OPTIMISING NEUROTHERAPY

Bayesian Optimisation of Deep Brain Stimulation

Oscar Lemmens Supervisor: Prof. Tim Denison

UNIVERSITY OF OXFORD

MRC Brain Network Dynamics Unit

DEPARTMENT OF ENGINEERING SCIENCE

Motivation & Objectives

Deep brain stimulation (DBS) is an increasingly adopted form of neurotherapy for brain disorders, a subcategory of neurological disorders. DBS involves the application of electrical neuromodulation, and works by altering neuronal activity by electrical impulses sent to targeted regions in the brain nuclei to alleviate symptoms.

Although DBS is an established method of treatment, some significant limitations remain, notably stim parameter selection. Stim parameters consist of stimulation variables such as stim amplitude and frequency. Currently, there is little understanding behind the effect of different stim parameters due to the poorly defined transfer functions between parameters and their effects. As a result, parameter optimisation involves a clinician manually setting the parameters post-implantation and then infrequently updating them. This is done through a laborious trial-and-error random searching of the parameter space which often leads to suboptimal results and side effects.

In order to improve treatment, closed-loop adaptive DBS (aDBS) has been developed where electrical impulses are delivered only when necessary, by using a feedback biosignal which is correlated with the patient's brain state. The recent development of aDBS devices and use of feedback provides a basis to support a continuous optimisation routine. DBS could greatly benefit from an adaptive learning optimisation algorithm which bypasses the current procedure and many of its associated limitations, ultimately improving treatment. This project focuses on the development of an automatic algorithm that optimises stim parameters according to their effects on the observed response. This allows for rapid optimal parameter choice by efficiently sampling the parameter space.

To be clinically viable, the algorithm should meet several requirements:

1. Generalisable to any disease and biomarker, allowing it to be expanded to a number of neurotherapies.
2. Hardware agnostic and universally compatible with any DBS or bioelectronic stimulator.
3. Reliably finds the true global extrema of an objective function even if it is a time-varying function.
4. Adjustable to the patient's needs for tailored treatment, such as the explorative or exploitative balance.

Brain Disorders & Deep Brain Stimulation

Many neurological disorders are degenerative and require a strong intervention and persistent treatment, which in turn usually imposes a significant monetary and social burden for both the patients and the healthcare system. Neurological disorders are at the forefront of research as major progress is needed to provide more effective treatment. Furthermore, the treatment for brain disorders are still at a very primitive stage and end up being costly due to the complexity of the problem at hand. In 2010, in Europe alone, brain disorders made up a total healthcare cost of €800 billion and directly impacted almost a quarter of Europeans [1].

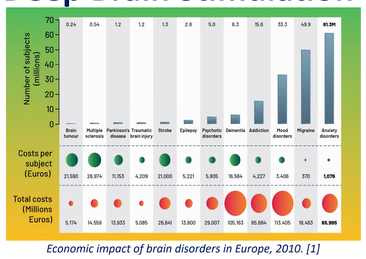

Economic impact of brain disorders in Europe, 2010. [1]

Many brain disorders can be 'treated' using DBS, which consists of a neurostimulator that produces electrical impulses and delivers them to specified areas of the brain using electrodes. However, there is still a lack of fundamental understanding of the mechanism of action of DBS and the effects of the different stim parameters. Stim parameters can be varied to control to what extent neuronal elements surrounding the electrode are recruited and activated. However, it is not entirely evident which regions of the surrounding area are being stimulated or blocked. As a result, unreliable past empirical and theoretical results are the only guidance clinicians have, leading to an ad hoc time-consuming task. There is also no standardised universal parameter setting for implanted patients, as each and every patient needs to have a unique set of parameters. Overall, there are many drawbacks to this empirical method of parameter selection, such as suboptimal settings and side effects.

[1] DiLuca, M., & Olesen, J. (2014). The cost of brain diseases: a burden or a challenge?. *Neuron*, *82*(6), 1205-1208.

Dynamics & Periodicity of Brain Disorders

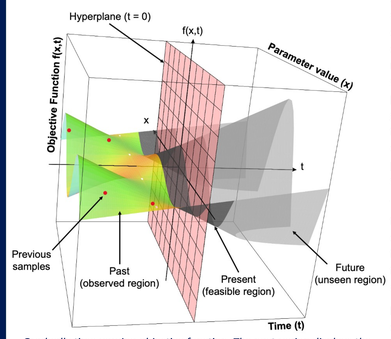

Gradually time-varying objective function. The past region displays the previously acquired samples, the darker central region and hyperplane illustrates the current feasible objective function, finally the lightly shaded region indicates the unobserved future region.

Objective functions are formed using biomarkers which often change over time as they are extracted from a dynamic neural network environment. This gives rise to a dynamic optimisation problem where the function has a spatial component, as well as a time component. In practice, these dynamics are due to gradual changes such as neuroplastic network reconfigurations and tissue-electrode interface variations, or sudden changes such as electrode movement and new oscillation emergence. This implies that stim parameter effects might not remain constant and that previous data points including the globally optimal parameter might not be fitting anymore. This poses a different question to that of a time-invariant problem where information is gathered from all the evaluations made in order to guide the search. Instead, in a dynamic optimisation problem the algorithm has two goals, firstly to find the minimum of the function, and secondly to track its development through time.

To deal with gradually varying dynamic functions, a time-varying GP is implemented to track the objective function's minimum through the solution space and ensure continuous optimal therapeutic delivery. This is done using a spatiotemporal kernel following a simple Markov model with a forgetting factor which places a weighting on the relevance of each sample based on the time they were acquired. As a result, older stale data is deemed irrelevant and discarded, thus leading to a smooth forgetting effect.

For rapid changes, a sudden change detection scheme is introduced, a statistical process control (SPC) $\bar{x}$ and R chart. SPC is particularly advantageous as it focuses on early change detection and problem prevention by providing a robust framework to distinguish between expected common cause variations (random, natural and intrinsic variation such as sensor noise) and special cause variations (non-random variation stemming from external changes). If special cause variations are detected (moving range or process average falls outside of the required bounds), the process is deemed to be out of statistical control and the algorithm's re-exploration function is triggered.

Biorhythms such as the circadian rhythm also have a significant effect on the stim response of a patient as it means that throughout the day, the objective function and consequently the optimal stimulation paradigm vary with a 24 hour time period. The quality of neurotherapy could be greatly enhanced by upgrading the algorithm to support the effects of chronobiology. By including more sophisticated spatiotemporal GP kernels, more complex dynamics such as biorhythms can be captured. This extra information in the prior would enable the device to track physiological time-based cyclic behaviour and improve parameter selection suitability by having multiple optimal control policies throughout the day. This is done through the addition of a periodic temporal GP kernel, where sample weighting is no longer just based on the age of the sample but also on what its phase in the cycle was when obtained. In other words, samples that are collected at similar points in the cycle will have greater impact on the current prediction, whereas samples collected at vastly different stages in the cycle are considered unrepresentative of the system.

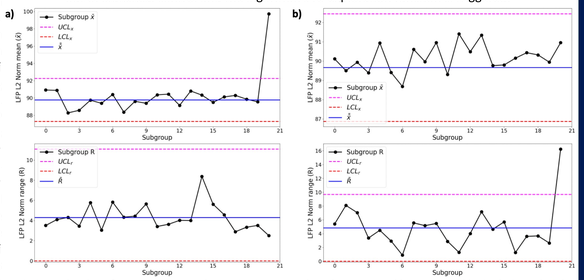

*SPC $\bar{x}$ and R chart for BayesOpt sample monitoring examples. The upper and lower control limits (UCL and LCL) provide a stability threshold of the system. **a)** The average rises above the $UCL_{\bar{x}}$ for the 20th subgroup indicating a significant change in the process and sub-par symptom suppression. The R chart does not indicate any irregular changes in statistical dispersion. **b)** Sample noise is greatly increased for subgroup 20. The subgroup's mean value is within bounds and appears to be reasonable, however the sample range is considerably greater than expected.*

Clinical Implementation

The aim of the optimisation algorithm is to find and continuously select the optimal stim parameters in as few trials as possible. With the increasing adoption of aDBS devices, optimisation routines are made feasible through disease-specific biosignal feedback which can be used to specify a biomarker that correlates with symptom severity. An objective function can then be formed using this biomarker and be evaluated by using noisy sensor readings. The argument of the objective function's minimum corresponds to the desired optimal parameter with greatest symptom suppression. However, in order to promptly provide effective symptom suppression it is important to minimise the number of samples made when finding the function's minimum. BayesOpt overcomes the challenge of hunting through the vast parameter space when faced with constraints on the search space. This removes the largely heuristic biannual process of manual parameter reprogramming which in turn lessens the burden and associated patient and clinical costs of regularly needing to go to a clinic for parameter adjustment.

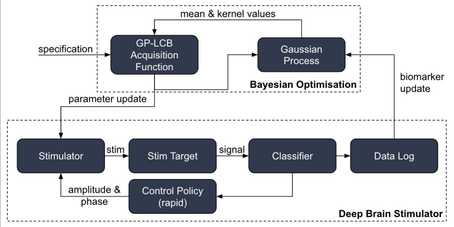

Bayesian optimisation and stimulator control flow. BayesOpt loop is the externally run algorithm which explores and selects stim parameters. The closed-loop stimulator component can function as a standalone aDBS device.

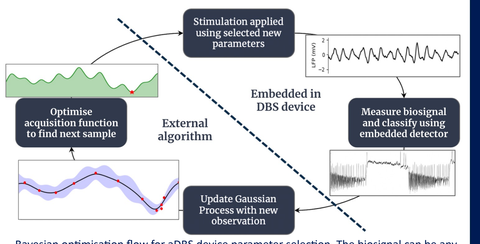

Bayesian optimisation flow for aDBS device parameter selection. The biosignal can be any signal that correlates well with the patient's state, such as specific bandwidth power from LFP, EEG, ECoG signals from sensor electrodes.

The BayesOpt algorithm proposed is implemented in Python using *scikit-learn* and *skopt* packages and can be deployed in conjunction with any aDBS device. However, a supporting architecture is needed in order to facilitate BayesOpt in a clinical setting, in particular, two crucial components are required to deploy it. First, an API is needed to provide remote connectivity and enable data relay between the DBS device and external computer, thus allowing the user to send and receive objective data. On the transmitting end, optimised stim parameters are pushed remotely to the DBS device. The resulting sensor and classification data are then returned wirelessly to the clinician's PC. This remote access interfaces the implanted device with the code and allows the user to test experimental routines by facilitating flexible algorithmic prototyping. Second, a classifier is needed to support the algorithm by mapping biomarker readings into state categories and keep track of pathological events. These classes can then be used as an actuation trigger in the control policy and determine which real-time stim setting should be selected. Additionally, the classifier and data log gather a virtual diary of states and episodes of a patient to aid in the remote monitoring of the patient's progress. These components comprise the fundamental blocks needed to build a tech stack for successful practical translation.

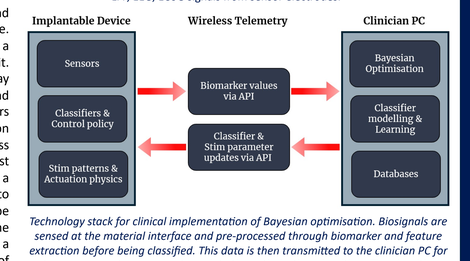

Technology stack for clinical implementation of Bayesian optimisation. Biosignals are sensed at the material interface and pre-processed through biomarker and feature extraction before being classified. This data is then transmitted to the clinician PC for storage and post-processing to update the optimization algorithm. The algorithm then telemeters optimized feedback actuation parameters to the implantable device.

Bayesian Optimisation

Bayesian optimisation (BayesOpt) is a global optimisation strategy and is particularly useful for expensive-to-evaluate black box functions with no functional form known a priori. BayesOpt efficiently finds extrema by minimising a given objective function through a systematic probabilistic search utilising Bayes' theorem. BayesOpt maintains a posterior distribution of the objective function as samples are made, and thus acts as a running learning algorithm. BayesOpt is an iterative process described as follows:

1. Estimate the true objective function by building a probabilistic surrogate model using a Gaussian process (GP).
2. Determine where to sample next by optimising an acquisition function.
3. Sample this point and update the surrogate model.
4. Repeat steps until the global minimum is found or resources are exhausted.

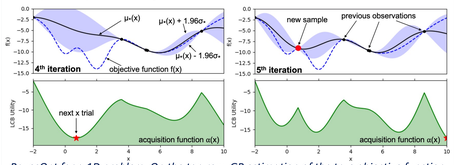

BayesOpt for a 1D problem. On the top row, GP estimation of the true objective function f(x) (blue dashed line). The purple shaded region indicates the posterior uncertainty of the surrogate model (95% confidence region) and the black line depicts the posterior mean, both determined through GP Regression. On the bottom row, the acquisition function determines which x value to sample next for best utility depicted by a red star.

This project proposes a sequential parameter optimisation algorithm for finding the extrema of objective functions using BayesOpt. By defining appropriate objective functions, BayesOpt can be implemented to find and select optimal parameters for maximum therapeutic benefit. This method is exceptionally effective when optimising costly black-box functions with unknown functional form, as is the case for most brain disorders. BayesOpt is chosen as it is effective and meets all the requirements set out: it is computationally efficient; it does not require expressions for the objective function or its derivatives; it reliably converges to the global minimum as opposed to local minima, even for non-convex and non-smooth functions.

Parkinson's Disease Study

Parkinson's disease (PD) is a common DBS application, PD symptoms manifest in a number of ways, including tremor, bradykinesia, postural instability and rigidity. In study, the STN is chosen as the stim target and LFP β-band (8-35 Hz) power measured from the GPi is used as the biomarker with the goal of minimisation, as LFP β power is correlated with symptom severity. A computational model of a basal ganglia-thalamocortical system (BGTCS) made up of nine neuronal populations is used in this study. As this is a mean-field model where neuronal populations are spatially averaged, values extracted from the populations can be interpreted as LFP signals and strongly resemble the measurements that can be sensed in practice. As for the stimulation policy, a combined β power and phase based strategy is used where the power and phase are extracted from the LFP measurement. In this approach, stimulation is applied when the β power is above a set threshold, and is delivered in the form of phase-locked pulses when the phase of the beta oscillation crosses a specified point. In this study, three parameters are optimised: oscillation phase trigger, power threshold and stim amplitude. These free variables are optimally found in an efficient automated manner using BayesOpt.

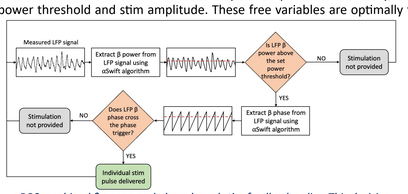

aDBS combined β power and phase based stim feedback policy. This decision tree diagram illustrates the key steps of the control process: the LFP signal is measured using an electrode sensor from which β power and phase are calculated using an aSwift algorithm. Stimulation is then only provided if both the power threshold and phase trigger (dashed red lines) criteria are met.

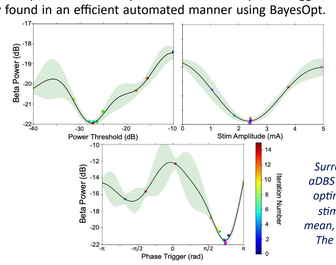

BGTCS model configuration. Black arrows represent excitatory connections, red arrows inhibitory ones. Stim is applied to the STN and the resulting LFP is measured from the GPi. Adapted from [2].

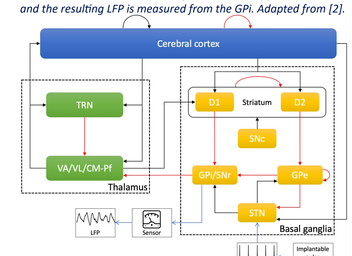

Surrogate model estimates of the biomarker value as a function of aDBS parameters using 1D BayesOpt. Individual Bayesian parameter optimisation for θ oscillation phase trigger, θ power threshold, and stimulation amplitude. The black line represents the GP posterior mean, the green shaded region illustrates the 95% confidence interval. The points represent the 15 total samples made, where the colour indicates the iteration number of the sample.

These parameters are optimised one at a time whilst the other two parameters are kept constant at their 'optimal' values. As illustrated, the algorithm is efficient at finding the minimum of the objective function in very few iterations. The algorithm starts off exploring the space before honing in and exploiting the minimum found. The samples do not necessarily lay on the surrogate function as noise is included during sample procurement to ensure noise robustness of the routine, as samples will be inherently noisy when implemented clinically. Even when testing the algorithm with large amounts of noise, the optimiser is able to factor this in and consistently find the minimum point of the function.

Now that a robust optimisation program has been created, one might think how beneficial BayesOpt actually is in comparison to simpler methods. Four different search strategies are compared in terms of convergence and regret ($f(x) - f_{min}$) which are both good quantitative indicators of optimiser performance. The other three searches tested are: random search; grid search, which generates a uniform grid from which points are tested; and Halton sequence, which is a low-discrepancy point generation method. BayesOpt vastly outperforms the other methods as it has the smallest total regret as well as the fastest convergence. This means optimal therapy can be derived and delivered faster leading to better treatment and an overall more productive parameter selection routine.

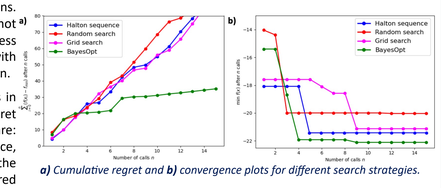

***a)** Cumulative regret and **b)** convergence plots for different search strategies.*

[2] van Albada, S. J., & Robinson, P. A. (2009). Mean-field modeling of the basal ganglia-thalamocortical system. Journal of theoretical biology, 257(4), 642-663.

Kuramoto Model Study

Objective functions tend not to be static as the brain is a dynamic system and the tissue response to stim parameters can vary over time due to inherent changes such as neuroplasticity. The complication of said time-varying objective functions is tackled using the Kuramoto model, a mathematical model used to describe systems that self organise and exhibit collective synchronisation. In particular, it can be used to examine rhythmic activity in large oscillator populations and describe the self organising behaviour behind sets of coupled oscillators, including the feedback loop between inhibitory and excitatory neurotransmitters in large neuronal networks. Excitatory oscillators want to conform and stay in phase, whereas inhibitory oscillators want to remain out of phase. This concept of inhibition and excitation often comes up when trying to understand the driving force behind oscillations present in brain disorders, such as the dominating β oscillation in PD. In healthy patients, STN neurons spike in an uncorrelated and desynchronised way whereas in PD, STN neurons form clusters of periodic and synchronous behaviour leading to tremors.

BayesOpt is deployed around this model using an objective function that is correlated to oscillator synchrony with gradually changing dynamics. aDBS is used again in this example, however only phase trigger is used and optimised. The numerical performance is compared for a optimisation strategies. It is apparent that the time-varying GP greatly outperforms a standard time-invariant GP as it is able to discard misleading old data. The time-invariant GP starts off well with a low regret in the early iterations, however as the dynamics change the model performance diminishes rapidly due to the stale data still contributing to the prediction. As time progresses, the time-invariant GP's performance is only slightly better than a random sampling strategy due to the large amount of unrepresentative data present in the model. These results confirm that a time-varying kernel is effective at dealing with dynamic systems with performance vastly exceeding standard techniques.

One key assumption in the time-varying model is that the objective function varies at a constant rate according to a simple Markov model. In practice, objective functions often follows underlying periodic rhythms such as the 24 hour circadian rhythm. To account for this, a periodic temporal GP is included in the model. To verify the utility of this solution, the system dynamics are periodically varied with a time period of 24 h to represent the effect of a circadian rhythm commonly found in brain disorders. The optimal parameter is evaluated at 600 time increments and plotted against time over 2 cycles. When comparing the algorithm's performance to other BayesOpt models, the periodic-forgetting strategy proves to be more effective in tracking the function's minimum, as highlighted when comparing the average regret of the routine using the three different methods.

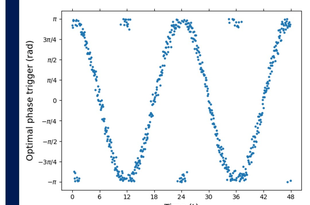

Optimal phase trigger determined for 600 BayesOpt iterations with periodic-forgetting spatiotemporal kernel.

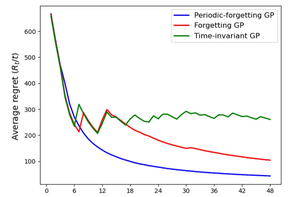

Average regret (R_t/t) of BayesOpt periodic search: periodic-forgetting, forgetting, and time-invariant GPs.

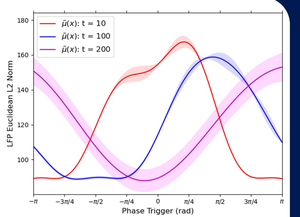

Time evolution of dynamic surrogate function. The surrogate function mean values and 95% confidence bounds are shown at times t = 10, 100 and 200.

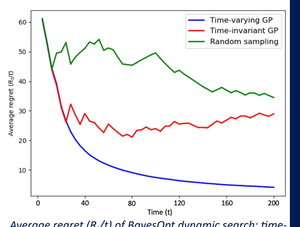

Average regret (R_t/t) of BayesOpt dynamic search: time-varying GP, time-invariant GP, random sampling.

Conclusion

The presented algorithm provides a flexible clinical framework for stimulation parameter optimisation and has the capability of continuously and autonomously calibrating DBS devices for optimal therapeutic delivery. Furthermore, the algorithm is able to determine and track the optimal DBS control policy allowing for effective treatment of dynamic and periodic brain disorders. By using continuous biofeedback, the control loop is never broken thus ensuring optimal symptom suppression and therapy. This not only greatly improves the device longevity and performance but also has the potential to drastically improve the quality of life of a patient. The algorithm's hyperparameters are all tunable and optimisable to provide maximum symptom suppression. The objectives initially laid out have all been met and in most cases exceeded through added functionality and levels of robustness.

lemmens.oscar@gmail.com
www.linkedin.com/in/oscarlemmens/

Acknowledgments: Prof. Tim Denison, Prof. Havriye Cagnan, Kei Landin, Prof. Tay Netoff, James Sermon