Additionally, tremor is well treated by simply increasing stimulation amplitude until the symptoms are reduced [29, 55]. As a result, the control problem in ET is very simply defined: when tremor is observed increase the stimulation amplitude. Alternatively, since tremor occurs during movement, we can simply increase stimulation amplitude during movement. In PD, by contrast, there are a much wider variety of symptoms that are non-trivially related to each other. These symptoms include bradykinesia, dyskinesia, tremor, gait changes, along with a number of non-motor neurological effects [123, 38, 24]. Moreover, these symptoms vary quite significantly between patients [123, 100]. Stimulation amplitude must also walk a tighter line in PD, since under stimulation leads to stiffness and tremor, while overstimulation can cause dyskinesia. Overall, this means that aDBS for PD is a much harder problem to specify the requirements.

4.1.3.2 Existing aDBS for PD

Recent research in PD aDBS has been generally focused on creating aDBS paradigms that act on a single dimension of the potential symptom space. Each of these has generally focused on a single biomarker associated with a particular subset of aDBS symptoms and has tuned aDBS stimulation to keep that biomarker within the 'healthy' regime.

Beta bursts are short term bursts of increased beta band activity, primarily visible in the STN that have been shown to be a biomarker of motor impairment [85, 84, 153, 154, 155]. Specifically, it has been shown that longer, more frequent and higher amplitude beta bursts in the STN all correlate with greater motor impairment as demonstrated by worse UPDRS scores[157, 84, 155]. Moreover, it has been shown that these pathological beta burst are suppressed with stimulation in a way that correlates with increased motor scores on stimulation and medication [153, 154, 34]. This has led several groups to develop aDBS paradigms based on rapid responses to beta bursts in the STN [87]. In this paradigm, the LFPs in the STN are monitored for power in the 13-30Hz band, and this band power was smoothed with boxcar average filter. Whenever this value crossed a pre-determined threshold, stimulation was rapidly ramped up to a pre-defined threshold, until the band vale returned below the threshold, at which time the stimulation was ramped back down. All studies using this family of approaches observed that this stimulation paradigm reduced both the number and duration of pathologically long beta bursts compared to the no stimulation condition [153, 85, 129]. The prevalence of long bursts was improved compared to cDBS [153]. Most importantly, this aDBS paradigm showed an improvement in UPDRS scores that exceeded No DBS, randomly varying DBS and even open loop DBS [157] Although very effective by all metrics, this approach does require very high sampling rates of LFP recordings (1kHZ), as well as very fast response times (on the order of 10s of ms) [130]. Given currently available hardware [141], this limits this approach to aDBS to the clinical setting, as that is the only environment where the bench-top recording hardware is practical. Hopefully future developments in implantable technologies can make this approach feasible in the at-home setting.

Another approach to controlling aDBS is based on a more smoothed out measurement of beta-band oscillations in the STN [163, 53]. In this approach, an implanted device, such as the Activa PC+S or the Summit RC+S, is responsible for performing measurements of the STN LFPs. Due to the resolution, these measurements are in many ways

a boxcar average of the rapid variations observed in the beta bursting literature. Two thresholds are set to control changes in amplitude in the aDBS paradigm. Whenever the measured beta amplitude increases past the upper threshold, then the stimulation amplitude is increased. Conversely, it is decreased whenever it drops below the lower amplitude. When the beta amplitude is between the two thresholds, then the amplitude is left constant. The effect of this algorithm is slow smooth variations in the amplitude of stimulation keeping the measured beta amplitude in the target band, much in the way a thermostat keeps temperature in the desired range. This study demonstrated that this dual threshold approach to PD aDBS is capable of effectively managing motor symptoms while reducing the average amplitude of stimulation.

Other aDBS studies have specifically focused on tubing DBS to reduce the cortical gamma, a biomarker associated with dyskinesia [148, 105]. This biomarker, referred to as either the 'finely tuned cortical gamma peak' or 'dyskinesia peak' refers to the sharp peak observed at half the stimulation frequency, that becomes visible in the cortex of patients who are experiencing dyskinesia [148]. In this approach, which can be implemented entirely on the implanted device, stimulation is switched to the low state whenever the dyskinesia peak is detected, and then switched to the high state whenever the dyskinesia peak disappears [105]. This approach allowed for a reduction in total stimulation power delivered with no adverse effects to the patients' clinical scores [105].

Recent studies have also worked to allow aDBS to take into account the changes that occur during sleep [50, 56]. These were driven by increasing evidence that not only is sleep disturbed in Parkinson's, but sleep disruption can accelerate the progression of the disease, and sleep affects the biomarkers often used to drive aDBS [106, 172, 14]. In this aDBS paradigm, the authors ad a second independent LD classifier to determine whether the patient is sleeping or not, and modify the behavior of the base aDBS paradigm, tuned to manage motor symptoms of PD, to account for sleep-related changes [50]. The trained sleep detector was able to detect sleep with an

average accuracy of 79%, and successfully managed unwanted changes in the aDBS algorithm due to sleep. This study serves as an important reminder of the significance of understanding the day-scale changes that affect aDBS.

4.1.4 Work in this chapter

In this chapter we aim to leverage all the data available to us thanks to the framework developed in the previous chapters to better understand the challenges and constraints that aDBS will have to overcome as it moves from a clinically demonstrated treatment to a chronic-at home therapy. As suggested above, these challenges take several different forms. There are multiple interacting timescales on changes in neural signals that need to be accounted for. Given the available data, we are able to investigate both the single second, day, and month-scale effects in detail in our one patient. Based on this, we describe a pipeline which can be used to predict how often outer-loop changes will need to be implemented. Additionally, we demonstrate the effectiveness of this analysis pipeline in our patient.

4.2 Methods

4.2.1 Data Sourcing and Availability

For the analysis in this chapter, we leveraged the remote data collection infrastructure described in detail in Chp 1. This system allowed us to collect multiple parallel data streams from the patient's home, during both pre-planned clinical tasks, and during the patient's regular daily routine. The infrastructure collected all these data streams independently to ensure robustness. Each data stream was then time aligned and prepared into a quick-to-analyze HDF5 format, and indexed using a SQL database for easy data lookup.

This chapter's analysis required time series neural data, accelerometer data from both wristwatches, as well as video and pose. This yielded a total of 77.5 hours of useable

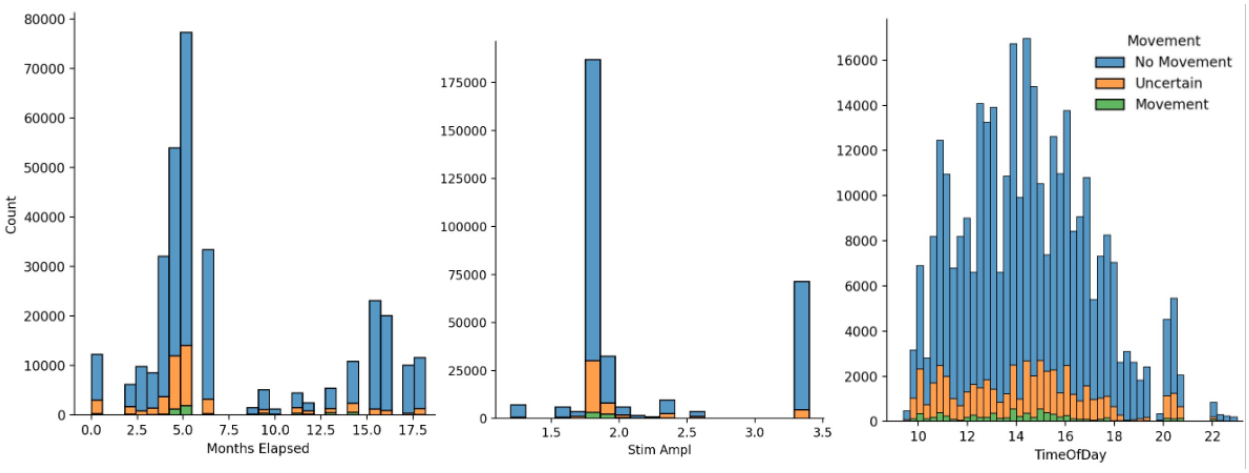

Figure 4.2: Summary histograms showing the distributions of the available data across the three most important continuous variables along with movement state. Although the amount of data collected on individual days varied, there is reasonably consistent coverage over the over 18 months during which recordings were carried out. Recordings were primarily conducted around the 1.9 mA and 3.4 mA levels, which were the patients preferred stimulation levels in each of the two stimulation electrode configurations used during the full-day recordings. Recordings cover the duration of the patients working and waking hours quite well, but no recordings were performed during sleep. It is important to note that although every condition included movement epochs, there are overall far more rest periods than movement periods, due to the patients tendency to remain extraordinarily still during most of the time they sat at their desk. These nuances of the data distributions mandated careful attention to the statistics used.

data distributed over 63 days, with a median duration of 30 minutes of usable data a day. The raw data is summarized in Fig 4.2. The data used was collected as part of three largely independent experiments conducted over the course of several days. These were the ‘Multi Stim’ experiment, the ‘move aDBS’ experiment, and ‘Full Day’ recordings.

The 'Multi Stim' experiment tested the effect of varied stimulation amplitude on the patients motor function and neural signals. For each day of this experiment, the patient was stimulated at three different levels in one of the possible permutation orders, and all permutation orders were tested. The stimulation levels were scaled relative to the patients preferred stimulation level at the time. Recordings were performed and preferred stimulation was determined in the stimulation group that allowed for the cleanest recordings. During each stimulation condition, the patient performed a series of clinical tasks (chest tapping, finger tapping, hand open-close, wrist rotations, sit-stand, walking) as well as typing a pre-selected paragraph. Two six day sets of this experiment were performed, first with 50%, 75%, 100%, and second with 50%, 100%, 110% of the patients preferred stimulation amplitude. Since all stimulation was delivered in an open loop format, all the available data could be used in our analysis.

The 'move aDBS' experiment tested a novel data-driven method of determining an optimal adaptive stimulation algorithm entirely driven by neural data [30]. Similarly, to the 'Multi Stim' experiments, three different stimulation paradigms were applied and the patient performed a pre-defined series of tasks during each stimulation block. Here however, the stimulation was delivered either as open loop DBS at just below the patients preferred stimulation, or during either optimized, or inverted aDBS. During optimized aDBS, the patient was stimulated at a higher level when movement intent was detected, and stimulation was reduced when the patient returned to rest. During the inverted condition, the patient was stimulated at a higher level during rest and at a lower level during movement. The evaluation tasks performed by the patient were a series of clinical tasks (chest tapping, finger tapping, nose tapping, and wrist rotations) followed by typing. For our analysis we only include the data recorded during the open-loop portion of these recordings.

The 'Full Day' recordings were designed to investigate the effect of DBS on the patient's daily life. As a result, the patient was instructed to stream data from all the available modalities while following their normal daily routine of working from home.

Recordings were considered to be opportunistic, and the patient streamed data once a week whenever they were not engaged in other experiments. Additionally, the patient was encouraged to stop the recording whenever they liked. This mostly involved the patient sitting in front of their computer typing and using the mouse, as well as occasionally walking around and eating. Stimulation during this time was delivered in an open-loop fashion with the patient occasionally adjusting the stimulation amplitude using the patient controller. At first, stimulation was delivered in a configuration optimized for the best STN recordings. However, after the patient reported that prolonged periods in this configuration were unpleasant, we switch stimulation to the patients preferred 'Dual Monopolar' configuration. Data from all of these recordings was considered a candidate for our analysis.

4.2.2 Nested Loop approach to aDBS

In an effort to ensure that the aDBS algorithms remain effective out of the clinic and for prolonged periods of time, we implement a nested-loop approach to aDBS algorithm tuning. Originally proposed in [**?**] and illustrated in Fig 4.3, this method uses two interacting adaptation loop each acting at different time scales to manage the various time scales affecting the neural signals in PD. In this section we review the generalized approach of this method as well as our particular implementation.

"Fast" Inner loop The fast inner loop is responsible for rapid stimulation changes on the scale of fractions of a second, and is entirely embedded on the implanted device. This algorithm uses the implanted hardware to measure power in pre-defined power bands, from up to 4 simultaneous channels from any of the 4 cortical and 4 stimulation lead electrodes. These power bands are fed into two threshold linear discriminant classifiers, each with two pre-defined thresholds and pre-defined feature weights. Each classifier can use its own power bands and operate independently of the other, resulting in a 9-state state table. Each state in this table can be configured to be responsible for

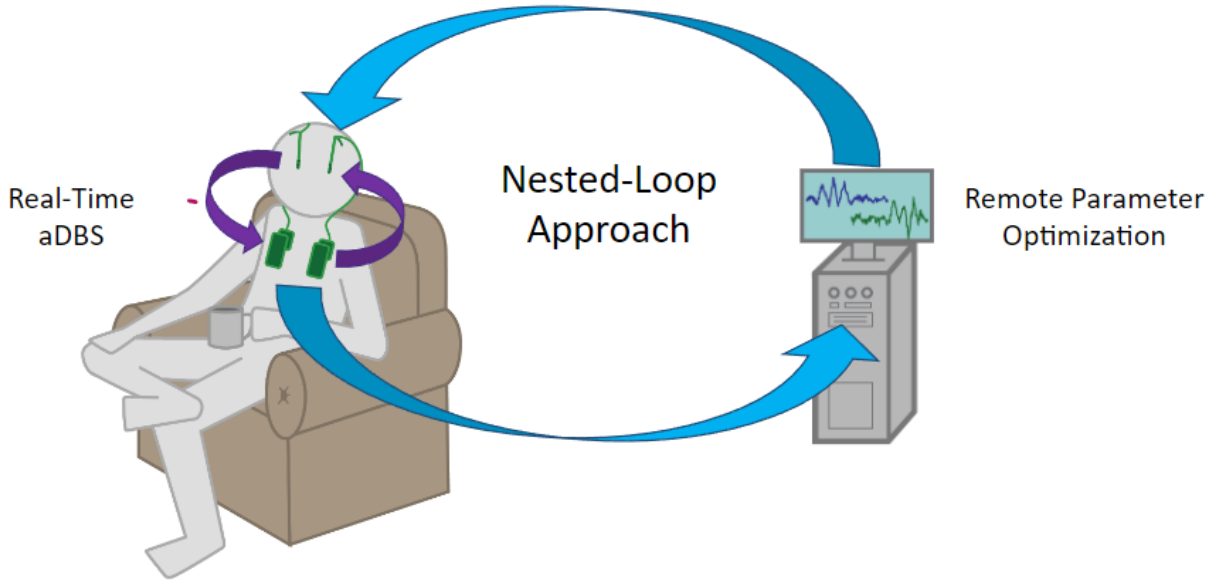

Figure 4.3: The nested loop approach to deployable chronic aDBS. A real-time aDBS algorithm is developed that can run in a fully-implanted configuration. This allows the patient to move freely. The performance of this algorithm is then asynchronously evaluated and re-tuned using a distributed system to ensure that the performance remains high on long timescales.

assigning either a new target stimulation parameter, or an adjustment to the current value of a stimulation parameter. This system can therefore be flexibly programmed to rapidly adjust the simulation delivered to optimally suppress symptoms at any given moment.

In our implementation, we used a simplified proof-of-concept approach for developing our aDBS algorithm. This approach was targeted to treat the most significant symptom our patient reported: stiffness/bradykinesia. Therefore, we aimed to develop a stimulation paradigm that would increase the amplitude of stimulation when the on-board device detected an intent to move, and reduce stimulation when the patient was at rest. More detail on the on-board algorithm training and evaluation is available in

[30].

"Slow" outer loop Despite the efficacy of the inner loop algorithm on short time scales, there are several issues which limit it's long-term effectiveness and mandate the usage of the outer loop. For a few months after implantation, there are sizeable drifts in the impedance on the recording electrodes, which result in a significant change in any power band metrics recorded. Furthermore, there are well documented neurological changes in many of the common aDBS control signals as a result of PD, which continue to increase as the disease progresses [32, 97].
In our implementation this outer loop is responsible for tuning the LD parameters to ensure that the embedded movement state decoding remains sufficiently accurate over the timescales of multiple months. This is accomplished with Gaussian Process Bayesian Optimization, which is repeated with additional data [30]. This allows the classifier thresholds to be adjusted to ensure that the classifier remains maximally effective at decoding movement states.

4.2.3 Development of movement labels

Since our recordings are performed primarily over the motor cortex it is essential to take movement into account. We expect that the well documented power band changes on movement will dominate the signal. Therefore, we need to be able to separate the movement periods from the rest periods by making labels of movement state for all of our data. Manual labels already exist for the clinical tasks, but still need to be made for all of the free behavior data. To make our analysis of the entire dataset feasible, we need an automated way to determine when these movements are occurring. To train this classifier, we take advantage of existing movement labels for clinical tasks from previous experiments. These movement labels were created by manual inspection of the video data, and precisely label the kind of movement the patient was performing as well as the time period that this movement lasted. The

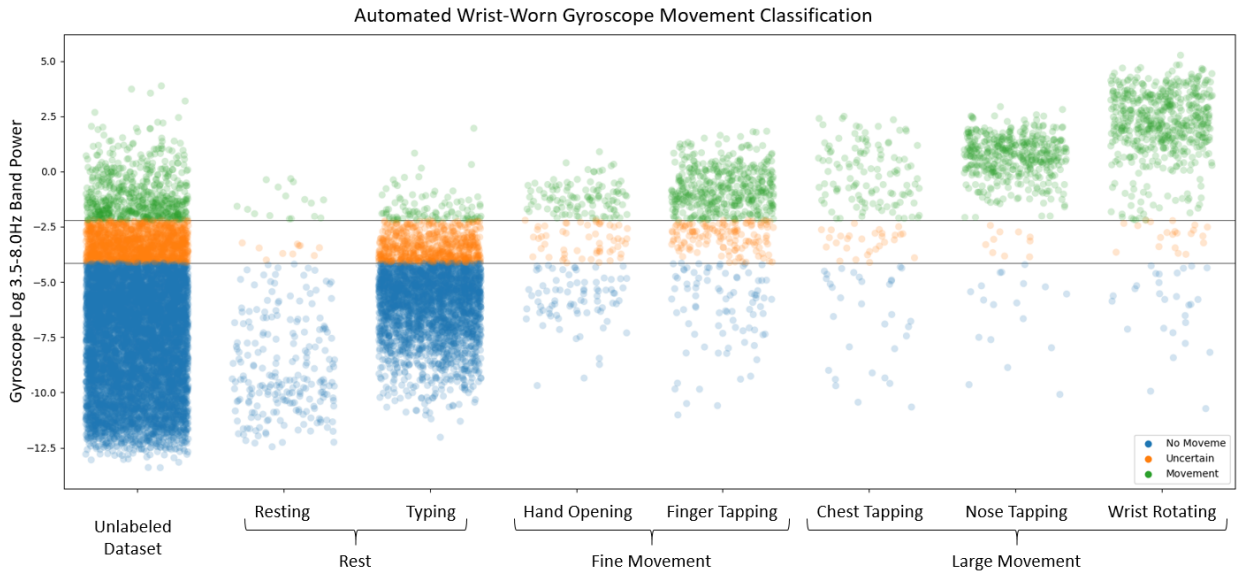

Figure 4.4: Demonstration of the trained movement classifier on the test data. The different categories of manually labeled movements are shown along the x axis. Each point corresponds to a single section of accelerometer data used to calculate a single value of the 3.5Hz - 8.0Hz band used for classification. The two decision boundaries are shown with the two thin horizontal lines, and the classification of each observation shown by the color of the point.

potential movement options in the labels were, 'Rest', and 'Typing' which were binned together into 'Rest'; 'Finger Tapping' and 'Hand Open Close' binned together into 'Fine Movement'; and 'Nose Tapping', 'Chest Tapping', and 'Wrist Rotation' binned together into 'Large Movement'.

Gyroscope data for 1.28 second (64 samples at 50Hz) intervals was used as the base data for calculation of a classifying feature. 378 power bands with evenly space limits between 0 and 15Hz, at least 2Hz wide were used as possible features for movement labeling. Power in each of these bands was calculated using a combination of the scipy welch and simps functions [167]. A separate 1-dimensional LD classifier was trained on the log of band power for each of these possible features, to optimally distinguish

between the 'Rest' and 'Large Movement labels [78]. These classifiers were compared by comparing their classification accuracy. Scores of all classifiers on all modalities of gyroscope data are shown in Fig 4.4. The best classifier, with an accuracy of 96.01% used gyroscope power in the 3.5-8.0Hz band for classification. We then used this classifier as a basis to generate high confidence labels for all the unlabeled data. Using the built in sklearn functions, we calculate confidence for each prediction in the training set. Points with a confidence of less than 80% were considered to be 'Uncertain'. This allows us to convert this into a simple two threshold classifier. We consider all observations with a higher band power than the highest uncertain band power to be "Movement" and all observations with a lower band power than the lowest uncertain observation to be "No Movement".

4.2.4 Neural Data Processing

aDBS has been suggested to have significant effects on the reduction in symptoms beyond the total amount of stimulation delivered [43, 129]. To avoid this confound we only include data recorded during open-loop stimulation in our analysis. Distributions of all the available neural data are visualized in Fig 4.2.

The power bands we chose to include in this analysis are chosen based on several factors. We include several canonical power bands, such as the alpha and beta bands. Many of the bands included are power bands that were found to be useful in LD decoding based on a data-driven discovery approach described in [30]. A few additional test bands such as broad-band power and stim frequency were included primarily for test and reference purposes. A full list of all the included power bands is available in Table 5. To compute the power in each of these bands, we first computed the full power spectrum using welches method with the scipy function welch [167, 137]. We used 512 sample intervals at 500Hz. Each window was centered on a corresponding movement label. All the individual power bands were extracted from the PSD using the simps function [167, 21].

As is apparent in Fig 4.2, there are several changes sources of concern that we need to be very careful about in performing our analysis. First of all, there are very different numbers of observations for the various parameter combinations. To manage this, we perform our analysis on a single summary metric on all observations with a common combination of values on all independent variables. Since the distribution of band powers is not normal, we take the log of band power. This brings the band power distribution closer, but still not quite normal. Therefore, we use the median to more accurately capture the peak of the distribution compared do the mean. More complicated is the fact that in stimulation settings and experimental conditions were changed over the time during the course of the recording. As can be seen in Fig 4.2 these changes correlate with both stimulation and date.

To handle this, we fit an ordinary linear model to the medians of the data in all dimensions. This model takes into account all the known variables in the data that could potentially affect the band power, specifically: date, time of day, stimulation amplitude, whether the patient is moving or not, the stimulation contacts, and kind of recording being collected on that day. We also include interaction terms between these factors. To perform this fitting, we use the ordinary linear system functionality in the statsmodel package, which fits equations of the form:

$$Y = \sum_{i}^{n} \beta_i x_i + \sum i \neq j \gamma_i j x_i x_j \tag{4.1}$$

Where x_i is the i-th variable, β_i is the coefficient measuring the effect of x_i, and $\gamma_i j$ is the coefficient of the interaction term between variables x_i and x_j. Once we fit the model to the data, we perform a combined ANOVA and ANCOVA analysis to find which of the parameters have significant effects. Then we regress out all the confounding effects to find just the effect of the variable of interest.

Once we have the effect coefficients on each axis, we would like to able to compare them to each other to get an intuitive sense of the relative effect sizes. This is a non-trivial since the units on the x-axis are all different, and therefore all the slopes

are completely different. Therefore, we use Pearson's R coefficient. This metric varies between -1 and 1, where more extreme values correspond to trends that more closely follow a line with the corresponding slope direction. Moreover, we measure Pearson's R on the data for a particular dimension after all other significant effects in the model have been regressed out. This ensures that we get a more precise idea of the size of each individual effect, since many of these effects counteract each other in the full dataset.

4.2.5 LD Prediction Reconstruction

Looking at the individual trends for individual power bands is insufficient for predicting how the onboard LD outputs and therefore the LD predictions will change over time. This is due to the fact that the onboard LDs are based on a linear combination of power bands. Therefore, we need to be able to reconstruct the predictions for the entire LD.

4.2.5.1 LD Function

We assume that the onboard classifier is a simple linear discriminant classifier in one dimension with two thresholds. Therefore, the classifier can be described as:

$$f = \begin{cases} \text{high} & y > t\_H \\ \text{low} & y < t\_L \\ \text{inBand} & \text{otherwise} \end{cases} \tag{4.2}$$

Where y is the variable that is used to make the classification decision, and t_H and t_L are the high and low decision thresholds respectively. The value of y is calculated as:

$$y = \vec{w} \cdot \vec{x}' \tag{4.3}$$

Where $\vec{w}$ is a 4-length vector of the LD weights, and $\vec{x}'$ is a 4 length vector of normalized power band values. This normalization for each component is calculated:

$$x_i' = (x_i - s_i) * r_i \tag{4.4}$$

Where s_i and r_i are the corresponding components of a 4D normalization vector that is pre-set, and x_i is the power band value in internal device units. As a reference, see the Summit RDK docs, pages 85-93, especially page 90 Trends in the LD In this analysis, we consider the effect of a change in band powers as a result of time. This means that x_i in the above equations is not drawn from a constant distribution, but rather takes the form:

$$\ln(x_i) = a_{ti}t + a_{mi}m + \epsilon + b \tag{4.5}$$

$$x_i = \exp\left[a_{ti}t + a_{mi}m + \epsilon + b\right] \tag{4.6}$$

$$x_i = e^{a_{ti}t} \cdot e^{a_{mi}m} \cdot e^{\epsilon+b} \tag{4.7}$$

where t is the number of months elapsed, m is a binary variable determining whether the patient is moving the corresponding arm, ϵ is normally distributed noise $\epsilon \in N(0, \sigma)$, and b is an arbitrary offset variable. This means that the time dependence of the LD would be:

$$y = \sum_{i=1}^{n=4} w_i r_i(a_{ti})t + w_i r_i(a_{mi} + \epsilon) \tag{4.8}$$

where we've substituted $\epsilon = \exp(\epsilon + b - s_i$ or alternatively $\epsilon \in N(b - s_i, \sigma)$ for the sake of simple notation However, if there is an interaction between movement and time, then power will take the form:

$$\ln(x_i) = a_{ti}t + a_{mi}m + a_{xi}mt + \epsilon + b \tag{4.9}$$

$$x_i = \exp\left[a_{ti}t + a_{mi}m + a_{xi}mt + \epsilon + b\right] \tag{4.10}$$

$$x_i = \exp\left[a_{ti}t + a_{mi}m + a_{xi}mt\right] \cdot \epsilon \tag{4.11}$$

It now becomes convienent to think about this equation in two cases, where $m = 1$

and the patient is moving, or $m = 0$ and the patient is at rest:

$$x_i = \begin{cases} \exp\left[(a_{ti} + a_x i)t + a_{mi} + \epsilon + b\right] & m = 1 \\ \exp\left[(a_{ti})t + \epsilon + b\right] & m = 0 \end{cases} \tag{4.12}$$

$$x_i = \begin{cases} A_{mi} E \exp\left[(a_{ti} + a_{xi})t\right] & m = 1 \\ \epsilon \exp\left[(a_{ti})t\right] & m = 0 \end{cases} \tag{4.13}$$

Where we have substituted $A_{mi} = e^{a_{mi}}$ and $E = e^{\epsilon + b}$ This allows us to write the resulting classifier results for the two cases:

$$y = \begin{cases} \sum_{i=1}^{n=4} w_i \left(r_i A_{mi} E \exp\left[(a_{ti} + a_{xi})t\right]\right) - s_i) & m = 1 \\ \sum_{i=1}^{n=4} w_i \left(r_i E \exp\left[(a_{ti})t\right] - s_i\right) & m = 0 \end{cases} \tag{4.14}$$

Note that the difference between the two cases will be:

$$\Delta y = \sum_{i=1}^{n=4} w_i r_i \epsilon \left(A_{mi} E\left[a_{xi} t\right] - 1\right) \exp\left[a_{ti} t\right] \tag{4.15}$$

We can rewrite the trends for the averages of the band power in each of these two cases as

$$y_m = y_o + \Delta y \text{ and } y_o = \sum_{i=1}^{n=4} w_i \left(r_i \epsilon \exp\left[(a_{ti})t\right] - s_i\right) \tag{4.16}$$

In this way of re-writing, it is clear that the maximum accuracy of the decoder will depend on only a few factors: the value of Δy, and the spread or error in measurement in each of the two movement conditions. The larger the spread, the greater the overlap in observations between the two conditions and the more difficult decoding between the two conditions will be. However, if we assume that this variance is relatively constant for the timescales of internet, then the ability to decode between movement and rest will depend only on the value of Δy: if this value is constant in time, then a simple linear discriminant decoder will continue to work, as long as thresholds are adjusted appropriately.

4.3 Results

In this section we cover the trends that we find on the individual power bands. We then combine these trends using a model of the particular LD classifier used in our experiments to provide predictions of the trends observed in the LD outputs over the timescales of months.

4.3.1 Band Specific Trends

We observe many significant effects in different bands. A full summary of all significant effects we observed in both hemispheres in our patient is shown in Appendix 5. It is important to note that there are a number of bands with significant effects which are always regressed out as un-interesting confounds. These effects were: 'isMonopolar', the experimental condition, and the interaction term between 'isMonopolar' and stimulation amplitude. The 'isMonopolar' column shows the effect of stimulation being delivered in an either Monopolar (stimulation applied between electrode 2 as the negative contact and IPG case as the positive contact) or Dual Monopolar (stimulation applied between electrodes 2 and 3 as the negative contacts and IPG case as the positive contact) configuration. As expected, these effects are the most pronounced in the STN channels, since during dual monopolar stimulation and recordings were performed on the same electrode, leading to an increase in electrical artifacts. We also consistently remove the effects of experimental condition, which is observed in a few bands, particularly in the cortex. This effect is difficult to pinpoint, but we hypothesize it is largely due to the different motor tasks the patient performed during these experiments. During both the Multi-Stim and aDBS experiments, the patient performed a series of slightly different clinical tasks, while during the full day recordings the patient spent most of their time working at their computer.

The observed effects of the difference between movement and rest are shown in Fig4.5. We observe the well documented beta band (12-30 Hz) desynchronization in all cortical

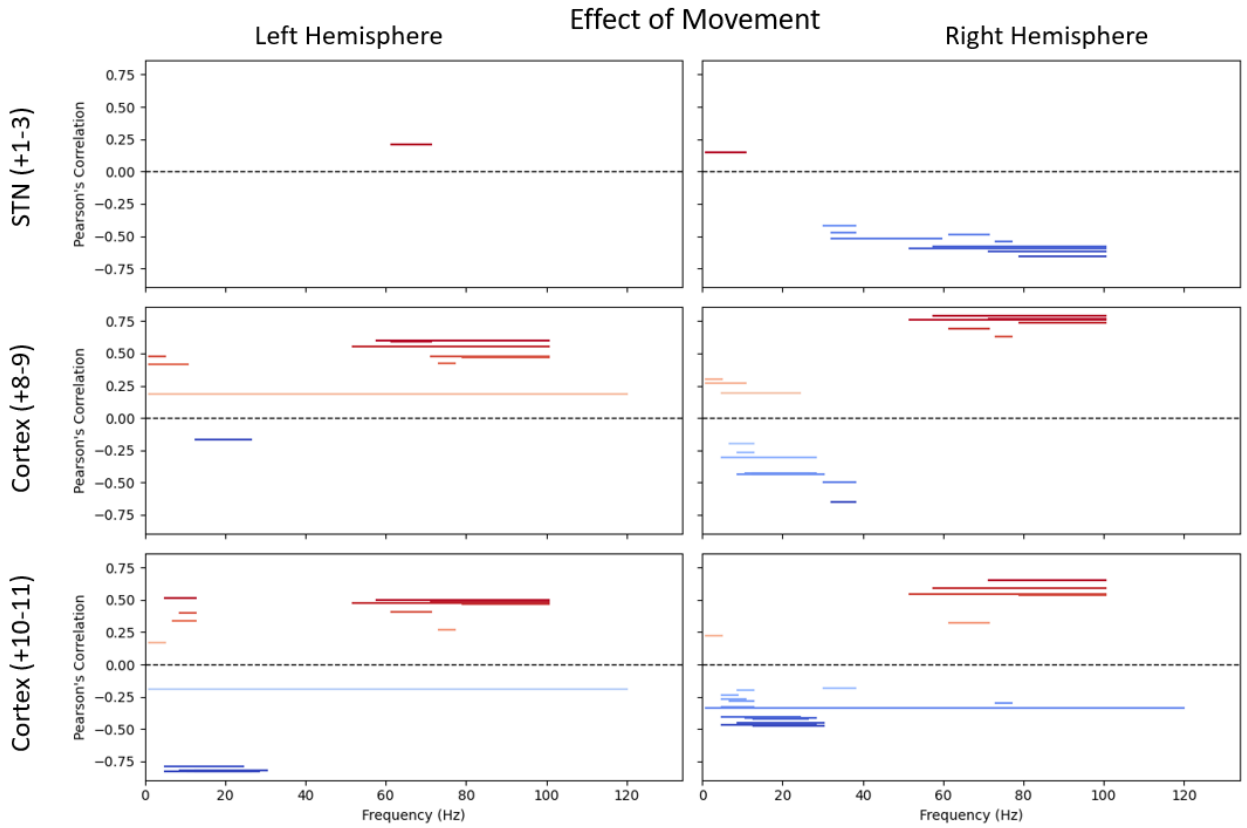

Figure 4.5: Significant effects of movement condition on the band power in all bands tested. Each horizontal line corresponds to a single trend. The horizontal extent shows the frequency extent of the power band, while the vertical location shows the effect size as a ratio between the band power variation due to the observed effect with the total variation observed. Only significant ($p<0.05$) effects are shown. Trends for each hemisphere are shown in the left and right columns respectively, and each row shows trends observed in a different recording location. We observe the expected beta band suppression in the cortex during movement along with an increase in cortical gama and alpha power.

recording locations, though it was most prominent in the +10-11 electrode pair. We also reliably observe increases in the alpha (4-12 Hz) and broad gamma (50-100 Hz) bands as is expected from the literature. Overall, a majority of the observed effects

were more pronounced in the right hemisphere, which corresponds to a higher average movement decoder accuracy for that side. IN addition to the expected effects, we also observe a few broadband effects, but these can be explained by the average of all the individual band effects adding up.

There are a variety of significant effects observed for stim amplitude, as illustrated in Fig 4.6. As expected, we observe a positive interaction coefficient between band power and stimulation amplitude in the STN for all bands. This makes sense as we are stimulating in this location and electrical effects are virtually guaranteed to bleed into the signal despite artifact mitigation strategies. This effect is particularly notable in the stim band (127-133 Hz) where it explains close to 100% of the observed variation. There is also a positive effect between the stimulation amplitude and the stim band power in the cortex of 0.756, 0.677, $p < 0.001$ in the left and right cortex respectively. There is also a strong broadband suppression of band power on the +8-9 recording contacts in the right hemisphere, which is especially pronounced around the beta range. A similar effect is observed on the left hemisphere for the +10-11 recording contacts. This observation aligns with the expectation that stimulation facilitates movement, which is partially facilitated by a reduction in the beta band power.

A summary of all the significant observed effects of Months Elapsed in all recordings is shown in Fig 4.7. For the most part, these are dominated by a gentle decrease in the power of most cortical bands in both the left and right hemisphere, with a few exceptions. On average, this corresponded to a 29.2% decrease in the broadband log power for the affected electrodes. On the timescale of months, we observe a decrease in band power across most bands in the +10 -11 recording contacts.

Our analysis also included the possibility of interaction terms. Specifically, we looked for the effect of an interaction between movement and months elapsed, which we only observed for three different high gamma bands in the cortex. These effects were relatively small compared to the relative effect sizes of both effects independently, only explaining an average of 11.0% of the observed variation in these bands. Moreover, the

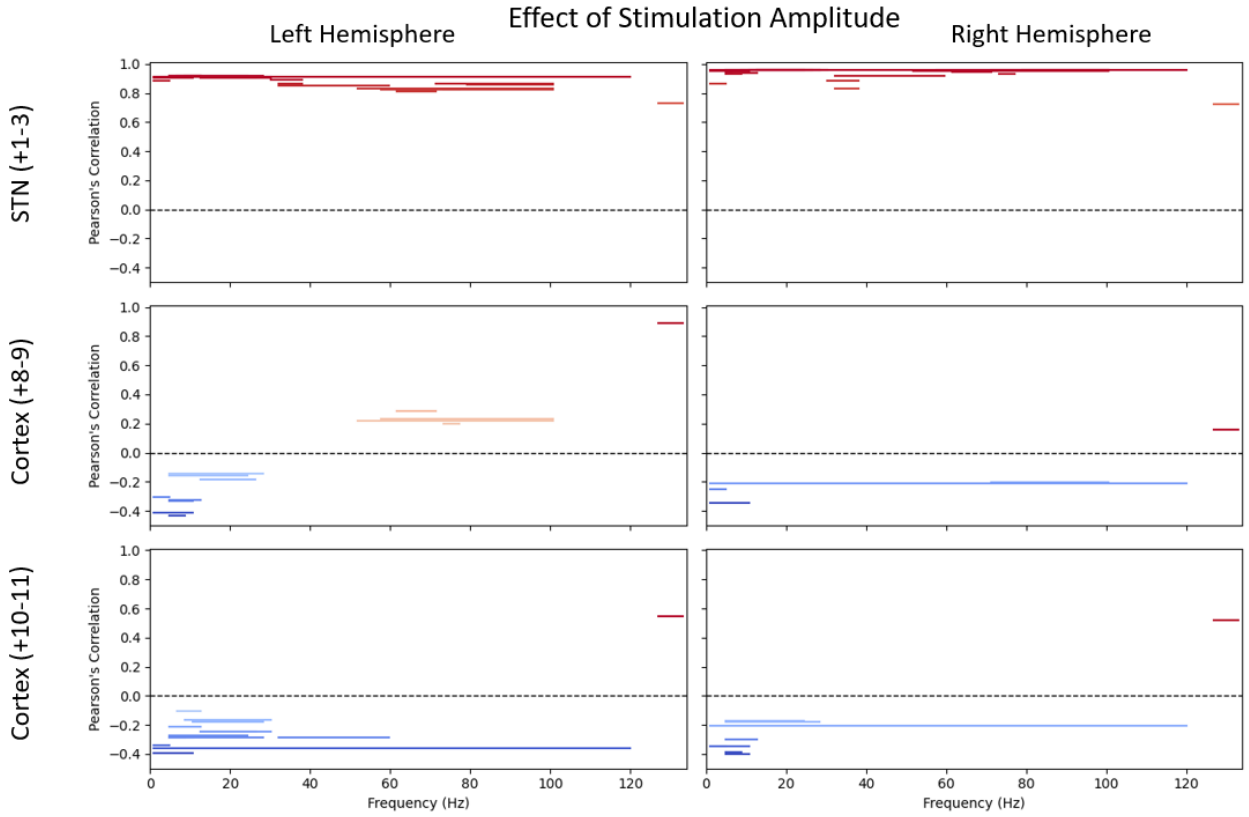

Figure 4.6: Significant effects of stimulation amplitude on the band power in all bands tested. Each horizontal line corresponds to a single trend. The horizontal extent shows the frequency extent of the power band, while the vertical location shows the effect size as a ratio between the band power variation due to the observed effect with the total variation observed. Only significant ($p<0.05$) effects are shown. Trends for each hemisphere are shown in the left and right columns respectively, and each row shows trends observed in a different recording location. As expected, stimulation has a dominating effect in the STN since that was the stimulation location, and a much smaller effect in the cortex.

coefficients were positive in the left hemisphere and negative in the right hemisphere.

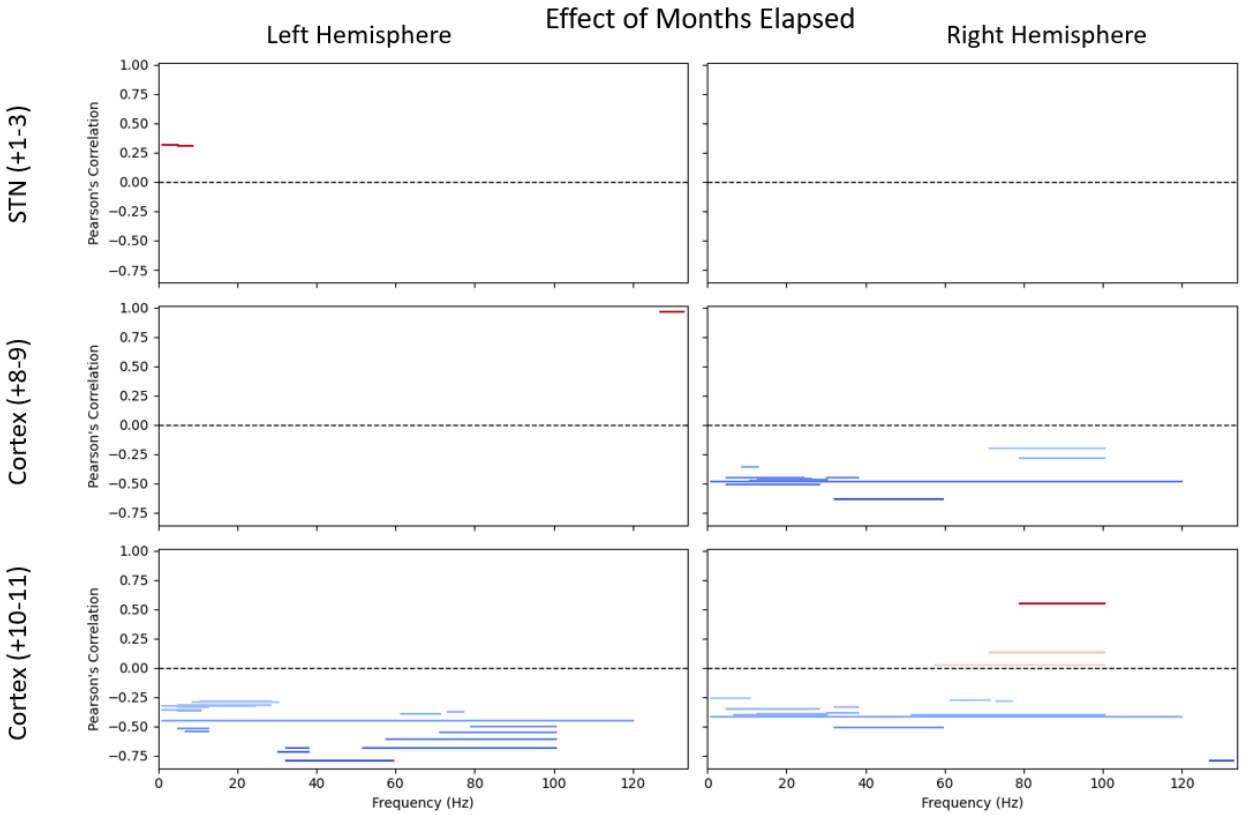

Figure 4.7: Significant effects of MonthsElapsed on the band power in all bands tested. Each horizontal line corresponds to a single trend. The horizontal extent shows the frequency extent of the power band, while the vertical location shows the effect size as a ratio between the band power variation due to the observed effect with the total variation observed. Only significant ($p<0.05$) effects are shown. Trends for each hemisphere are shown in the left and right columns respectively, and each row shows trends observed in a different recording location. Overall, there is a decrease in broadband power in most cortical recording locations.

4.3.2 LD Performance Predictions

For this analysis, we were primarily concerned with the effects on LD output over time. As a result, we only considered the time effects from the model above, as described in section 4.14. We extract the time dependent trends along with the effect of movement

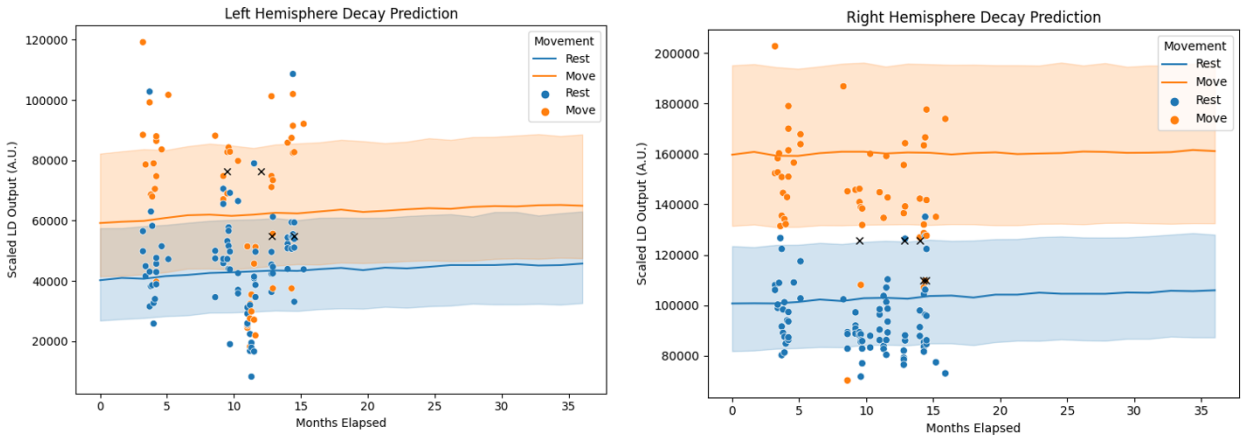

Figure 4.8: Linear discriminant output predictions. These are model predictions based on linear fits of individual band trends, for trends that were used in our example implementation of a LD for movement-driven aDBS for PD. Each solid line represents the median prediction for the LD output in the corresponding movement state over a series of months elapsed, with the shaded region indicating a quartile above and below. Dots represent logged LD outputs from the implanted device. Note the difference in separation between the classifiers used for the left and right side. Thanks to the design of the LD the consistent band decrease is mostly canceled out and the classifier remains reasonably stable.

to arrive at a model of how LD predictions will change over time. These predictions are based entirely on the LD's used during a move adaptive DBS trial we conducted [30]. Predictions based on this model are shown in Fig 4.8. Since the classifier was a differential classifier between primarily the beta and gamma bands, the effect of time on the scale of months was largely canceled out in the classifier outputs, leading to a stable LD output over time. The interaction terms between the movement condition and the number of months elapsed is close to zero, meaning that the two conditions do not approach each other at a significant rate. Importantly for our scenario, this means that the LD should continue working indefinitely, so long as the LD thresholds are adjusted to distinguish between the two conditions.

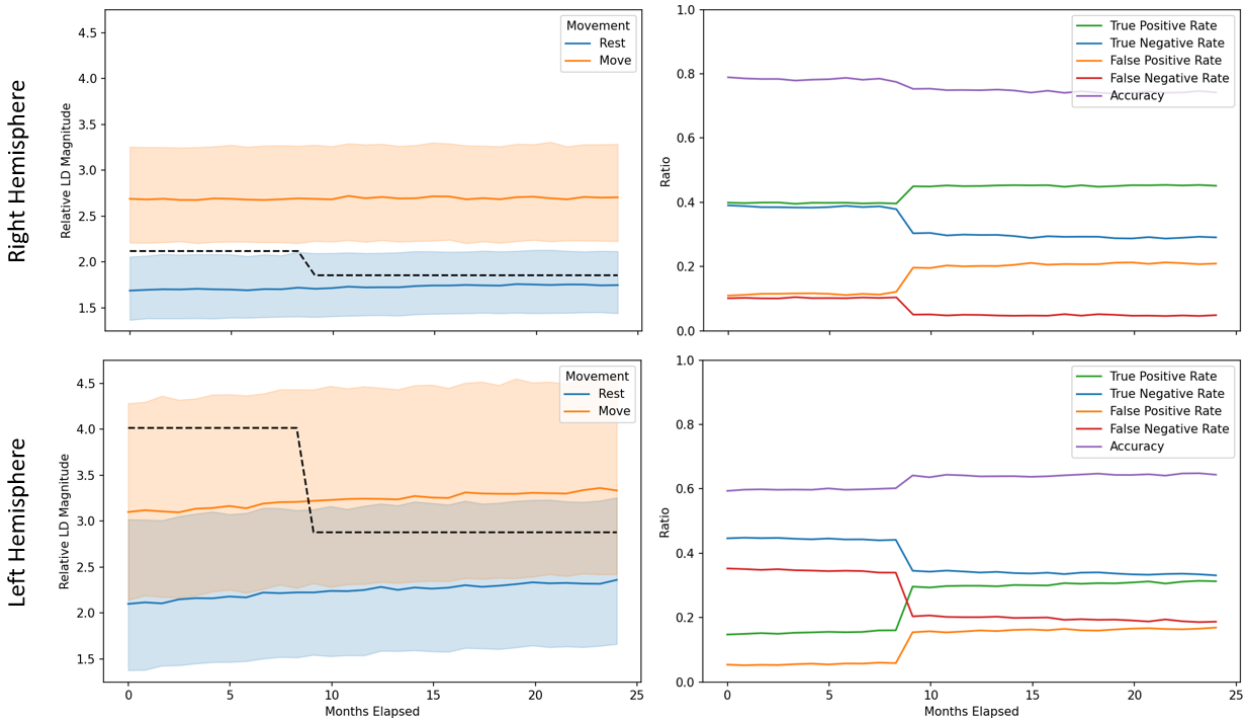

Figure 4.9: Performance of the implemented aDBS algorithm modeled over time with the real thresholds for both the left and right hemisphere. In the left two plots the the solid orange and blue lines show the median modeled LD output in the movement and rest cases. The shaded regions show the lower and upper quartiles. The decision threshold between movement and rest is shown with the dotted black line. The effect of the decision threshold on the overall accuracy and the individual positive and negative rates is shown on the right side. For both the left and the right side, the change in threshold decreased the false negative rate and increased the true positive rate. In the right hemisphere the adjustment came at the cost of overall accuracy, but still improved performance, as false negatives significantly detriment the patient's experience while false positive often have a minimal effect on the patient.

To test these predictions, we repeat the fitting process on a reduced dataset, and compare our predictions to our aDBS experiments. Our model is able to predict the

LD output, shown as the lines, when compared to LD logged from the device, and compared to LD values reconstructed from logged time domain values by the rcssim package. This particularly true for the right hemisphere, which in general presented better separation between the move and rest conditions. However, the change in the LD threshold used in the experiments cannot be explained by this model as it is shown here.

To more better understand the effect of the threshold and threshold change on the performance of the algorithm, we directly use the threshold on the modeled LD predictions to look at accuracy over time. These results for both the left and right hemisphere are shown in Fig 4.9. The decision threshold used for each discriminator over time is shown in the dotted line. The overall performance metrics, shown in two plots on the right, are much higher for the right hemisphere decoder as the move and rest conditions were better separable with the bands in used.

4.4 Discussion

4.4.1 Specific case of these LD predictions

The variation in predictions here could potentially be quite significant and depends entirely on the time-dependent trends observed in the data. In our case, all the bands that were included in the LD had either none or a small negative dependence on the number of months elapsed. It could theoretically be possible that some bands in a differently tuned LD could have both positive and negative coefficients, leading to a much more complicated picture. For some of these parameter combinations, it even becomes possible for the LD outputs to cross over, mandating complete re-training of the LD. In our case, this should be impossible, as beta desynchronization and gamma band power changes relative to movement are well documented and consistent phenomena.

4.4.2 Recommendations for aDBS

Our model can make some predictions on how long different kinds of power-band LD classifier based aDBS paradigms may remain effective. In particular, the reason the specific classifier used in this work remained as accurate for so long was because the classifier used a combination of the changes in the beta and gamma bands. Since the movement related change in these bands was opposite of each other, the weights in the classifier were opposite as well. This means that since all the bands had a similar decrease in band power over time, this effect largely canceled out in the LD outputs. As a result, some of the LD classifiers that are based on a single band, would lose accuracy much faster. In our patient, a classifier based purely on the beta band would have its accuracy decrease to half its original accuracy in about 10 months. This could be countered by incorporating an exponential decay parameter on the LD threshold. Such an LD could remain effective for much longer periods of time according to the predictions of our model.

It is likely that the relative long term stability of these two particular aDBS algorithms is due to the specific data they were trained on. Training was performed using a dataset that had been collected in two week-long intensive recording periods that were separated by a few months. When the Gaussian Process Bayesian Optimization process was used to find the best linear combination of power bands to separate movement and rest epochs, the result was naturally weighted to a linear combination of power bands that canceled out the wide-band decrease in power across all bands. Future analysis will hopefully be able to verify this hypothesis. This could also be used to bias the aDBS paradigm training process so that it could use a robust combination of power bands even with a more limited training data.

4.4.3 Limitations

While our model is able to predict the median LD outputs, it fails to account for the fact that a decrease in the threshold value led to an increase in prediction accuracy on both sides. This threshold decrease was made after a decrease in the LD performance was observed. The new threshold was set by re-running the training algorithm to include the open loop data collected in the months since the classifier was originally trained. Our model however does not predict that the threshold needed to be changed. It is possible that this was due to shorter term variation that was averaged out in our model. More work is needed to better understand this discrepancy.

This study is limited primarily by the number of patients analyzed. Due to logistical constraints, we were only able to collect data from one patient in the time available. This prevents us from being able to make any confident statements about how the observations with this patient would generalize to the population at large. This is especially true as in many ways our patient was rather unusual. They only took a tiny dose of medication each day, meaning that a day-scale trends as a result of medication wash in and out are likely much smaller than in a conventional patient. They were also very technically adept, and comfortable interacting with the technical components of the DBS systems. This meant they were a wonderful patient to test the remote data collection platform with since they thoroughly enjoyed the debugging process. This allowed us to fully streamline the deployment pipeline so future deployments are much simpler and smoother. We are currently working on instrumenting a second patient, and hope that future work will be able to repeat this analysis on more patient yielding greater statistical power.

Overall our work here is able to predict the changes in LD outputs based on independently recorded neural data. This method opens the door to developing automated methods to ensure that aDBS paradigms can be retrained automatically and remain effective for many months.

Chapter 5

OVERALL CONCLUSION

In this work we have successfully demonstrated the methods needed to understand the long term efficacy of aDBS. We have developed a framework for collecting, processing, and analyzing longitudinal data from patients with deep brain stimulation. Thanks to the particular implementations of the framework, we are able to track both neurological and symptoms changes in the patients as time progresses. We have shown that motor driven aDBS paradgims can be implemented in a way that is robust to longitudinal changes in the control signals. However, there remains a great amount of work that needs to be done.

One of the greatest outstanding questions is how well our results for both the ET and PD work generalize to the respective populations. To better understand this, the data collection framework we have built needs to be scaled up to collect from many more patients. A more comprehensive sample of the DBS patient population will allow the field to understand the level of inter-patient variation and common trends. Moreover, data should be collected from each patient for multiple years to more clearly see the effects of disease progression. To distinguish between longitudinal effects as a result of disease progression from any effects resulting from aging or the long-term influence of DBS, comparisons could be made across disease conditions: changes only robustly observed in PD patients are likely the cause of PD progression.

A data collection effort of this scale will require a significant investment of resources and a dedicated data management team members. Thanks to the design of the framework we developed here, the system can scale up by leveraging cloud architectures with relatively little effort. However, the cost of the storage space and processing

power needed to manage this data is a significant factor. Based on our work, a patient recording two days a week for 8 hours is expected to generate on the order of 50TB of data a year. Thanks to the relatively cheap price of cloud storage, storing this amount of data for years with a large number of patients is feasible with proper planning. However, processing this amount of data and preparing it for analysis will require a significant amount of computational power as well as constant dedicated oversight. It is our hope that this effort will be undertaken in the coming years to complete the dataset we have started here.

A large, at-home, cross-indication longitudinal dataset would prove invaluable to understanding the effects of aDBS in patient's daily lives as well as the long term effects of both stimulation and neurodegenerative disorders.

BIBLIOGRAPHY

[1] Ftc online archive, Jun 2023.

[2] Hippa online case violation examples, Jun 2023.

[3] Moshaddique Al Ameen, Jingwei Liu, and Kyungsup Kwak. Security and privacy issues in wireless sensor networks for healthcare applications. *Journal of medical systems*, 36:93–101, 2012.

[4] Bassam Al-Fatly, Siobhan Ewert, Dorothee Kübler, Daniel Kroneberg, Andreas Horn, and Andrea A Kühn. Connectivity profile of thalamic deep brain stimulation to effectively treat essential tremor. *Brain*, 142(10):3086–3098, 2019.

[5] Roger L Albin, Anne B Young, and John B Penney. The functional anatomy of basal ganglia disorders. *Trends in neurosciences*, 12(10):366–375, 1989.

[6] Anneke Alkemade, Gilles de Hollander, Steven Miletic, Max C Keuken, Rawien Balesar, Onno de Boer, Dick F Swaab, and Birte U Forstmann. The functional microscopic neuroanatomy of the human subthalamic nucleus. *Brain Structure and Function*, 224:3213–3227, 2019.

[7] Mattia Arlotti, Manuela Rosa, Sara Marceglia, Sergio Barbieri, and Alberto Priori. The adaptive deep brain stimulation challenge. *Parkinsonism and related disorders*, 28:12–17, 2016.

[8] Hayden Basinger and Joe Joseph. Neuroanatomy, subthalamic nucleus. 2020.

[9] Ishita Basu, Ali Yousefi, Britni Crocker, Rina Zelmann, Angelique C Paulk, Noam Peled, Kristen K Ellard, Daniel S Weisholtz, G Rees Cosgrove, Thilo Deckersbach, et al. Closed-loop enhancement and neural decoding of cognitive control in humans. *Nature biomedical engineering*, 7(4):576–588, 2023.

[10] Julian Benito-Leon. Essential tremor: a neurodegenerative disease? *Tremor and other hyperkinetic movements*, 4, 2014.

[11] Julian Benito-Leon and Elan D Louis. Essential tremor: emerging views of a common disorder. *Nature clinical practice Neurology*, 2(12):666–678, 2006.

[12] Dominik K Biezonski, Pierre Trifilieff, Jozsef Meszaros, Jonathan A Javitch, and Christoph Kellendonk. Evidence for limited d1 and d2 receptor coexpression and colocalization within the dorsal striatum of the neonatal mouse. *Journal of Comparative Neurology*, 523(8):1175–1189, 2015.

[13] Rafal Bogacz and Kevin Gurney. The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, 19(2):442–477, 2007.

[14] Nicolaas I Bohnen and Michele Hu. Sleep disturbance as potential risk and progression factor for parkinson's disease. *Journal of Parkinson's disease*, 9(3):603–614, 2019.

[15] Tamara Bonaci, Ryan Calo, and Howard Jay Chizeck. App stores for the brain: Privacy and security in brain-computer interfaces. In *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*, pages 1–7, 2014.

[16] Tamara Bonaci, Jeffrey Herron, Charlie Matlack, and Howard Jay Chizeck. Securing the exocortex: A twenty-first century cybernetics challenge. In *2014 IEEE Conference on Norbert Wiener in the 21st Century (21CW)*, pages 1–8, 2014.

[17] Walid Bouthour, Pierre Mégevand, John Donoghue, Christian Lüscher, Niels Birbaumer, and Paul Krack. Biomarkers for closed-loop deep brain stimulation in parkinson disease and beyond. *Nature Reviews Neurology*, 15(6):343–352, 2019.

[18] Timothy Brown. Building intricate partnerships with neurotechnology: Deep brain stimulation and relational agency. *IJFAB: International Journal of Feminist Approaches to Bioethics*, 13(1):134–154, 2020.

[19] Timothy Brown, Margaret C Thompson, Jeffrey Herron, Andrew Ko, Howard Chizeck, and Sara Goering. Controlling our brains–a case study on the implications of brain-computer interface-triggered deep brain stimulation for essential tremor. *Brain-Computer Interfaces*, 3(4):165–170, 2016.

[20] Timothy Emmanuel Brown. *Modulating Agency: the Moral & Aesthetic Import of Closed-Loop Deep-Brain Stimulation*. PhD thesis, 2019.

[21] Kenneth V Cartwright. Simpson's rule cumulative integration with ms excel and irregularly-spaced data.

[22] Sebastián Castaño-Candamil, Benjamin I Ferleger, Andrew Haddock, Sarah S Cooper, Jeffrey Herron, Andrew Ko, Howard J Chizeck, and Michael Tangermann. A pilot study on data-driven adaptive deep brain stimulation in chronically implanted essential tremor patients. *Frontiers in Human Neuroscience*, 14:541625, 2020.

[23] Y Chandu, K. S. Rakesh Kumar, Ninad Vivek Prabhukhanolkar, A N Anish, and Sushma Rawal. Design and implementation of hybrid encryption for security of iot data. In *2017 International Conference On Smart Technologies For Smart Nation (SmartTechCon)*, pages 1228–1231, 2017.

[24] K Ray Chaudhuri and Per Odin. The challenge of non-motor symptoms in parkinson's disease. *Progress in brain research*, 184:325–341, 2010.

[25] Witney Chen, Lowry Kirkby, Miro Kotzev, Patrick Song, Ro'ee Gilron, and Brian Pepin. The role of large-scale data infrastructure in developing next-generation deep brain stimulation therapies. *Frontiers in Human Neuroscience*, 15:717401, 2021.

[26] Tanusree Das, Jaclyn J Hwang, and Kathleen L Poston. Episodic recognition memory and the hippocampus in parkinson's disease: a review. *Cortex*, 113:191–209, 2019.

[27] Till A Dembek, Michael T Barbe, Mattias Åström, Mauritius Hoevels, Veerle Visser-Vandewalle, Gereon R Fink, and Lars Timmermann. Probabilistic mapping of deep brain stimulation effects in essential tremor. *NeuroImage: Clinical*, 13:164–173, 2017.

[28] Günther Deuschl and Rodger Elble. Essential tremor—neurodegenerative or nondegenerative disease towards a working definition of et. *Movement disorders: official journal of the Movement Disorder Society*, 24(14):2033–2041, 2009.

[29] Günther Deuschl, Jan Raethjen, Helge Hellriegel, and Rodger Elble. Treatment of patients with essential tremor. *The Lancet Neurology*, 10(2):148–161, 2011.

[30] Tanner Dixon, Gabrielle Strandquist, Alicia Zeng, Tomek Fraczek, Raphael Bechtold, Daryl Lawrence, Shravanan Ravi, Philip Starr, Jack Gallant, Jeffrey Herron, and Simon Little. Fully automated at-home optimization of movement-responsive deep brain stimulation for parkinson's disease. *In Preparation*, 2023.

[31] Jason A Donenfeld. Wireguard: next generation kernel network tunnel. In *NDSS*, pages 1–12, 2017.

[32] Sarah Donley, Sadie McGregor, Catherine Wielinski, and Martha Nance. Use and perceived effectiveness of complementary therapies in parkinson's disease. *Parkinsonism & related disorders*, 58:46–49, 2019.

[33] Osama Dorgham, Banan Al-Rahamneh, Ammar Almomani, Khalaf F Khatatneh, et al. Enhancing the security of exchanging and storing dicom medical images on the cloud. *International Journal of Cloud Applications and Computing (IJCAC)*, 8(1):154–172, 2018.

[34] Benoit Duchet, Filippo Ghezzi, Gihan Weerasinghe, Gerd Tinkhauser, Andrea A Kühn, Peter Brown, Christian Bick, and Rafal Bogacz. Average beta burst duration profiles provide a signature of dynamical changes between the on and off medication states in parkinson's disease. *PLoS computational biology*, 17(7):e1009116, 2021.

[35] Benoit Duchet, Gihan Weerasinghe, Hayriye Cagnan, Peter Brown, Christian Bick, and Rafal Bogacz. Phase-dependence of response curves to deep brain stimulation and their relationship: from essential tremor patient data to a wilson–cowan model. *The Journal of Mathematical Neuroscience*, 10:1–39, 2020.

[36] Nora El-Rashidy, Shaker El-Sappagh, SM Riazul Islam, Hazem M. El-Bakry, and Samir Abdelrazek. Mobile health in remote patient monitoring for chronic diseases: Principles, trends, and challenges. *Diagnostics*, 11(4):607, 2021.

[37] Khaled El Emam. Heuristics for de-identifying health data. *IEEE Security and Privacy*, 6(4):58–61, 2008.

[38] Alberto J Espay, Francesca Morgante, Aristide Merola, Alfonso Fasano, Luca Marsili, Susan H Fox, Erwan Bezard, Barbara Picconi, Paolo Calabresi, and Anthony E Lang. Levodopa-induced dyskinesia in parkinson disease: current and evolving concepts. *Annals of Neurology*, 84(6):797–811, 2018.

[39] S Fahn, Elton RMembers of the UPDRS Development Committee, et al. Recent developments in parkinson's disease. 1987; vol 2 florham park.

[40] Stanley Fahn, Eduardo Tolosa, Concepcíon Marín, et al. Clinical rating scale for tremor. *Parkinson's disease and movement disorders*, 2:271–280, 1993.

[41] Huang Fang. Managing data lakes in big data era: What's a data lake and why has it became popular in data management ecosystem. In *2015 IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*, pages 820–824, 2015.

[42] Frederico Arriaga Criscuoli de Farias, Carolina Matté Dagostini, Yan de Assunção Bicca, Vincenzo Fin Falavigna, and Asdrubal Falavigna. Remote patient monitoring: a systematic review. *Telemedicine and e-Health*, 26(5):576–583, 2020.

[43] Xiao-jiang Feng, Brian Greenwald, Herschel Rabitz, Eric Shea-Brown, and Robert Kosut. Toward closed-loop optimization of deep brain stimulation for parkinson's disease: concepts and lessons from a computational model. *Journal of neural engineering*, 4(2):L14, 2007.

[44] BI Ferleger, B Houston, MC Thompson, SS Cooper, KS Sonnet, AL Ko, JA Herron, and HJ Chizeck. Fully implanted adaptive deep brain stimulation in freely moving essential tremor patients. *Journal of Neural Engineering*, 17(5):056026, 2020.

[45] Pavel Filip, Ovidiu V Lungu, Mario-Ubaldo Manto, and Martin Bareš. Linking essential tremor to the cerebellum: physiological evidence. *The Cerebellum*, 15(6):774–780, 2016.

[46] Tomasz M Fraczek, Benjamin I Ferleger, Timothy E Brown, Margaret C Thompson, Andrew J Haddock, Brady C Houston, Jeffrey G Ojemann, Andrew L Ko, Jeffrey A Herron, and Howard J Chizeck. Closing the loop with cortical sensing: the development of adaptive deep brain stimulation for essential tremor using the activa pc+ s. *Frontiers in Neuroscience*, 15:749705, 2021.

[47] Charles R Gerfen. Segregation of d1 and d2 dopamine receptors in the striatal direct and indirect pathways: An historical perspective. *Frontiers in Synaptic Neuroscience*, 14:1002960, 2023.

[48] Frederic Gilbert, John Noel M Viaña, and Christian Ineichen. Deflating the "dbs causes personality changes" bubble. *Neuroethics*, 14(Suppl 1):1–17, 2021.

[49] Ro'ee Gilron, Simon Little, Randy Perrone, Robert Wilt, Coralie de Hemptinne, Maria S Yaroshinsky, Caroline A Racine, Sarah S Wang, Jill L Ostrem, Paul S Larson, et al. Long-term wireless streaming of neural recordings for circuit discovery and adaptive stimulation in individuals with parkinson's disease. *Nature biotechnology*, 39(9):1078–1085, 2021.

[50] Ro'ee Gilron, Simon Little, Robert Wilt, Randy Perrone, Juan Anso, and Philip A Starr. Sleep-aware adaptive deep brain stimulation control: chronic use at home with dual independent linear discriminate detectors. *Frontiers in Neuroscience*, 15:732499, 2021.

[51] Alexandre Gironell. The gaba hypothesis in essential tremor: lights and shadows. *Tremor and other hyperkinetic movements*, 4, 2014.

[52] Christopher G Goetz, Barbara C Tilley, Stephanie R Shaftman, Glenn T Stebbins, Stanley Fahn, Pablo Martinez-Martin, Werner Poewe, Cristina Sampaio, Matthew B Stern, Richard Dodel, et al. Movement disorder society-sponsored revision of the unified parkinson's disease rating scale (mds-updrs): scale presentation and clinimetric testing results. *Movement disorders: official journal of the Movement Disorder Society*, 23(15):2129–2170, 2008.

[53] Logan L Grado, Matthew D Johnson, and Theoden I Netoff. Bayesian adaptive dual control of deep brain stimulation in a computational model of parkinson's disease. *PLoS computational biology*, 14(12):e1006606, 2018.

[54] Henk J Groenewegen et al. The basal ganglia and motor control. *Neural plasticity*, 10(1-2):107–120, 2003.

[55] Andrew Haddock, Kyle T Mitchell, Andrew Miller, Jill L Ostrem, Howard J Chizeck, and Svjetlana Miocinovic. Automated deep brain stimulation programming for tremor. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(8):1618–1625, 2018.

[56] Ahnaf Rashik Hassan and Mohammed Imamul Hassan Bhuiyan. Automatic sleep stage classification. In *2015 2nd International Conference on Electrical Information and Communication Technologies (EICT)*, pages 211–216. IEEE, 2015.

[57] Dietrich Haubenberger and Mark Hallett. Essential tremor. *New England Journal of Medicine*, 378(19):1802–1810, 2018.

[58] Michael J Hawrylycz, Ed S Lein, Angela L Guillozet-Bongaarts, Elaine H Shen, Lydia Ng, Jeremy A Miller, Louie N Van De Lagemaat, Kimberly A Smith, Amanda Ebbert, Zackery L Riley, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*, 489(7416):391–399, 2012.

[59] Shenghong He, Fahd Baig, Abteen Mostofi, Alek Pogosyan, Jean Debarros, Alexander L Green, Tipu Z Aziz, Erlick Pereira, Peter Brown, and Huiling Tan. Closed-loop deep brain stimulation for essential tremor based on thalamic local field potentials. *Movement Disorders*, 36(4):863–873, 2021.

[60] Eva Henriksen, Tatjana M Burkow, Elin Johnsen, and Lars K Vognild. Privacy and information security risks in a technology platform for home-based chronic disease rehabilitation and education. *BMC medical informatics and decision making*, 13:1–13, 2013.

[61] Jeffrey Herron, Tim Denison, and Howard Jay Chizeck. Closed-loop dbs with movement intention. In *2015 7th international IEEE/EMBS conference on neural engineering (NER)*, pages 844–847. IEEE, 2015.

[62] Jeffrey A Herron, Margaret C Thompson, Timothy Brown, Howard J Chizeck, Jeffrey G Ojemann, and Andrew L Ko. Chronic electrocorticography for sensing movement intention and closed-loop deep brain stimulation with wearable sensors in an essential tremor patient. *Journal of neurosurgery*, 127(3):580–587, 2016.

[63] Jeffrey A Herron, Margaret C Thompson, Timothy Brown, Howard Jay Chizeck, Jeffrey G Ojemann, and Andrew L Ko. Cortical brain–computer interface for closed-loop deep brain stimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(11):2180–2187, 2017.

[64] Franziska Hopfner and Günther Deuschl. Is essential tremor a single entity? *European journal of neurology*, 25(1):71–82, 2018.

[65] Brady Houston, Margaret Thompson, Andrew Ko, and Howard Chizeck. A machine-learning approach to volitional control of a closed-loop deep brain stimulation system. *Journal of neural engineering*, 16(1):016004, 2019.

[66] Naroa Ibarretxe-Bilbao, Blanca Ramirez-Ruiz, Eduardo Tolosa, M Jose Marti, Francesc Valldeoriola, Nuria Bargallo, and Carme Junque. Hippocampal head atrophy predominance in parkinson's disease with hallucinations and with dementia. *Journal of neurology*, 255:1324–1331, 2008.

[67] Mohamed Fasil Ibrahim, Jessica C Beevis, and Ruth M Empson. Essential tremor–a cerebellar driven disorder? *Neuroscience*, 462:262–273, 2021.

[68] Nismat Javed and Marco Cascella. Neuroanatomy, globus pallidus. 2020.

[69] Nancy S Jecker and Andrew L Ko. Is that the same person? case studies in neurosurgery. *AJOB Neuroscience*, 8(3):160–170, 2017.

[70] Petr Ježek, Jeffery L Teeters, and Friedrich T Sommer. Nwb query engines: tools to search data stored in neurodata without borders format. *Frontiers in Neuroinformatics*, 14:27, 2020.

[71] Lyric A Jorgenson, William T Newsome, David J Anderson, Cornelia I Bargmann, Emery N Brown, Karl Deisseroth, John P Donoghue, Kathy L Hudson, Geoffrey SF Ling, Peter R MacLeish, et al. The brain initiative: developing technology to catalyse neuroscience discovery. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668):20140164, 2015.

[72] Yeun-Ho Joung. Development of implantable medical devices: from an engineering perspective. *International neurourology journal*, 17(3):98, 2013.

[73] Bonnie Kaplan. Phi protection under hipaa: An overall analysis. *Kaplan, B.(with appendix by Monteiro, APL)," PHI Protection under HIPAA: An Overall Analysis," LGPD na Saúde (LGPD Applicable to Health), Dallari, AB, Monaco, GFC, ed., São Paulo: Editora Revista dos Tribunais (Thomsom Reuters)*, 2021:61–88, 2020.

[74] Taha Khan, Ali Zeeshan, and Mark Dougherty. A novel method for automatic classification of parkinson gait severity using front-view video analysis. *Technology and Health Care*, 29(4):643–653, 2021.

[75] Eran Klein, Sara Goering, Josh Gagne, Conor V Shea, Rachel Franklin, Samuel Zorowitz, Darin D Dougherty, and Alik S Widge. Brain-computer interface-based control of closed-loop brain stimulation: attitudes and ethical considerations. *Brain-Computer Interfaces*, 3(3):140–148, 2016.

[76] Dylan L. Knowles, Kevin G. Stanley, and Nathaniel D. Osgood. A field-validated architecture for the collection of health-relevant behavioural data. In *2014 IEEE International Conference on Healthcare Informatics*, pages 79–88, 2014.

[77] Adam M Koss, Ron L Alterman, Michele Tagliati, and Jay L Shils. Calculating total electrical energy delivered by deep brain stimulation systems. *Annals of neurology*, 58(1):168, 2005.

[78] Oliver Kramer and Oliver Kramer. Scikit-learn. *Machine learning for evolution strategies*, pages 45–53, 2016.

[79] José L Lanciego, Natasha Luquin, and José A Obeso. Functional neuroanatomy of the basal ganglia. *Cold Spring Harbor perspectives in medicine*, 2(12), 2012.

[80] Abdeldjalil Ledmi, Hakim Bendjenna, and Sofiane Mounine Hemam. Fault tolerance in distributed systems: A survey. In *2018 3rd International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, pages 1–5, 2018.

[81] Abdeldjalil Ledmi, Hakim Bendjenna, and Sofiane Mounine Hemam. Fault tolerance in distributed systems: A survey. In *2018 3rd International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, pages 1–5, 2018.

[82] Rongheng Lin, Zezhou Ye, Hao Wang, and Budan Wu. Chronic diseases and health monitoring big data: A survey. *IEEE Reviews in Biomedical Engineering*, 11:275–288, 2018.

[83] Rongheng Lin, Zezhou Ye, Hao Wang, and Budan Wu. Chronic diseases and health monitoring big data: A survey. *IEEE Reviews in Biomedical Engineering*, 11:275–288, 2018.

[84] Simon Little, James Bonaiuto, Gareth Barnes, and Sven Bestmann. Human motor cortical beta bursts relate to movement planning and response errors. *PLoS biology*, 17(10):e3000479, 2019.

[85] Simon Little and Peter Brown. The functional role of beta oscillations in parkinson's disease. *Parkinsonism & related disorders*, 20:S44–S48, 2014.

[86] Simon Little and Peter Brown. Debugging adaptive deep brain stimulation for parkinson's disease. *Movement Disorders*, 35(4):555–561, 2020.

[87] Simon Little, Alex Pogosyan, Spencer Neal, Baltazar Zavala, Ludvic Zrinzo, Marwan Hariz, Thomas Foltynie, Patricia Limousin, Keyoumars Ashkan, James FitzGerald, et al. Adaptive deep brain stimulation in advanced parkinson disease. *Annals of neurology*, 74(3):449–457, 2013.

[88] Kenneth H Louie, Matthew N Petrucci, Logan L Grado, Chiahao Lu, Paul J Tuite, Andrew G Lamperski, Colum D MacKinnon, Scott E Cooper, and Theoden I Netoff. Semi-automated approaches to optimize deep brain stimulation parameters in parkinson's disease. *Journal of NeuroEngineering and Rehabilitation*, 18(1):83, 2021.

[89] Elan D Louis. Essential tremor. *New England Journal of Medicine*, 345(12):887–891, 2001.

[90] Elan D Louis. Re-thinking the biology of essential tremor: from models to morphology. *Parkinsonism & related disorders*, 20:S88–S93, 2014.

[91] Elan D Louis and Joaquim J Ferreira. How common is the most common adult movement disorder? update on the worldwide prevalence of essential tremor. *Movement Disorders*, 25(5):534–541, 2010.

[92] Elan D Louis and Jean Paul G Vonsattel. The emerging neuropathology of essential tremor. *Movement Disorders*, 23(2):174–182, 2008.

[93] Mark K Lyons. Deep brain stimulation: current and future clinical applications. In *Mayo Clinic Proceedings*, volume 86, pages 662–672. Elsevier, 2011.

[94] Lakmini P Malasinghe, Naeem Ramzan, and Keshav Dahal. Remote patient monitoring: a comprehensive study. *Journal of Ambient Intelligence and Humanized Computing*, 10:57–76, 2019.

[95] C Marras, JC Beck, JH Bower, E Roberts, B Ritz, GW Ross, RD Abbott, R Savica, SK Van Den Eeden, AW Willis, et al. Prevalence of parkinson's disease across north america. *NPJ Parkinson's disease*, 4(1):21, 2018.

[96] Pablo Martinez-Martin, Carmen Rodriguez-Blazquez, Mario Alvarez-Sanchez, Tomoko Arakaki, Alberto Bergareche-Yarza, Anabel Chade, Nelida Garretto, Oscar Gershanik, Monica M Kurtis, Juan Carlos Martinez-Castrillo, et al. Expanded and independent validation of the movement disorder society–unified parkinson's disease rating scale (mds-updrs). *Journal of neurology*, 260:228–236, 2013.

[97] Matthew M McGregor and Alexandra B Nelson. Circuit mechanisms of parkinson's disease. *Neuron*, 101(6):1042–1056, 2019.

[98] Cameron C McIntyre, Marc Savasta, Benjamin L Walter, and Jerrold L Vitek. How does deep brain stimulation work? present understanding and future questions. *Journal of clinical neurophysiology*, 21(1):40–50, 2004.

[99] Giulio Mecacci and WFG Haselager. Stimulating the self: The influence of conceptual frameworks on reactions to deep brain stimulation. *AJOB neuroscience*, 5(4):30–39, 2014.

[100] Anders Christian Meidahl, Gerd Tinkhauser, Damian Marc Herz, Hayriye Cagnan, Jean Debarros, and Peter Brown. Adaptive deep brain stimulation for movement disorders: the long road to clinical therapy. *Movement disorders*, 32(6):810–819, 2017.

[101] Anders Christian Meidahl, Gerd Tinkhauser, Damian Marc Herz, Hayriye Cagnan, Jean Debarros, and Peter Brown. Adaptive deep brain stimulation for movement disorders: the long road to clinical therapy. *Movement disorders*, 32(6):810–819, 2017.

[102] Luiz EAM Mello and João Villares. Neuroanatomy of the basal ganglia. *Psychiatric Clinics of North America*, 20(4):691–704, 1997.

[103] EH Middlebrooks, RA Domingo, T Vivas-Buitrago, L Okromelidze, T Tsuboi, JK Wong, RS Eisinger, L Almeida, MR Burns, A Horn, et al. Neuroimaging advances in deep brain stimulation: review of indications, anatomy, and brain connectomics. *American Journal of Neuroradiology*, 41(9):1558–1568, 2020.

[104] Erik H Middlebrooks, Lela Okromelidze, Joshua K Wong, Robert S Eisinger, Mathew R Burns, Ayushi Jain, Hsin-Pin Lin, Jun Yu, Enrico Opri, Andreas Horn, et al. Connectivity correlates to predict essential tremor deep brain stimulation outcome: Evidence for a common treatment pathway. *NeuroImage: Clinical*, 32:102846, 2021.

[105] Svjetlana Miocinovic, Nicole C Swann, Coralie de Hemptinne, Andrew Miller, Jill L Ostrem, and Philip A Starr. Cortical gamma oscillatrions in isolated dystonia. *Parkinsonism and Related Disorders*, 49:104–105, 2018.

[106] Aviv D Mizrahi-Kliger, Alexander Kaplan, Zvi Israel, Marc Deffains, and Hagai Bergman. Basal ganglia beta oscillations during sleep underlie parkinsonian insomnia. *Proceedings of the National Academy of Sciences*, 117(29):17359–17368, 2020.

[107] Erwin B Montgomery Jr and Kenneth B Baker. Mechanisms of deep brain stimulation and future technical developments. *Neurological research*, 22(3):259–266, 2000.

[108] Wilnellys Moore and Sarah Frye. Review of hipaa, part 1: history, protected health information, and privacy and security rules. *Journal of nuclear medicine technology*, 47(4):269–272, 2019.

[109] Wilnellys Moore and Sarah Frye. Review of hipaa, part 2: limitations, rights, violations, and role for the imaging technologist. *Journal of nuclear medicine technology*, 48(1):17–23, 2020.

[110] Fatemeh Nargesian, Erkang Zhu, Renée J Miller, Ken Q Pu, and Patricia C Arocena. Data lake management: challenges and opportunities. *Proceedings of the VLDB Endowment*, 12(12):1986–1989, 2019.

[111] Rachel Nosowsky and Thomas J Giordano. The health insurance portability and accountability act of 1996 (hipaa) privacy rule: implications for clinical research. *Annu. Rev. Med.*, 57:575–590, 2006.

[112] Michal Novotny, Tereza Tykalova, Hana Ruzickova, Evzen Ruzicka, Petr Dusek, and Jan Rusz. Automated video-based assessment of facial bradykinesia in de-novo parkinson's disease. *NPJ digital medicine*, 5(1):98, 2022.

[113] Ewan S Nurse, Sam E John, Dean R Freestone, Thomas J Oxley, Hoameng Ung, Samuel F Berkovic, Terence J O'Brien, Mark J Cook, and David B Grayden. Consistency of long-term subdural electrocorticography in humans. *IEEE Transactions on Biomedical Engineering*, 65(2):344–352, 2017.

[114] Movement Disorder Society Task Force on Rating Scales for Parkinson's Disease. The unified parkinson's disease rating scale (updrs): status and recommendations. *Movement Disorders*, 18(7):738–750, 2003.

[115] Enrico Opri, Stephanie Cernera, Rene Molina, Robert S Eisinger, Jackson N Cagle, Leonardo Almeida, Timothy Denison, Michael S Okun, Kelly D Foote, and Aysegul Gunduz. Chronic embedded cortico-thalamic closed-loop deep brain stimulation for the treatment of essential tremor. *Science translational medicine*, 12(572):eaay7680, 2020.

[116] Enrico Opri, Jonathan Shute, Rene Molina, Kelly Foote, Michael Okun, and Aysegul Gunduz. Closing the loop in deep brain stimulation: a responsive treatment for essential tremor (s27. 005), 2016.

[117] Ashwini Oswal, Chunyan Cao, Chien-Hung Yeh, Wolf-Julian Neumann, James Gratwicke, Harith Akram, Andreas Horn, Dianyou Li, Shikun Zhan, Chao Zhang, et al. Neural signatures of hyperdirect pathway activity in parkinson's disease. *Nature communications*, 12(1):5185, 2021.

[118] Ming-Kai Pan and Sheng-Han Kuo. Essential tremor: clinical perspectives and pathophysiology. *Journal of the Neurological Sciences*, 435:120198, 2022.

[119] Ming-Kai Pan, Yong-Shi Li, Shi-Bing Wong, Chun-Lun Ni, Yi-Mei Wang, Wen-Chuan Liu, Liang-Yin Lu, Jye-Chang Lee, Etty P Cortes, Jean-Paul G Vonsattel, et al. Cerebellar oscillations driven by synaptic pruning deficits of cerebellar climbing fibers contribute to tremor pathophysiology. *Science translational medicine*, 12(526):eaay1769, 2020.

[120] Susan K Patrick, Allen A Denington, Michel JA Gauthier, Deborah M Gillard, and Arthur Prochazka. Quantification of the updrs rigidity scale. *IEEE transactions on neural systems and rehabilitation engineering*, 9(1):31–41, 2001.

[121] Luis R Peraza, Kirsi M Kinnunen, Roisin McNaney, Ian J Craddock, Alan L Whone, Catherine Morgan, Richard Joules, and Robin Wolz. An automatic gait analysis pipeline for wearable sensors: A pilot study in parkinson's disease. *Sensors*, 21(24):8286, 2021.

[122] Alfredo J Perez, Sherali Zeadally, and Jonathan Cochran. A review and an empirical analysis of privacy policy and notices for consumer internet of things. *Security and Privacy*, 1(3):e15, 2018.

[123] Ronald F Pfeiffer. Non-motor symptoms in parkinson's disease. *Parkinsonism & related disorders*, 22:S119–S122, 2016.

[124] Gert Pfurtscheller and A Aranibar. Event-related cortical desynchronization detected by power measurements of scalp eeg. *Electroencephalography and clinical neurophysiology*, 42(6):817–826, 1977.

[125] Uyen Pham, Anne-Kristin Solbakk, Inger-Marie Skogseid, Mathias Toft, Are Hugo Pripp, Ane Eidahl Konglund, Stein Andersson, Ira Ronit Haraldsen, Dag Aarsland, Espen Dietrichs, et al. Personality changes after deep brain stimulation in parkinson's disease. *Parkinson's Disease*, 2015, 2015.

[126] Jan Raethjen and Günther Deuschl. The oscillating central network of essential tremor. *Clinical neurophysiology*, 123(1):61–64, 2012.

[127] Fernando Rodriguez, Shenghong He, and Huiling Tan. The potential of convolutional neural networks for identifying neural states based on electrophysiological signals: experiments on synthetic and real patient data. *Frontiers in Human Neuroscience*, 17:1134599, 2023.

[128] Carlos Oberdan Rolim, Fernando Luiz Koch, Carlos Becker Westphall, Jorge Werner, Armando Fracalossi, and Giovanni Schmitt Salvador. A cloud computing solution for patient's data collection in health care institutions. In *2010 Second International Conference on eHealth, Telemedicine, and Social Medicine*, pages 95–99, 2010.

[129] Manuela Rosa, Mattia Arlotti, Sara Marceglia, Filippo Cogiamanian, Gianluca Ardolino, Alessio Di Fonzo, Leonardo Lopiano, Emma Scelzo, Aristide Merola, Marco Locatelli, et al. Adaptive deep brain stimulation controls levodopa-induced side effects in parkinsonian patients. *Movement Disorders*, 32(4):628, 2017.

[130] Lorenzo Rossi, Guglielmo Foffani, SARA Marceglia, F Bracchi, S Barbieri, and A Priori. An electronic device for artefact suppression in human local field potential recordings during deep brain stimulation. *Journal of neural engineering*, 4(2):96, 2007.

[131] Oliver Rübel, Andrew Tritt, Benjamin Dichter, Thomas Braun, Nicholas Cain, Nathan Clack, Thomas J Davidson, Max Dougherty, Jean-Christophe Fillion-Robin, Nile Graddis, et al. Nwb: N 2.0: an accessible data standard for neurophysiology. *BioRxiv*, page 523035, 2019.

[132] Oliver Rübel, Andrew Tritt, Ryan Ly, Benjamin K Dichter, Satrajit Ghosh, Lawrence Niu, Pamela Baker, Ivan Soltesz, Lydia Ng, Karel Svoboda, et al. The neurodata without borders ecosystem for neurophysiological data science. *Elife*, 11:e78362, 2022.

[133] Pegdwendé Sawadogo and Jérôme Darmont. On data lake architectures and metadata management. *Journal of Intelligent Information Systems*, 56:97–120, 2021.

[134] Mohamed Seliem, Khalid Elgazzar, and Kasem Khalil. Towards privacy preserving iot environments: a survey. *Wireless Communications and Mobile Computing*, 2018:1–15, 2018.

[135] Jahanzeb Shahid, Rizwan Ahmad, Adnan K Kiani, Tahir Ahmad, Saqib Saeed, and Abdullah M Almuhaideb. Data protection and privacy of the internet of healthcare things (iohts). *Applied Sciences*, 12(4):1927, 2022.

[136] Jung Hwan Shin, Ri Yu, Jed Noel Ong, Chan Young Lee, Seung Ho Jeon, Hwanpil Park, Han-Joon Kim, Jehee Lee, and Beomseok Jeon. Quantitative gait analysis using a pose-estimation algorithm with a single 2d-video of parkinson's disease patients. *Journal of Parkinson's Disease*, 11(3):1271–1283, 2021.

[137] Otis M Solomon Jr. Psd computations using welch's method. *NASA STI/Recon Technical Report N*, 92:23584, 1991.

[138] James Sonne, Vamsi Reddy, and Morris R Beato. Neuroanatomy, substantia nigra. In *StatPearls [Internet]*. StatPearls Publishing, 2022.

[139] María Carolina Sepúlveda Soto and Alfonso Fasano. Essential tremor: new advances. *Clinical Parkinsonism & Related Disorders*, 3:100031, 2020.

[140] John A Stankovic, Quihia Cao, T Doan, Lei Fang, Zhijun He, R Kiran, S Lin, S Son, R Stoleru, and A Wood. Wireless sensor networks for in-home healthcare: Potential and challenges. In *High confidence medical device software and systems (HCMDSS) workshop*, volume 2. Penn Library Wolverhampton, UK, 2005.

[141] Scott Stanslaski, Pedram Afshar, Peng Cong, Jon Giftakis, Paul Stypulkowski, Dave Carlson, Dave Linde, Dave Ullestad, Al-Thaddeus Avestruz, and Timothy Denison. Design and validation of a fully implantable, chronic, closed-loop neuromodulation device with concurrent sensing and stimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):410–421, 2012.

[142] Alessandro Stefani, Rocco Cerroni, Paolo Mazzone, Claudio Liguori, Giuseppe Di Giovanni, Mariangela Pierantozzi, and Salvatore Galati. Mechanisms of action underlying the efficacy of deep brain stimulation of the subthalamic nucleus in parkinson's disease: central role of disease severity. *European Journal of Neuroscience*, 49(6):805–816, 2019.

[143] Gabrielle Strandquist, Tanner Dixon, Tomasz Frączek, Shravanan Ravi, Alicia Zeng, Raphael Bechtold, Daryl Lawrence, Simon Little, Jack Gallant, and Jeffrey Herron. In-home video and imu kinematics of self guided tasks correlate with clinical bradykinesia scores. In *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 1–6. IEEE, 2023.

[144] Gabrielle Strandquist, Tomasz Frączek, Tanner Dixon, Shravanan Ravi, Raphael Bechtold, Daryl Lawrence, Alicia Zeng, Jack Gallant, Simon Little, and Jeffrey Herron. Bringing the clinic home using an at-home multi-modal data collection ecosystem to support adaptive deep brain stimulation. *Journal of Video Essays*, 2023.

[145] Lynn K Struck, Robert L Rodnitzky, and Judith K Dobson. Circadian fluctuations of contrast sensitivity in parkinson's disease. *Neurology*, 40(3 Part 1):467–467, 1990.

[146] Wencheng Sun, Zhiping Cai, Yangyang Li, Fang Liu, Shengqun Fang, and Guoyan Wang. Security and privacy in the medical internet of things: a review. *Security and Communication Networks*, 2018:1–9, 2018.

[147] Wencheng Sun, Zhiping Cai, Yangyang Li, Fang Liu, Shengqun Fang, and Guoyan Wang. Security and privacy in the medical internet of things: a review. *Security and Communication Networks*, 2018:1–9, 2018.

[148] Nicole C Swann, Coralie De Hemptinne, Svjetlana Miocinovic, Salman Qasim, Sarah S Wang, Nathan Ziman, Jill L Ostrem, Marta San Luciano, Nicholas B Galifianakis, and Philip A Starr. Gamma oscillations in the hyperkinetic state detected with chronic human brain recordings in parkinson's disease. *Journal of Neuroscience*, 36(24):6445–6458, 2016.

[149] Lo'ai Tawalbeh, Fadi Muheidat, Mais Tawalbeh, and Muhannad Quwaider. Iot privacy and security: Challenges and solutions. *Applied Sciences*, 10(12):4102, 2020.

[150] Jeffery L Teeters, Keith Godfrey, Rob Young, Chinh Dang, Claudia Friedsam, Barry Wark, Hiroki Asari, Simon Peron, Nuo Li, Adrien Peyrache, et al. Neurodata without borders: creating a common data format for neurophysiology. *Neuron*, 88(4):629–634, 2015.

[151] Jeffery L Teeters, Keith Godfrey, Rob Young, Chinh Dang, Claudia Friedsam, Barry Wark, Hiroki Asari, Simon Peron, Nuo Li, Adrien Peyrache, et al. Neurodata without borders: creating a common data format for neurophysiology. *Neuron*, 88(4):629–634, 2015.

[152] Bhomraj Thanvi, Nelson Lo, and Tom Robinson. Essential tremor—the most common movement disorder in older people. *Age and ageing*, 35(4):344–349, 2006.

[153] Gerd Tinkhauser, Alek Pogosyan, Simon Little, Martijn Beudel, Damian M Herz, Huiling Tan, and Peter Brown. The modulatory effect of adaptive deep brain stimulation on beta bursts in parkinson's disease. *Brain*, 140(4):1053–1067, 2017.

[154] Gerd Tinkhauser, Alek Pogosyan, Huiling Tan, Damian M Herz, Andrea A Kühn, and Peter Brown. Beta burst dynamics in parkinson's disease off and on dopaminergic medication. *Brain*, 140(11):2968–2981, 2017.

[155] Gerd Tinkhauser, Flavie Torrecillos, Yann Duclos, Huiling Tan, Alek Pogosyan, Petra Fischer, Romain Carron, Marie-Laure Welter, Carine Karachi, Wim Vandenberghe, et al. Beta burst coupling across the motor circuit in parkinson's disease. *Neurobiology of disease*, 117:217–225, 2018.

[156] Camilo Toro, Günther Deuschl, Robert Thatcher, Susumu Sato, Conrad Kufta, and Mark Hallett. Event-related desynchronization and movement-related cortical potentials on the ecog and eeg. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 93(5):380–389, 1994.

[157] Flavie Torrecillos, Gerd Tinkhauser, Petra Fischer, Alexander L Green, Tipu Z Aziz, Thomas Foltynie, Patricia Limousin, Ludvic Zrinzo, Keyoumars Ashkan, Peter Brown, et al. Modulation of beta bursts in the subthalamic nucleus predicts motor performance. *Journal of neuroscience*, 38(41):8905–8917, 2018.

[158] Tyler J Torrico and Sunil Munakomi. Neuroanatomy, thalamus. 2019.

[159] Kaviraja Udupa and Robert Chen. The mechanisms of action of deep brain stimulation and ideas for the future development. *Progress in neurobiology*, 133:27–49, 2015.

[160] Ata Ullah, Muhammad Azeem, Humaira Ashraf, Abdulellah A. Alaboudi, Mamoona Humayun, and NZ Jhanjhi. Secure healthcare data aggregation and transmission in iot—a survey. *IEEE Access*, 9:16849–16865, 2021.

[161] Ata Ullah, Muhammad Azeem, Humaira Ashraf, Abdulellah A. Alaboudi, Mamoona Humayun, and NZ Jhanjhi. Secure healthcare data aggregation and transmission in iot—a survey. *IEEE Access*, 9:16849–16865, 2021.

[162] Julian Unterweger, Martin Seeber, Stavros Zanos, Jeffrey G Ojemann, and Reinhold Scherer. Ecog beta suppression and modulation during finger extension and flexion. *Frontiers in neuroscience*, 14:35, 2020.

[163] A Velisar, J Syrkin-Nikolau, Z Blumenfeld, MH Trager, MF Afzal, V Prabhakar, and H Bronte-Stewart. Dual threshold neural closed loop deep brain stimulation in parkinson disease patients. *Brain stimulation*, 12(4):868–876, 2019.

[164] Nataraj Venkataramanan and Ashwin Shriram. *Data privacy: principles and practice*. CRC Press, 2016.

[165] Martina Vettoretti and Andrea Facchinetti. Combining continuous glucose monitoring and insulin pumps to automatically tune the basal insulin infusion in diabetes therapy: a review. *Biomedical engineering online*, 18:1–17, 2019.

[166] Aleksandar Videnovic and Diego Golombek. Circadian and sleep disorders in parkinson's disease. *Experimental neurology*, 243:45–56, 2013.

[167] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.

[168] Thomas Welton, Francisco Cardoso, Jonathan A Carr, Ling-Ling Chan, Günther Deuschl, Joseph Jankovic, and Eng-King Tan. Essential tremor. *Nature Reviews Disease Primers*, 7(1):83, 2021.

[169] Diane Whitmer, Camille De Solages, Bruce Hill, Hong Yu, Jaimie M Henderson, and Helen Bronte-Stewart. High frequency deep brain stimulation attenuates subthalamic and cortical rhythms in parkinson's disease. *Frontiers in human neuroscience*, 6:155, 2012.

[170] Takamitsu Yamamoto, Yoichi Katayama, Junichi Ushiba, Hiroko Yoshino, Toshiki Obuchi, Kazutaka Kobayashi, Hideki Oshima, and Chikashi Fukaya. On-demand control system for deep brain stimulation for treatment of intention tremor. *Neuromodulation: Technology at the Neural Interface*, 16(3):230–235, 2013.

[171] Nada Yousif, Michael Mace, Nicola Pavese, Roman Borisyuk, Dipankar Nandi, and Peter Bain. A network model of local field potential activity in essential tremor and the impact of deep brain stimulation. *PLoS computational biology*, 13(1):e1005326, 2017.

[172] Hengameh Zahed, Jose Rafael Pantoja Zuzuarregui, Ro'ee Gilron, Timothy Denison, Philip A Starr, and Simon Little. The neurophysiology of sleep in parkinson's disease. *Movement Disorders*, 36(7):1526–1542, 2021.

APPENDIX A

Summary of all significant (p>0.05) trends observed in the longitudinal data collected as described in Chapter 4. Units of the slope value depend on the x-variable. MonthsElapsed is measured in months, StimAmpl is measured in milliamps, HoursOfDay is measured in hours. The is—- variables are all measured as the difference between 0 and 1. Combo variables are the multiples of the two individual units. The y-axis is always band power. Ranges and distributions for variable are shown in Fig 4.2.

Hemisphere	Electrode	Power Band	X-Variable	Pearson R	Raw Slope
Left	"(1, 3)"	0.98Hz - 4.88Hz	MonthsElapsed	0.0503	0.0647
Left	"(1, 3)"	4.88Hz - 8.79Hz	MonthsElapsed	0.2797	0.0718
Right	"(1, 3)"	73.24Hz - 77.14Hz	MonthsElapsed	-0.2618	-0.0307
Right	"(8, 9)"	0.98Hz - 120Hz	MonthsElapsed	-0.5678	-0.0270
Right	"(8, 9)"	4.88Hz - 24.41Hz	MonthsElapsed	-0.4396	-0.0272
Right	"(8, 9)"	4.88Hz - 28.32Hz	MonthsElapsed	-0.4971	-0.0316
Right	"(8, 9)"	8.79Hz - 12.7Hz	MonthsElapsed	-0.3479	-0.0240
Right	"(8, 9)"	8.79Hz - 30.27Hz	MonthsElapsed	-0.4681	-0.0318
Right	"(8, 9)"	10.74Hz - 28.32Hz	MonthsElapsed	-0.4508	-0.0313
Right	"(8, 9)"	12.7Hz - 30.27Hz	MonthsElapsed	-0.4474	-0.0331
Right	"(8, 9)"	12.7Hz - 26.37Hz	MonthsElapsed	-0.4395	-0.0335
Right	"(8, 9)"	30.27Hz - 38.09Hz	MonthsElapsed	-0.4648	-0.0263
Right	"(8, 9)"	32.23Hz - 59.57Hz	MonthsElapsed	-0.6463	-0.0269
Right	"(8, 9)"	79.1Hz - 100.59Hz	MonthsElapsed	-0.2120	-0.0099
Left	"(8, 9)"	127Hz - 133Hz	MonthsElapsed	0.9681	0.8937
Left	"(10, 11)"	0.98Hz - 4.88Hz	MonthsElapsed	-0.3348	-0.0391
Left	"(10, 11)"	0.98Hz - 10.74Hz	MonthsElapsed	-0.3865	-0.0294
Right	"(10, 11)"	0.98Hz - 10.74Hz	MonthsElapsed	-0.2376	-0.0151
Left	"(10, 11)"	0.98Hz - 120Hz	MonthsElapsed	-0.3443	-0.0319
Right	"(10, 11)"	0.98Hz - 120Hz	MonthsElapsed	-0.4027	-0.0278
Left	"(10, 11)"	4.88Hz - 8.79Hz	MonthsElapsed	-0.3441	-0.0268
Right	"(10, 11)"	4.88Hz - 8.79Hz	MonthsElapsed	-0.2299	-0.0184
Left	"(10, 11)"	4.88Hz - 10.74Hz	MonthsElapsed	-0.3788	-0.0366
Left	"(10, 11)"	4.88Hz - 12.7Hz	MonthsElapsed	-0.3925	-0.0396
Left	"(10, 11)"	4.88Hz - 24.41Hz	MonthsElapsed	-0.3204	-0.0251
Right	"(10, 11)"	4.88Hz - 24.41Hz	MonthsElapsed	-0.3712	-0.0346
Left	"(10, 11)"	4.88Hz - 28.32Hz	MonthsElapsed	-0.3165	-0.0248
Right	"(10, 11)"	4.88Hz - 28.32Hz	MonthsElapsed	-0.3385	-0.0299
Left	"(10, 11)"	6.84Hz - 12.7Hz	MonthsElapsed	-0.5616	-0.0502
Left	"(10, 11)"	8.79Hz - 12.7Hz	MonthsElapsed	-0.3594	-0.0536

Left	"(10, 11)"	8.79Hz - 30.27Hz	MonthsElapsed	-0.2884	-0.0281
Left	"(10, 11)"	10.74Hz - 28.32Hz	MonthsElapsed	-0.2723	-0.0283
Left	"(10, 11)"	12.7Hz - 30.27Hz	MonthsElapsed	-0.2369	-0.0258
Left	"(10, 11)"	12.7Hz - 26.37Hz	MonthsElapsed	-0.2504	-0.0290
Left	"(10, 11)"	30.27Hz - 38.09Hz	MonthsElapsed	-0.6969	-0.0569
Right	"(10, 11)"	30.27Hz - 38.09Hz	MonthsElapsed	-0.3748	-0.0259
Left	"(10, 11)"	32.23Hz - 59.57Hz	MonthsElapsed	-0.7997	-0.0567
Right	"(10, 11)"	32.23Hz - 59.57Hz	MonthsElapsed	-0.4929	-0.0244
Left	"(10, 11)"	32.23Hz - 38.09Hz	MonthsElapsed	-0.6539	-0.0615
Right	"(10, 11)"	32.23Hz - 38.09Hz	MonthsElapsed	-0.3202	-0.0250
Left	"(10, 11)"	51.76Hz - 100.59Hz	MonthsElapsed	-0.6702	-0.0404
Right	"(10, 11)"	51.76Hz - 100.59Hz	MonthsElapsed	-0.3499	-0.0033
Left	"(10, 11)"	57.62Hz - 100.59Hz	MonthsElapsed	-0.6083	-0.0383
Right	"(10, 11)"	57.62Hz - 100.59Hz	MonthsElapsed	-0.0776	0.0105
Left	"(10, 11)"	61.52Hz - 71.29Hz	MonthsElapsed	-0.3070	-0.0245
Left	"(10, 11)"	71.29Hz - 100.59Hz	MonthsElapsed	-0.5673	-0.0282
Right	"(10, 11)"	71.29Hz - 100.59Hz	MonthsElapsed	0.3908	0.0497
Left	"(10, 11)"	73.24Hz - 77.14Hz	MonthsElapsed	-0.3773	-0.0333
Right	"(10, 11)"	73.24Hz - 77.14Hz	MonthsElapsed	-0.1996	-0.0031
Left	"(10, 11)"	79.1Hz - 100.59Hz	MonthsElapsed	-0.4917	-0.0246
Right	"(10, 11)"	79.1Hz - 100.59Hz	MonthsElapsed	0.8584	0.0699
Right	"(10, 11)"	127Hz - 133Hz	MonthsElapsed	-0.9154	-0.3242
Left	"(1, 3)"	0.98Hz - 4.88Hz	StimAmpl	0.8764	2.5033
Right	"(1, 3)"	0.98Hz - 4.88Hz	StimAmpl	0.8696	1.3262
Left	"(1, 3)"	0.98Hz - 10.74Hz	StimAmpl	0.8952	2.7750
Right	"(1, 3)"	0.98Hz - 10.74Hz	StimAmpl	0.9321	2.0077
Left	"(1, 3)"	0.98Hz - 120Hz	StimAmpl	0.9057	2.9040
Right	"(1, 3)"	0.98Hz - 120Hz	StimAmpl	0.9412	2.3368
Left	"(1, 3)"	4.88Hz - 8.79Hz	StimAmpl	0.9089	3.2512
Right	"(1, 3)"	4.88Hz - 8.79Hz	StimAmpl	0.9293	2.2023
Left	"(1, 3)"	4.88Hz - 10.74Hz	StimAmpl	0.9077	3.2853
Right	"(1, 3)"	4.88Hz - 10.74Hz	StimAmpl	0.9299	2.3324
Left	"(1, 3)"	4.88Hz - 12.7Hz	StimAmpl	0.9104	3.3567
Right	"(1, 3)"	4.88Hz - 12.7Hz	StimAmpl	0.9373	2.3287
Left	"(1, 3)"	4.88Hz - 24.41Hz	StimAmpl	0.9105	3.2365
Right	"(1, 3)"	4.88Hz - 24.41Hz	StimAmpl	0.9584	2.6082
Left	"(1, 3)"	4.88Hz - 28.32Hz	StimAmpl	0.9102	3.2303
Right	"(1, 3)"	4.88Hz - 28.32Hz	StimAmpl	0.9585	2.6076
Left	"(1, 3)"	6.84Hz - 12.7Hz	StimAmpl	0.9169	3.5019
Right	"(1, 3)"	6.84Hz - 12.7Hz	StimAmpl	0.9343	2.5339
Left	"(1, 3)"	8.79Hz - 12.7Hz	StimAmpl	0.9184	3.6572
Right	"(1, 3)"	8.79Hz - 12.7Hz	StimAmpl	0.9327	2.5547

Left	"(1, 3)"	8.79Hz - 30.27Hz	StimAmpl	0.9096	3.2348
Right	"(1, 3)"	8.79Hz - 30.27Hz	StimAmpl	0.9552	2.7376
Left	"(1, 3)"	10.74Hz - 28.32Hz	StimAmpl	0.9077	3.2102
Right	"(1, 3)"	10.74Hz - 28.32Hz	StimAmpl	0.9556	2.7775
Left	"(1, 3)"	12.7Hz - 30.27Hz	StimAmpl	0.9030	3.0553
Right	"(1, 3)"	12.7Hz - 30.27Hz	StimAmpl	0.9546	2.8347
Left	"(1, 3)"	12.7Hz - 26.37Hz	StimAmpl	0.9010	3.0543
Right	"(1, 3)"	12.7Hz - 26.37Hz	StimAmpl	0.9541	2.8603
Left	"(1, 3)"	30.27Hz - 38.09Hz	StimAmpl	0.8759	2.4928
Right	"(1, 3)"	30.27Hz - 38.09Hz	StimAmpl	0.8516	1.6858
Left	"(1, 3)"	32.23Hz - 59.57Hz	StimAmpl	0.8697	2.3750
Right	"(1, 3)"	32.23Hz - 59.57Hz	StimAmpl	0.9149	2.0261
Left	"(1, 3)"	32.23Hz - 38.09Hz	StimAmpl	0.8785	2.4732
Right	"(1, 3)"	32.23Hz - 38.09Hz	StimAmpl	0.8222	1.6761
Left	"(1, 3)"	51.76Hz - 100.59Hz	StimAmpl	0.8324	1.9649
Right	"(1, 3)"	51.76Hz - 100.59Hz	StimAmpl	0.9527	1.9359
Left	"(1, 3)"	57.62Hz - 100.59Hz	StimAmpl	0.8250	1.8705
Right	"(1, 3)"	57.62Hz - 100.59Hz	StimAmpl	0.9558	1.8974
Left	"(1, 3)"	61.52Hz - 71.29Hz	StimAmpl	0.8137	1.8358
Right	"(1, 3)"	61.52Hz - 71.29Hz	StimAmpl	0.9462	1.9453
Left	"(1, 3)"	71.29Hz - 100.59Hz	StimAmpl	0.9262	3.0227
Right	"(1, 3)"	71.29Hz - 100.59Hz	StimAmpl	0.9435	1.6236
Left	"(1, 3)"	73.24Hz - 77.14Hz	StimAmpl	0.8253	1.9415
Right	"(1, 3)"	73.24Hz - 77.14Hz	StimAmpl	0.9180	1.9061
Left	"(1, 3)"	79.1Hz - 100.59Hz	StimAmpl	0.8552	2.0127
Right	"(1, 3)"	79.1Hz - 100.59Hz	StimAmpl	0.9450	1.6304
Left	"(1, 3)"	127Hz - 133Hz	StimAmpl	0.7435	1.7403
Right	"(1, 3)"	127Hz - 133Hz	StimAmpl	0.7141	0.1853
Left	"(8, 9)"	0.98Hz - 4.88Hz	StimAmpl	-0.3005	-0.2291
Right	"(8, 9)"	0.98Hz - 4.88Hz	StimAmpl	-0.2665	-0.1978
Left	"(8, 9)"	0.98Hz - 10.74Hz	StimAmpl	-0.4011	-0.1993
Right	"(8, 9)"	0.98Hz - 10.74Hz	StimAmpl	-0.3645	-0.0748
Left	"(8, 9)"	0.98Hz - 120Hz	StimAmpl	-0.1824	-0.1096
Right	"(8, 9)"	0.98Hz - 120Hz	StimAmpl	-0.1828	-0.0466
Left	"(8, 9)"	4.88Hz - 8.79Hz	StimAmpl	-0.3532	-0.2460
Left	"(8, 9)"	4.88Hz - 10.74Hz	StimAmpl	-0.3301	0.1516
Left	"(8, 9)"	4.88Hz - 12.7Hz	StimAmpl	0.3964	0.1936
Left	"(8, 9)"	4.88Hz - 24.41Hz	StimAmpl	-0.2508	-0.1174
Left	"(8, 9)"	4.88Hz - 28.32Hz	StimAmpl	-0.2383	-0.1191
Left	"(8, 9)"	8.79Hz - 30.27Hz	StimAmpl	-0.1880	-0.1257
Left	"(8, 9)"	10.74Hz - 28.32Hz	StimAmpl	-0.1818	-0.1262
Left	"(8, 9)"	12.7Hz - 30.27Hz	StimAmpl	-0.1695	-0.1242

Left	"(8, 9)"	12.7Hz - 26.37Hz	StimAmpl	-0.1907	-0.1447
Left	"(8, 9)"	51.76Hz - 100.59Hz	StimAmpl	0.1614	0.1324
Left	"(8, 9)"	57.62Hz - 100.59Hz	StimAmpl	0.1782	0.1133
Left	"(8, 9)"	61.52Hz - 71.29Hz	StimAmpl	0.2649	0.1867
Right	"(8, 9)"	71.29Hz - 100.59Hz	StimAmpl	-0.0352	-0.0094
Right	"(8, 9)"	79.1Hz - 100.59Hz	StimAmpl	-0.1603	-0.0476
Left	"(8, 9)"	127Hz - 133Hz	StimAmpl	0.9283	3.5142
Right	"(8, 9)"	127Hz - 133Hz	StimAmpl	0.4931	0.6572
Left	"(10, 11)"	0.98Hz - 4.88Hz	StimAmpl	-0.3082	-0.2167
Right	"(10, 11)"	0.98Hz - 4.88Hz	StimAmpl	-0.3518	-0.2221
Left	"(10, 11)"	0.98Hz - 10.74Hz	StimAmpl	-0.3965	-0.1920
Right	"(10, 11)"	0.98Hz - 10.74Hz	StimAmpl	-0.3541	-0.1493
Left	"(10, 11)"	0.98Hz - 120Hz	StimAmpl	-0.3215	-0.1288
Right	"(10, 11)"	0.98Hz - 120Hz	StimAmpl	-0.1735	-0.0713
Left	"(10, 11)"	4.88Hz - 8.79Hz	StimAmpl	-0.2563	-0.1178
Right	"(10, 11)"	4.88Hz - 8.79Hz	StimAmpl	-0.3129	-0.1641
Left	"(10, 11)"	4.88Hz - 10.74Hz	StimAmpl	-0.2514	-0.1132
Right	"(10, 11)"	4.88Hz - 10.74Hz	StimAmpl	-0.3844	-0.1947
Left	"(10, 11)"	4.88Hz - 12.7Hz	StimAmpl	-0.1248	-0.0645
Right	"(10, 11)"	4.88Hz - 12.7Hz	StimAmpl	-0.2505	-0.1035
Left	"(10, 11)"	4.88Hz - 24.41Hz	StimAmpl	-0.2656	-0.1244
Left	"(10, 11)"	4.88Hz - 28.32Hz	StimAmpl	-0.2749	-0.1292
Right	"(10, 11)"	4.88Hz - 28.32Hz	StimAmpl	-0.1474	-0.0792
Left	"(10, 11)"	6.84Hz - 12.7Hz	StimAmpl	-0.0738	-0.0333
Left	"(10, 11)"	8.79Hz - 30.27Hz	StimAmpl	-0.1765	-0.1018
Left	"(10, 11)"	10.74Hz - 28.32Hz	StimAmpl	-0.1859	-0.1152
Left	"(10, 11)"	12.7Hz - 30.27Hz	StimAmpl	-0.1753	-0.1144
Left	"(10, 11)"	12.7Hz - 26.37Hz	StimAmpl	-0.1798	-0.1246
Left	"(10, 11)"	32.23Hz - 59.57Hz	StimAmpl	-0.2784	-0.0750
Left	"(10, 11)"	61.52Hz - 71.29Hz	StimAmpl	0.1997	0.0941
Right	"(10, 11)"	71.29Hz - 100.59Hz	StimAmpl	0.3376	0.2110
Right	"(10, 11)"	79.1Hz - 100.59Hz	StimAmpl	0.7276	0.2832
Left	"(10, 11)"	127Hz - 133Hz	StimAmpl	0.5021	-1.0971
Right	"(10, 11)"	127Hz - 133Hz	StimAmpl	0.0903	0.9093
Left	"(1, 3)"	0.98Hz - 4.88Hz	HourOfDay	0.1912	0.0743
Left	"(1, 3)"	0.98Hz - 10.74Hz	HourOfDay	0.1815	0.0686
Left	"(1, 3)"	0.98Hz - 120Hz	HourOfDay	0.1757	0.0656
Left	"(1, 3)"	4.88Hz - 8.79Hz	HourOfDay	0.1503	0.0598
Left	"(1, 3)"	4.88Hz - 10.74Hz	HourOfDay	0.1428	0.0617
Left	"(1, 3)"	4.88Hz - 12.7Hz	HourOfDay	0.1470	0.0614
Left	"(1, 3)"	4.88Hz - 24.41Hz	HourOfDay	0.1515	0.0618
Left	"(1, 3)"	4.88Hz - 28.32Hz	HourOfDay	0.1513	0.0617

Left	"(1, 3)"	6.84Hz - 12.7Hz	HourOfDay	0.1528	0.0623
Left	"(1, 3)"	8.79Hz - 12.7Hz	HourOfDay	0.1687	0.0713
Left	"(1, 3)"	8.79Hz - 30.27Hz	HourOfDay	0.1578	0.0637
Left	"(1, 3)"	10.74Hz - 28.32Hz	HourOfDay	0.1589	0.0647
Left	"(1, 3)"	12.7Hz - 30.27Hz	HourOfDay	0.1513	0.0601
Left	"(1, 3)"	12.7Hz - 26.37Hz	HourOfDay	0.1545	0.0622
Right	"(1, 3)"	51.76Hz - 100.59Hz	HourOfDay	-0.1681	-0.0268
Right	"(1, 3)"	57.62Hz - 100.59Hz	HourOfDay	-0.1625	-0.0245
Right	"(1, 3)"	61.52Hz - 71.29Hz	HourOfDay	-0.1518	-0.0260
Left	"(1, 3)"	71.29Hz - 100.59Hz	HourOfDay	0.5499	0.2109
Right	"(1, 3)"	71.29Hz - 100.59Hz	HourOfDay	-0.1853	-0.0325
Left	"(1, 3)"	79.1Hz - 100.59Hz	HourOfDay	0.1641	0.0546
Right	"(1, 3)"	79.1Hz - 100.59Hz	HourOfDay	-0.1702	-0.0279
Left	"(1, 3)"	127Hz - 133Hz	HourOfDay	0.1857	0.0722
Left	"(8, 9)"	0.98Hz - 10.74Hz	HourOfDay	0.1438	0.0268
Left	"(8, 9)"	0.98Hz - 120Hz	HourOfDay	0.1680	0.0181
Left	"(8, 9)"	4.88Hz - 12.7Hz	HourOfDay	0.5586	0.0762
Left	"(8, 9)"	4.88Hz - 24.41Hz	HourOfDay	0.1448	0.0255
Left	"(8, 9)"	4.88Hz - 28.32Hz	HourOfDay	0.1180	0.0217
Left	"(8, 9)"	6.84Hz - 12.7Hz	HourOfDay	0.1616	0.0196
Left	"(8, 9)"	8.79Hz - 12.7Hz	HourOfDay	0.1949	0.0267
Left	"(8, 9)"	32.23Hz - 59.57Hz	HourOfDay	0.1612	0.0173
Left	"(10, 11)"	0.98Hz - 120Hz	HourOfDay	0.1420	0.0116
Left	"(10, 11)"	4.88Hz - 8.79Hz	HourOfDay	0.2073	0.0238
Left	"(10, 11)"	4.88Hz - 10.74Hz	HourOfDay	0.2241	0.0214
Left	"(10, 11)"	4.88Hz - 12.7Hz	HourOfDay	0.2049	0.0187
Left	"(10, 11)"	4.88Hz - 24.41Hz	HourOfDay	0.1309	0.0151
Left	"(10, 11)"	4.88Hz - 28.32Hz	HourOfDay	0.0927	0.0106
Left	"(10, 11)"	8.79Hz - 12.7Hz	HourOfDay	0.1959	0.0280
Left	"(10, 11)"	8.79Hz - 30.27Hz	HourOfDay	0.0442	0.0063
Left	"(10, 11)"	10.74Hz - 28.32Hz	HourOfDay	0.0362	0.0056
Left	"(10, 11)"	12.7Hz - 30.27Hz	HourOfDay	0.0171	0.0028
Left	"(10, 11)"	12.7Hz - 26.37Hz	HourOfDay	0.0213	0.0037
Left	"(10, 11)"	57.62Hz - 100.59Hz	HourOfDay	0.1572	0.0122
Left	"(10, 11)"	71.29Hz - 100.59Hz	HourOfDay	0.1521	0.0104
Left	"(10, 11)"	127Hz - 133Hz	HourOfDay	-0.2526	-0.0910
Left	"(1, 3)"	0.98Hz - 10.74Hz	isMonopolar	0.1205	0.8436
Left	"(1, 3)"	0.98Hz - 120Hz	isMonopolar	0.2241	1.0368
Left	"(1, 3)"	4.88Hz - 8.79Hz	isMonopolar	0.5388	3.6193
Left	"(1, 3)"	4.88Hz - 10.74Hz	isMonopolar	0.3123	3.2425
Left	"(1, 3)"	4.88Hz - 12.7Hz	isMonopolar	0.3374	1.4718
Left	"(1, 3)"	4.88Hz - 24.41Hz	isMonopolar	0.3373	1.3912

Right	"(1, 3)"	4.88Hz - 24.41Hz	isMonopolar	0.9835	6.2300
Left	"(1, 3)"	4.88Hz - 28.32Hz	isMonopolar	0.3332	1.3705
Right	"(1, 3)"	4.88Hz - 28.32Hz	isMonopolar	0.9837	6.2528
Left	"(1, 3)"	6.84Hz - 12.7Hz	isMonopolar	0.3794	0.9907
Left	"(1, 3)"	8.79Hz - 12.7Hz	isMonopolar	0.4198	1.1534
Left	"(1, 3)"	8.79Hz - 30.27Hz	isMonopolar	0.3503	0.9173
Right	"(1, 3)"	8.79Hz - 30.27Hz	isMonopolar	0.9803	6.2176
Left	"(1, 3)"	10.74Hz - 28.32Hz	isMonopolar	0.3456	0.9157
Right	"(1, 3)"	10.74Hz - 28.32Hz	isMonopolar	0.9810	6.3928
Left	"(1, 3)"	12.7Hz - 30.27Hz	isMonopolar	0.3221	0.8294
Right	"(1, 3)"	12.7Hz - 30.27Hz	isMonopolar	0.9844	7.3170
Left	"(1, 3)"	12.7Hz - 26.37Hz	isMonopolar	0.3212	0.8397
Right	"(1, 3)"	12.7Hz - 26.37Hz	isMonopolar	0.9845	7.4502
Right	"(1, 3)"	32.23Hz - 59.57Hz	isMonopolar	0.9780	6.1822
Left	"(1, 3)"	32.23Hz - 38.09Hz	isMonopolar	0.2037	0.4406
Right	"(1, 3)"	51.76Hz - 100.59Hz	isMonopolar	0.9808	4.5903
Right	"(1, 3)"	57.62Hz - 100.59Hz	isMonopolar	0.9811	4.3669
Right	"(1, 3)"	61.52Hz - 71.29Hz	isMonopolar	0.9794	4.7607
Left	"(1, 3)"	71.29Hz - 100.59Hz	isMonopolar	0.3609	0.7938
Left	"(1, 3)"	79.1Hz - 100.59Hz	isMonopolar	0.3493	0.7586
Right	"(1, 3)"	127Hz - 133Hz	isMonopolar	0.1620	-1.1518
Left	"(8, 9)"	0.98Hz - 120Hz	isMonopolar	0.2349	0.1879
Right	"(8, 9)"	4.88Hz - 8.79Hz	isMonopolar	0.2785	0.1682
Right	"(8, 9)"	4.88Hz - 10.74Hz	isMonopolar	0.3411	0.2343
Right	"(8, 9)"	4.88Hz - 12.7Hz	isMonopolar	0.3905	0.2674
Left	"(8, 9)"	32.23Hz - 59.57Hz	isMonopolar	0.3416	0.2746
Left	"(8, 9)"	32.23Hz - 38.09Hz	isMonopolar	0.2212	0.1891
Left	"(8, 9)"	71.29Hz - 100.59Hz	isMonopolar	0.2013	0.2225
Left	"(8, 9)"	73.24Hz - 77.14Hz	isMonopolar	0.2193	0.2473
Left	"(8, 9)"	79.1Hz - 100.59Hz	isMonopolar	0.1879	0.2081
Left	"(8, 9)"	127Hz - 133Hz	isMonopolar	0.8909	4.0607
Right	"(8, 9)"	127Hz - 133Hz	isMonopolar	0.8946	3.4262
Left	"(10, 11)"	0.98Hz - 4.88Hz	isMonopolar	-0.4170	-0.4515
Left	"(10, 11)"	0.98Hz - 10.74Hz	isMonopolar	-0.3397	-0.2239
Left	"(10, 11)"	6.84Hz - 12.7Hz	isMonopolar	-0.3365	-0.2364
Left	"(10, 11)"	30.27Hz - 38.09Hz	isMonopolar	-0.4273	-0.2339
Left	"(10, 11)"	32.23Hz - 59.57Hz	isMonopolar	-0.5855	-0.2750
Left	"(10, 11)"	32.23Hz - 38.09Hz	isMonopolar	-0.3915	-0.3084
Left	"(10, 11)"	51.76Hz - 100.59Hz	isMonopolar	-0.3245	-0.1373
Left	"(10, 11)"	57.62Hz - 100.59Hz	isMonopolar	-0.2667	-0.1239
Right	"(10, 11)"	79.1Hz - 100.59Hz	isMonopolar	0.5032	0.2295
Left	"(10, 11)"	127Hz - 133Hz	isMonopolar	-0.8035	-5.4306

Right	"(1, 3)"	0.98Hz - 10.74Hz	isMoving	0.1397	0.1537
Right	"(1, 3)"	32.23Hz - 59.57Hz	isMoving	-0.5317	-0.8005
Right	"(1, 3)"	32.23Hz - 38.09Hz	isMoving	-0.4491	-0.8577
Right	"(1, 3)"	51.76Hz - 100.59Hz	isMoving	-0.6182	-0.6928
Right	"(1, 3)"	57.62Hz - 100.59Hz	isMoving	-0.6044	-0.6316
Left	"(1, 3)"	61.52Hz - 71.29Hz	isMoving	0.2042	0.3851
Right	"(1, 3)"	61.52Hz - 71.29Hz	isMoving	-0.5113	-0.5646
Right	"(1, 3)"	73.24Hz - 77.14Hz	isMoving	-0.5153	-0.7018
Left	"(8, 9)"	0.98Hz - 4.88Hz	isMoving	0.4760	0.5541
Right	"(8, 9)"	0.98Hz - 4.88Hz	isMoving	0.3023	0.3238
Left	"(8, 9)"	0.98Hz - 10.74Hz	isMoving	0.4109	0.7060
Right	"(8, 9)"	0.98Hz - 10.74Hz	isMoving	0.2819	0.1757
Right	"(8, 9)"	4.88Hz - 8.79Hz	isMoving	0.1345	0.1095
Right	"(8, 9)"	4.88Hz - 24.41Hz	isMoving	0.1975	0.1022
Right	"(8, 9)"	4.88Hz - 28.32Hz	isMoving	-0.3015	-0.1589
Right	"(8, 9)"	6.84Hz - 12.7Hz	isMoving	-0.1816	-0.1353
Right	"(8, 9)"	8.79Hz - 12.7Hz	isMoving	-0.2461	-0.1845
Left	"(8, 9)"	8.79Hz - 30.27Hz	isMoving	-0.1555	-0.1456
Right	"(8, 9)"	8.79Hz - 30.27Hz	isMoving	-0.4278	-0.2593
Left	"(8, 9)"	10.74Hz - 28.32Hz	isMoving	-0.1606	-0.1564
Right	"(8, 9)"	10.74Hz - 28.32Hz	isMoving	-0.4092	-0.2535
Left	"(8, 9)"	12.7Hz - 30.27Hz	isMoving	-0.1867	-0.1934
Right	"(8, 9)"	12.7Hz - 30.27Hz	isMoving	-0.4311	-0.2887
Left	"(8, 9)"	12.7Hz - 26.37Hz	isMoving	-0.1896	-0.2026
Right	"(8, 9)"	12.7Hz - 26.37Hz	isMoving	-0.4254	-0.2932
Right	"(8, 9)"	30.27Hz - 38.09Hz	isMoving	-0.4865	-0.1151
Right	"(8, 9)"	32.23Hz - 38.09Hz	isMoving	-0.3161	-0.1541
Left	"(8, 9)"	51.76Hz - 100.59Hz	isMoving	0.5496	0.3959
Right	"(8, 9)"	51.76Hz - 100.59Hz	isMoving	0.7602	0.4126
Left	"(8, 9)"	57.62Hz - 100.59Hz	isMoving	0.5948	0.4842
Right	"(8, 9)"	57.62Hz - 100.59Hz	isMoving	0.7920	0.4564
Left	"(8, 9)"	61.52Hz - 71.29Hz	isMoving	0.5913	0.6354
Right	"(8, 9)"	61.52Hz - 71.29Hz	isMoving	0.6055	0.1800
Left	"(8, 9)"	71.29Hz - 100.59Hz	isMoving	0.4732	0.3686
Right	"(8, 9)"	71.29Hz - 100.59Hz	isMoving	0.7768	0.4718
Left	"(8, 9)"	73.24Hz - 77.14Hz	isMoving	0.4140	0.1666
Right	"(8, 9)"	73.24Hz - 77.14Hz	isMoving	0.6248	0.5430
Left	"(8, 9)"	79.1Hz - 100.59Hz	isMoving	0.4814	0.3727
Right	"(8, 9)"	79.1Hz - 100.59Hz	isMoving	0.7432	0.4646
Left	"(10, 11)"	0.98Hz - 4.88Hz	isMoving	0.1803	0.1726
Right	"(10, 11)"	0.98Hz - 4.88Hz	isMoving	0.2468	0.2147
Left	"(10, 11)"	0.98Hz - 120Hz	isMoving	-0.7556	-0.5311

Right	"(10, 11)"	0.98Hz - 120Hz	isMoving	-0.3197	-0.1947
Right	"(10, 11)"	4.88Hz - 8.79Hz	isMoving	-0.2386	-0.1746
Right	"(10, 11)"	4.88Hz - 10.74Hz	isMoving	-0.2425	-0.1667
Left	"(10, 11)"	4.88Hz - 12.7Hz	isMoving	0.5498	0.3378
Right	"(10, 11)"	4.88Hz - 12.7Hz	isMoving	-0.3113	-0.1868
Left	"(10, 11)"	4.88Hz - 24.41Hz	isMoving	-0.7857	-0.8078
Right	"(10, 11)"	4.88Hz - 24.41Hz	isMoving	-0.3897	-0.3340
Left	"(10, 11)"	4.88Hz - 28.32Hz	isMoving	-0.8290	-0.9430
Right	"(10, 11)"	4.88Hz - 28.32Hz	isMoving	-0.4556	-0.3879
Left	"(10, 11)"	6.84Hz - 12.7Hz	isMoving	0.3465	0.2339
Right	"(10, 11)"	6.84Hz - 12.7Hz	isMoving	-0.2492	-0.1887
Left	"(10, 11)"	8.79Hz - 12.7Hz	isMoving	0.3613	0.2891
Right	"(10, 11)"	8.79Hz - 12.7Hz	isMoving	-0.1931	-0.1604
Left	"(10, 11)"	8.79Hz - 30.27Hz	isMoving	-0.8318	-1.1982
Right	"(10, 11)"	8.79Hz - 30.27Hz	isMoving	-0.4001	-0.4163
Left	"(10, 11)"	10.74Hz - 28.32Hz	isMoving	-0.8355	-1.3038
Right	"(10, 11)"	10.74Hz - 28.32Hz	isMoving	-0.4120	-0.4482
Left	"(10, 11)"	12.7Hz - 30.27Hz	isMoving	-0.8342	-1.3689
Right	"(10, 11)"	12.7Hz - 30.27Hz	isMoving	-0.4217	-0.4737
Left	"(10, 11)"	12.7Hz - 26.37Hz	isMoving	-0.8355	-1.4591
Right	"(10, 11)"	12.7Hz - 26.37Hz	isMoving	-0.4140	-0.4924
Right	"(10, 11)"	30.27Hz - 38.09Hz	isMoving	-0.1606	-0.0929
Left	"(10, 11)"	51.76Hz - 100.59Hz	isMoving	0.4641	0.2006
Right	"(10, 11)"	51.76Hz - 100.59Hz	isMoving	0.4306	0.0591
Left	"(10, 11)"	57.62Hz - 100.59Hz	isMoving	0.4888	0.2401
Right	"(10, 11)"	57.62Hz - 100.59Hz	isMoving	0.5793	0.2506
Left	"(10, 11)"	61.52Hz - 71.29Hz	isMoving	0.4089	0.2912
Right	"(10, 11)"	61.52Hz - 71.29Hz	isMoving	0.3488	0.2132
Left	"(10, 11)"	71.29Hz - 100.59Hz	isMoving	0.4865	0.2904
Right	"(10, 11)"	71.29Hz - 100.59Hz	isMoving	0.6327	0.2840
Left	"(10, 11)"	73.24Hz - 77.14Hz	isMoving	0.2677	0.1946
Right	"(10, 11)"	73.24Hz - 77.14Hz	isMoving	0.3893	0.2812
Left	"(10, 11)"	79.1Hz - 100.59Hz	isMoving	0.4576	0.1406
Right	"(10, 11)"	79.1Hz - 100.59Hz	isMoving	0.5771	0.2692
Left	"(1, 3)"	0.98Hz - 4.88Hz	isFullDay	-0.1003	-0.4512
Left	"(1, 3)"	4.88Hz - 8.79Hz	isFullDay	-0.1481	-0.4520
Right	"(1, 3)"	71.29Hz - 100.59Hz	isFullDay	0.2709	0.3059
Right	"(1, 3)"	79.1Hz - 100.59Hz	isFullDay	0.2867	0.3214
Right	"(8, 9)"	0.98Hz - 120Hz	isFullDay	0.1537	0.0673
Left	"(8, 9)"	6.84Hz - 12.7Hz	isFullDay	0.2425	0.2086
Right	"(8, 9)"	8.79Hz - 12.7Hz	isFullDay	0.2195	0.1900
Right	"(8, 9)"	30.27Hz - 38.09Hz	isFullDay	0.1780	0.1186

Left	"(8, 9)"	51.76Hz - 100.59Hz	isFullDay	0.3119	0.2111
Left	"(8, 9)"	57.62Hz - 100.59Hz	isFullDay	0.2921	0.2071
Left	"(8, 9)"	61.52Hz - 71.29Hz	isFullDay	0.2680	0.2756
Right	"(10, 11)"	0.98Hz - 10.74Hz	isFullDay	0.2469	0.1737
Right	"(10, 11)"	4.88Hz - 8.79Hz	isFullDay	0.2990	0.2698
Right	"(10, 11)"	4.88Hz - 10.74Hz	isFullDay	0.3152	0.2684
Right	"(10, 11)"	4.88Hz - 12.7Hz	isFullDay	0.2704	0.1941
Right	"(10, 11)"	6.84Hz - 12.7Hz	isFullDay	0.1995	0.1810
Left	"(10, 11)"	32.23Hz - 59.57Hz	isFullDay	0.2924	0.1393
Right	"(1, 3)"	127Hz - 133Hz	MonthsElapsed:StimAmpl	-0.1052	0.1151
Right	"(8, 9)"	57.62Hz - 100.59Hz	MonthsElapsed:StimAmpl	-0.2902	-0.0032
Right	"(8, 9)"	71.29Hz - 100.59Hz	MonthsElapsed:StimAmpl	-0.3079	-0.0041
Left	"(8, 9)"	127Hz - 133Hz	MonthsElapsed:StimAmpl	-0.9554	-0.2251
Right	"(10, 11)"	57.62Hz - 100.59Hz	MonthsElapsed:StimAmpl	-0.3128	-0.0100
Right	"(10, 11)"	71.29Hz - 100.59Hz	MonthsElapsed:StimAmpl	-0.7248	-0.0218
Right	"(10, 11)"	79.1Hz - 100.59Hz	MonthsElapsed:StimAmpl	-0.9100	-0.0275
Left	"(10, 11)"	127Hz - 133Hz	MonthsElapsed:StimAmpl	-0.0983	0.1402
Right	"(10, 11)"	127Hz - 133Hz	MonthsElapsed:StimAmpl	0.8197	0.0557
Left	"(1, 3)"	71.29Hz - 100.59Hz	HourOfDay:StimAmpl	-0.6765	-0.0695
Left	"(8, 9)"	4.88Hz - 12.7Hz	HourOfDay:StimAmpl	-0.6794	-0.0250
Right	"(8, 9)"	6.84Hz - 12.7Hz	HourOfDay:StimAmpl	-0.1707	-0.0052
Right	"(8, 9)"	51.76Hz - 100.59Hz	HourOfDay:StimAmpl	-0.2773	-0.0042
Right	"(10, 11)"	8.79Hz - 12.7Hz	HourOfDay:StimAmpl	-0.1448	-0.0049
Right	"(10, 11)"	127Hz - 133Hz	HourOfDay:StimAmpl	-0.2929	-0.0704
Right	"(1, 3)"	0.98Hz - 10.74Hz	isMonopolar:StimAmpl	-0.0596	-2.5453
Right	"(1, 3)"	0.98Hz - 120Hz	isMonopolar:StimAmpl	0.0342	-2.9353
Right	"(1, 3)"	4.88Hz - 10.74Hz	isMonopolar:StimAmpl	0.0268	-3.1749
Right	"(1, 3)"	4.88Hz - 24.41Hz	isMonopolar:StimAmpl	-0.9821	-3.2300
Right	"(1, 3)"	4.88Hz - 28.32Hz	isMonopolar:StimAmpl	-0.9823	-3.2408
Right	"(1, 3)"	6.84Hz - 12.7Hz	isMonopolar:StimAmpl	0.0810	-2.6787
Right	"(1, 3)"	8.79Hz - 30.27Hz	isMonopolar:StimAmpl	-0.9784	-3.2001
Right	"(1, 3)"	10.74Hz - 28.32Hz	isMonopolar:StimAmpl	-0.9790	-3.2828
Right	"(1, 3)"	12.7Hz - 30.27Hz	isMonopolar:StimAmpl	-0.9829	-3.7671
Right	"(1, 3)"	12.7Hz - 26.37Hz	isMonopolar:StimAmpl	-0.9830	-3.8348
Left	"(1, 3)"	32.23Hz - 59.57Hz	isMonopolar:StimAmpl	0.2290	0.2647
Right	"(1, 3)"	32.23Hz - 59.57Hz	isMonopolar:StimAmpl	-0.9757	-3.1765
Right	"(1, 3)"	51.76Hz - 100.59Hz	isMonopolar:StimAmpl	-0.9786	-2.3394
Right	"(1, 3)"	57.62Hz - 100.59Hz	isMonopolar:StimAmpl	-0.9786	-2.2129
Right	"(1, 3)"	61.52Hz - 71.29Hz	isMonopolar:StimAmpl	-0.9762	-2.3889
Right	"(1, 3)"	73.24Hz - 77.14Hz	isMonopolar:StimAmpl	-0.0279	-2.2979
Left	"(1, 3)"	127Hz - 133Hz	isMonopolar:StimAmpl	0.4783	0.6048
Right	"(8, 9)"	32.23Hz - 38.09Hz	isMonopolar:StimAmpl	0.3417	0.1376

Right	"(8, 9)"	127Hz - 133Hz	isMonopolar:StimAmpl	-0.8127	-1.2894
Right	"(10, 11)"	4.88Hz - 12.7Hz	isMonopolar:StimAmpl	0.2804	0.0931
Left	"(10, 11)"	127Hz - 133Hz	isMonopolar:StimAmpl	0.8116	2.2594
Right	"(1, 3)"	71.29Hz - 100.59Hz	isMoving:MonthsElapsed	-0.3846	-0.0345
Right	"(1, 3)"	79.1Hz - 100.59Hz	isMoving:MonthsElapsed	-0.3886	-0.0381
Right	"(8, 9)"	0.98Hz - 120Hz	isMoving:MonthsElapsed	0.3624	0.0128
Right	"(8, 9)"	32.23Hz - 59.57Hz	isMoving:MonthsElapsed	0.4571	0.0141
Right	"(1, 3)"	0.98Hz - 120Hz	isMoving:StimAmpl	0.1187	0.0593
Right	"(1, 3)"	30.27Hz - 38.09Hz	isMoving:StimAmpl	0.1868	0.3788
Right	"(1, 3)"	32.23Hz - 59.57Hz	isMoving:StimAmpl	0.6342	0.4278
Right	"(1, 3)"	32.23Hz - 38.09Hz	isMoving:StimAmpl	0.5732	0.4754
Right	"(1, 3)"	51.76Hz - 100.59Hz	isMoving:StimAmpl	0.7230	0.3770
Right	"(1, 3)"	57.62Hz - 100.59Hz	isMoving:StimAmpl	0.7220	0.3553
Right	"(1, 3)"	61.52Hz - 71.29Hz	isMoving:StimAmpl	0.6948	0.3749
Right	"(1, 3)"	71.29Hz - 100.59Hz	isMoving:StimAmpl	0.4864	0.3866
Right	"(1, 3)"	73.24Hz - 77.14Hz	isMoving:StimAmpl	0.6438	0.3977
Right	"(1, 3)"	79.1Hz - 100.59Hz	isMoving:StimAmpl	0.4808	0.2094
Right	"(8, 9)"	0.98Hz - 120Hz	isMoving:StimAmpl	-0.4101	-0.0657
Right	"(8, 9)"	4.88Hz - 24.41Hz	isMoving:StimAmpl	-0.4348	-0.1002
Right	"(8, 9)"	32.23Hz - 59.57Hz	isMoving:StimAmpl	-0.4428	-0.0617
Right	"(8, 9)"	61.52Hz - 71.29Hz	isMoving:StimAmpl	0.1914	0.1089
Left	"(10, 11)"	4.88Hz - 8.79Hz	isMoving:StimAmpl	-0.1849	-0.0527
Left	"(10, 11)"	4.88Hz - 10.74Hz	isMoving:StimAmpl	-0.1987	-0.0467
Left	"(10, 11)"	4.88Hz - 12.7Hz	isMoving:StimAmpl	-0.6664	-0.2044
Left	"(10, 11)"	6.84Hz - 12.7Hz	isMoving:StimAmpl	-0.5157	-0.1707
Left	"(10, 11)"	8.79Hz - 12.7Hz	isMoving:StimAmpl	-0.5616	-0.2332
Left	"(10, 11)"	32.23Hz - 59.57Hz	isMoving:StimAmpl	-0.1739	-0.0288
Right	"(10, 11)"	51.76Hz - 100.59Hz	isMoving:StimAmpl	0.1631	0.0708
Left	"(10, 11)"	0.98Hz - 120Hz	isMoving:HourOfDay	0.6843	0.0307
Left	"(10, 11)"	4.88Hz - 24.41Hz	isMoving:HourOfDay	0.7158	0.0463
Left	"(10, 11)"	4.88Hz - 28.32Hz	isMoving:HourOfDay	0.7682	0.0542
Left	"(10, 11)"	8.79Hz - 30.27Hz	isMoving:HourOfDay	0.7657	0.0676
Left	"(10, 11)"	10.74Hz - 28.32Hz	isMoving:HourOfDay	0.7713	0.0738
Left	"(10, 11)"	12.7Hz - 30.27Hz	isMoving:HourOfDay	0.7689	0.0773
Left	"(10, 11)"	12.7Hz - 26.37Hz	isMoving:HourOfDay	0.7772	0.0842