



(51) International Patent Classification:

A61N 1/36 (2006.01) A61N 1/05 (2006.01)

(21) International Application Number:

PCT/GB2024/051935

(22) International Filing Date:

23 July 2024 (23.07.2024)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

2311440.8 26 July 2023 (26.07.2023) GB

(71) Applicant: **OXFORD UNIVERSITY INNOVATION LIMITED** [GB/GB]; Buxton Court, 3 West Way, Oxford OX2 0JB (GB).

(72) Inventors: **CAGNAN, Hayriye**; Nuffield Department of Clinical Neurosciences, Level 6, West Wing, John Radcliffe Hospital, Oxford Oxfordshire OX3 9DU (GB). **DENISON,**

**Timothy**; Department of Engineering Science, University of Oxford, Parks Road, Oxford Oxfordshire OX1 3PJ (GB).


(74) Agent: **J A KEMP LLP**; 80 Turnmill Street, London Greater London EC1M 5QU (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, MG, MK, MN, MU, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, CV, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, SD, SL, ST,

(54) Title: SYSTEMS AND METHODS FOR GENERATING A STIMULATION SIGNAL FOR BRAIN OR NERVE STIMULATION AND FOR PERFORMING BRAIN OR NERVE STIMULATION

Fig. 1



(57) Abstract: Systems and methods for generating stimulation signals for brain or nerve stimulation are disclosed. In one arrangement, a signal generation unit generates a stimulation signal and transmits the stimulation signal to a stimulation unit. The stimulation unit applies the stimulation to a biological system. The stimulation signal is defined by at least one stimulation parameter. A data receiving unit receives time series biomarker data comprising measurement samples representing one or more biomarkers indicative of a condition of the biological system that is affected by applied stimulation based on the stimulation signal. A controller controls generation of the stimulation signal based on the received biomarker data. The controller uses a model of the biological system to estimate an optimal value of the at least one stimulation parameter using the biomarker data. The model defines strengths of contribution of measurement samples to the estimate based on the times at which the measurement samples were obtained.

SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, ME, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

— *with international search report (Art. 21(3))*

SYSTEMS AND METHODS FOR GENERATING A STIMULATION SIGNAL FOR  
BRAIN OR NERVE STIMULATION AND FOR PERFORMING BRAIN OR  
NERVE STIMULATION

5           The present disclosure relates to generating stimulation signals for brain or nerve stimulation.

          Over the last four decades neuromodulation has been adopted as an adjunct therapy for a variety of treatment resistant neurological and psychiatric disorders, whereby electrical or magnetic stimulation of certain brain regions is used to reduce disease  
10       symptoms. Invasive and non-invasive stimulation-based therapies (e.g., Deep Brain Stimulation (DBS), Transcranial Magnetic Stimulation (TMS), etc.) have so far relied on clinical expertise to empirically select patient specific stimulation parameters during clinical appointments. Due to the potentially very high dimensionality of stimulation parameterization, manual identification of effective therapy parameters is a difficult and  
15       time-consuming process. Electrode contacts and stimulation parameters (e.g., amplitude and frequency) must be tuned to provide optimal therapeutic outcomes. Therapy performance is guided by both symptom severity and side-effects, where an ideal parameter set is chosen as one that minimizes symptoms whilst also avoiding stimulation induced side-effects. In recent years, there also has been an increasing acknowledgement  
20       of the limitations of titrating therapies based on the limited snapshot of disease characteristics available from daytime clinical assessments.

          Optimization and machine-learning techniques have received considerable interest for more rapid and efficient identification of patient-specific stimulation parameters. These techniques often focus on the sequential evaluation of samples from the stimulation  
25       parameter space to characterize and guide future exploration of the parameter space or to identify an optimal parameter set for therapy. Implementation of these approaches in practice requires objective quantification of each sample selected from the parameter space in terms of its effect on patient's symptom severity. Biomarkers (i.e., signals that correlate with specific states of health or disease) can be derived from data recorded invasively or  
30       non-invasively and then utilized to assess the relationship between therapy settings and performance. A variety of such signals have been explored in the literature for a range of

neurological and psychiatric disorders, but selection of optimal therapy settings remains challenging.

It is an object of the present disclosure to improve generation of stimulation signals, by supporting selection of stimulation parameters that provide improved reduction of  
5 patient symptoms and/or side-effects induced due to stimulation.

In neurodegenerative disorders such as Parkinson's disease, therapy efficacy is lost over time due to disease progression. To counter act this, either the medication dose is increased or for those patients treated with deep brain stimulation the stimulation  
10 amplitude is increased to capture a larger portion of the target nucleus. Other variations which influence stimulation efficacy include concurrent medication intake, which in a subset of Parkinson's disease patients may induce dyskinesias when used on its own or in conjunction with deep brain stimulation. Beyond disease progression and medication intake, symptom variations over the course of the day, which may reflect circadian dependencies (e.g., Parkinson's disease, Essential Tremor, etc), could also influence  
15 therapy efficacy.

Brain stimulation-based therapies currently take a time-invariant approach to stimulation parametrisation. Optimal stimulation settings are chosen during day-time clinical visits through a process of trial and error and these settings remain the same until the next clinic visit. Adaptive brain stimulation approaches which switch stimulation ON  
20 when symptom severity or a representative neural biomarker exceeds a threshold, also assume that the subject specific thresholds do not change over time. Other adaptive stimulation approaches, such as phase-specific stimulation, also do not consider dynamic variations in effective stimulation settings and use the same stimulation phase. However, taking a time invariant approach to stimulation parametrisation may either over stimulate  
25 or under stimulate patients, resulting in stimulation induced side-effects and sub-optimal symptom management, respectively.

Although a variety of optimization approaches have been explored for neuromodulation therapy in both clinical and computational settings, a limitation of the approaches suggested to date is that the therapy parameter spaces are assumed to be time-  
30 invariant (i.e., they are fixed). In practice, patient's symptom severity can vary across multiple timescales due to factors such as changes over several hours arising due to

medication intake or the sleep/wake circadian rhythm, or fluctuations at a much longer timescale due to disease progression. For this reason, optimization approaches which focus on the identification of a time-invariant, or stationary, optimal parameter set may provide suboptimal therapy as patient symptoms, or the mapping between stimulation settings and performance, drifts in response to these temporal variations over time.

Arrangements of the present disclosure provide improved performance by addressing these limitations. A time varying controller is provided for the delivery of brain stimulation that can maintain optimal stimulation performance in the face of shifting physiological demands that might be expected from the factors introduced above.

According to an aspect of the invention, there is provided a system for generating a stimulation signal for brain or nerve stimulation, comprising: a signal generation unit configured to generate a stimulation signal for performing brain or nerve stimulation on a biological system and to transmit the stimulation signal to a stimulation unit configured to apply the stimulation to the biological system, wherein the stimulation signal is defined by at least one stimulation parameter; a data receiving unit configured to receive time series biomarker data comprising measurement samples obtained by measurements performed on the biological system at different respective times, the biomarker data representing one or more biomarkers indicative of a condition of the biological system that is affected by brain or nerve stimulation applied by the stimulation unit based on the stimulation signal; and a controller configured to control generation of the stimulation signal by the signal generation unit based on the received biomarker data, wherein: the controller is configured to use a model of the biological system to estimate an optimal value of the at least one stimulation parameter using the biomarker data, wherein the model is configured to define strengths of contribution of measurement samples to the estimate based on the times at which the measurement samples were obtained.

Configuring the model to define the strengths of contribution of samples to the estimate based on the times when the samples were obtained has been shown to improve how stimulation is parametrised at a later time, thereby improving the process of generating stimulation signals. Stimulation signals can be generated that take account of variations in the state of the biological system over time, for example allowing a medical treatment performed by the stimulation signals to remain effective over time. An approach

to therapy optimisation is thus enabled that can consider variations in therapy efficacy over time, due to factors such as medication intake, circadian rhythms and disease progression. By learning these variations, the time-varying controller can sustain therapy efficacy in contrast to time-invariant approaches which show degraded performance.

5           In an embodiment, the model comprises a forgetting component configured to progressively decrease strengths of contribution of measurement samples as a function of increasing age of the measurement samples. Use of such a forgetting component is demonstrated to provide effective tracking of changes in the biological system over time, particularly where a strength of the forgetting component (represented for example by a  
10           forgetting factor  $\epsilon$ ) is selected carefully to provide an appropriate balance between considering sufficient samples to gather information and lowering the strength of contribution of older samples.

          In an embodiment, the model further comprises a structured temporal component comprising one or more periodic components. The inventors have found that combining a  
15           forgetting component and one or more periodic components provides particularly robust tracking of changes in the biological system over time.

          As described above, while use of a forgetting component alone can provide good performance when the strength of the forgetting component is carefully selected, combining the forgetting component with a periodic component increases the range of  
20           strengths of the forgetting component for which good performance is achieved, thereby increasing ease of setup, robustness, and performance stability. Requirements for accurate selection of the strength of the forgetting component are relaxed. Good performance can be achieved, for example, even when the strength of the forgetting component is too low for the forgetting component to provide good performance when used on its own (see Figures  
25           13 and 14 and accompanying discussion below).

          Furthermore, as demonstrated and described below with reference to Figure 14, the inventors have additionally found that combining the forgetting component and the one or more periodic components provides performance that is more robust to variations in the period (or periods) of the one or more periodic components.

30           In an embodiment, the structured temporal component further comprises an aperiodic component. The aperiodic component may comprise a progressive drift. The

progressive drift may correspond to progression of a disease state of the biological system. Providing such an aperiodic component may improve the model's performance in tracking aperiodic changes in the state of the biological system (which may be represented by aperiodic changes in the shape of the objective function in cases where the model  
5 comprises a Gaussian process).

Embodiments of the disclosure will now be further described, merely by way of example, with reference to the accompanying drawings.

Figure 1 depicts a system for generating a stimulation signal and a stimulation unit for applying stimulation.

10 Figure 2 depicts an example framework for methods of generating a stimulation signal and performing stimulation using the signal.

Figures 3 and 4 depicts parameter space exploration (Figure 3) in comparison to exploitation (Figure 4) for a sequential selection of samples during Bayesian Optimization. The top panels illustrate the true (dotted black line) and estimated (solid line) shape of the underlying objective function. The confidence bounds for the estimated shape of the  
15 objective function are highlighted in grey while the associated samples used for the estimation are represented as circles. Selection of the next parameter value to be tested is subsequently determined as the minimum of an acquisition function calculated from the mean and confidence bounds of the current estimated shape of the objective function. The  
20 bottom panels illustrate a corresponding lower confidence bound acquisition function, where the value  $k$  determines whether the algorithm prioritizes exploring regions of the parameter space where there is greater uncertainty (exploration) or prioritizes selecting regions of the parameter space where there is less uncertainty (exploitation). The location of the minimum value in the acquisition function (circle) is subsequently selected as the  
25 next parameter to be tested by the Bayesian Optimization.

Figures 5 to 7 depict time-varying objective functions and their respective covariance functions. Various explored drifts are illustrated for a periodic objective function with respect to a parameter of interest (phase trigger). The objective function is allowed to drift according to three cases: gradual (Figure 5(a)), periodic (Figure 5(b)), and  
30 the superposition of the two (Figure 5(c)). At each time-step in the optimisation, samples from the objective function are weighed by their respective temporal covariance function

(depicted in Figures 6(a)-(c)). This resultant weighing is illustrated in Figures 7(a)-(c) in the “colour” (represented in greyscale) gradient of the drifting objective function, with brighter colours (lighter greys) representing a stronger sample weight.

Figure 8 depicts a Kuramoto deep brain stimulation (DBS) model response to phase-locked stimulation. Individual oscillators in the population are represented as small solid circles distributed around the unit circle where the position of each oscillator represents its phase value at the associated timestep. The population order parameter at different simulation timesteps is illustrated in Figure 8(b). In the top row of Figure 8(b) the radial line within the unit circle illustrates the magnitude of the order parameter,  $\rho$ , and the mean phase coherence of the population,  $\psi$ . The radial line points in the direction of  $\psi$  and the length of the line represents the synchronization level of the population  $\rho$  (a short line represents low synchrony, and a long line represents high synchrony). Figure 8(a) represents the distribution of oscillators prior to stimulation when the model is at its steady state behaviour. In Figure 8(b) the rows illustrate the change in population synchrony in response to stimulation at a specific phase value,  $\psi_{target}$ . Figure 8(c) summarizes the change in population synchrony,  $\Delta\rho$ , as a function of phase-locked stimulation applied between  $\psi_{target} = (-\pi, \pi)$  radians.  $\Delta\rho < 0$  corresponded to phase-locked stimulation at the specified  $\psi_{target}$  values desynchronizing the population, while  $\Delta\rho > 0$  corresponded to the phase-locked stimulation increasing population synchrony.

Figure 9 is an illustration of dynamic variation of a population ARC and PRC. The shape of the initial population ARC corresponded to a PRC of  $Z(\theta_j) = -\sin(\theta_i)$  at the first optimization step ( $t = 1$ ). By incrementally adding an offset value,  $\Delta\theta$ , to the PRF, the shape of the ARC and location of the optimum phase value for phase-locked stimulation,  $\psi^*$ , were gradually varied over the optimization process.

Figures 10(a)-(e) depict TV-BayesOpt algorithm performance for tracking a gradual drift in the optimal stimulation phase for phase-locked stimulation,  $\psi^*$ . Figure 10(a) illustrates the tracking performance of the TV-BayesOpt algorithm (upper series of dots) at locating the true optimal phase value (lower series of dots) for population desynchronization. Figure 10(b) illustrates the associated average regret for the TV-BayesOpt algorithm (upper line) at tracking the true optimum phase value in comparison to



when static BayesOpt was implemented alone (lower line). Figures 10(c)-(e) illustrate the true location of the optimal stimulation phase (black dot), the minimum predicted by the TV-BayesOpt algorithm (black circle and cross) and the shape of the GPR predicted by the TV-BayesOpt algorithm at 0.01-, 17.62- and 20.00-day time points respectively.

5        Figures 11(a)-(e) depict TV-BayesOpt algorithm performance for tracking a periodic drift in the optimal stimulation phase for phase-locked stimulation,  $\psi^*$ . Figure 11(a) illustrates the tracking performance of the TV-BayesOpt algorithm (dark grey dots) at locating the true optimal phase value (black dots) for population desynchronization. The dark grey dots and the black dots overlaid each other closely. Figure 11(b) illustrates the associated average regret for the TV-BayesOpt algorithm (upper line) at tracking the true optimum phase value in comparison to when static BayesOpt was implemented alone (lower line). Figures 11(c)-(e) illustrate the true location of the optimal stimulation phase (black dot), the minimum predicted by the TV-BayesOpt algorithm (black circle and cross) and the shape of the GPR predicted by the TV-BayesOpt algorithm at 0.01-, 17.62- and 20.00-day time points respectively.

10        Figures 12(a)-(e) depict TV-BayesOpt algorithm performance for tracking a gradual and periodic drift in the optimal stimulation phase for phase-locked stimulation,  $\psi^*$ . Figure 12(a) illustrates the tracking performance of the TV-BayesOpt algorithm (dark grey dots) at locating the true optimal phase value (black dots) for population desynchronization. The dark grey dots and the black dots overlaid each other closely. Figure 12(b) illustrates the associated average regret for the TV-BayesOpt algorithm (upper line) at tracking the true optimum phase value in comparison to when static BayesOpt was implemented alone (lower line). Figures 12(c)-(e) illustrate the true location of the optimal stimulation phase (black dot), the minimum predicted by the TV-BayesOpt algorithm (black circle and cross) and the shape of the GPR predicted by the TV-BayesOpt algorithm at 0.01-, 17.62- and 20.00-day time points respectively.

25        Figure 13 depicts a sensitivity analysis of the TV-BayesOpt algorithm with a forgetting (upper line) or a forgetting-periodic (lower line) covariance function for a range of  $\epsilon$  values. Incorporation of prior knowledge of the temporal variation in the objective function optimum value (lower line) resulted in improved TV-BayesOpt algorithm performance than implementing a forgetting covariance function (upper line) alone. Best

algorithm performance was observed for an  $\epsilon$  value of 0.22; above this value the algorithm begins to forget previous samples too quickly to accurately track the optimal value in the objective function.

Figure 14 depicts a sensitivity analysis of the TV-BayesOpt algorithm performance at tracking a periodic temporal drift when the period of the drift is offset from the covariance function period anticipated by the TV-BayesOpt algorithm. The performance of a scheduler (uppermost line), the TV-BayesOpt with a periodic temporal covariance function (middle line line) and the TV-BayesOpt with a periodic and smooth forgetting covariance algorithm (lowermost line) were estimated by calculating the AUC value for the associated regret plot for each implementation at different offset in the true temporal period.

Methods of the present disclosure are computer-implemented. Each step of the disclosed methods may therefore be performed by a computer. The computer may comprise various combinations of computer hardware, including for example CPUs, RAM, SSDs, motherboards, network connections, firmware, software, and/or other elements known in the art that allow the computer hardware to perform the required computing operations. The required computing operations may be defined by one or more computer programs. The one or more computer programs may be provided in the form of media, optionally non-transitory media, storing computer readable instructions. When the computer readable instructions are read by the computer, the computer performs the required method steps. The computer may consist of a self-contained unit, such as a general-purpose desktop computer, laptop, tablet, mobile telephone, smart device, etc. Alternatively, the computer may consist of a distributed computing system having plural different computers connected to each other via a network such as the internet or an intranet.

As exemplified schematically in Figure 1, the present disclosure relates to a system 2 for generating a stimulation signal for brain or nerve stimulation. The system 2 may perform all or part of a method according to the framework depicted in Figure 2. The method comprises receiving biomarker data from a biomarker data source in step S1, generating a stimulation signal in step S2, and applying the stimulation signal in step S3.

Measurements made during the applying of the stimulation signal in step S3 may lead to new data being added to the biomarker data in the biomarker data source.

The system 2 comprises a signal generation unit 4. The signal generation unit 4 generates a stimulation signal 5 (step S2) for performing brain or nerve stimulation on a biological system 12 (step S3). The signal generation unit 4 transmits the stimulation signal 5 to a stimulation unit 10. The stimulation unit 10 applies the stimulation to the biological system 12. In some arrangements, the stimulation unit 10 comprises one or more electrodes 14 and applies the stimulation to the biological system 12 via the one or more electrodes 14. The stimulation unit 10 may comprise a power source for applying the stimulation based on the stimulation signal. The stimulation may be achieved, for example, by driving the one or more electrodes 14 with an electrical signal provided by the power source. The biological system 12 may comprise the brain or a portion of the nervous system of a human or animal other than the brain.

The stimulation signal is defined by at least one stimulation parameter. In some arrangements, as exemplified in the detailed example described below, the stimulation signal is configured to be phase-locked to a biomarker of the biological system and the at least one stimulation parameter comprises a stimulation phase relative to the biomarker. The biomarker may, for example, comprise a disease biomarker having rhythmicity, which in the case of Parkinson's disease may range from 14-30 Hz, and for tremor may range from 3-8 Hz.

Alternatively, or additionally, the stimulation parameter may comprise one or more of the following: a stimulation amplitude; a stimulation pulse width; and a stimulation pulse frequency.

The system 2 further comprises a data receiving unit 6. The data receiving unit 6 receives time series biomarker data 7 (step S1). The biomarker data 7 comprises measurement samples (which may be referred to herein as samples) obtained by measurements performed on the biological system 12 at different respective times. The biomarker data 7 is thus derived from measurements performed on the biological system 12. The measurements may be performed using the one or more electrodes 14 and/or by other means. The biomarker data represents one or more biomarkers indicative of a condition of the biological system that is affected by brain or nerve stimulation applied by

the stimulation unit 10 based on the stimulation signal 5. The one or more biomarkers may comprise a disease biomarker having rhythmicity, such as a biomarker affected by Parkinson's disease and/or tremor. The biomarker data may indicate a degree to which the stimulation applied by the stimulation unit 10 is effective in suppressing a disease biomarker. The condition of the biological system 12 that is affected by the stimulation may thus be a disease condition such as Parkinson's and/or tremor.

The system further comprises a controller 8. The controller 8 controls generation of the stimulation signal by the signal generation unit 4 based on the received biomarker data 5. The controller 8 uses a model of the biological system 12 to estimate an optimal value of the at least one stimulation parameter using the biomarker data 5. The model is configured to define strengths of contribution (which may be referred to herein as weightings) of measurement samples to the estimate based on the times at which the measurement samples were obtained.

In some arrangements, as exemplified in the detailed example described below, the model comprises a Gaussian process. The Gaussian process may be configured to estimate an objective function representing an expected variation of the biomarker data as a function of the at least one stimulation parameter of the stimulation signal. The objective function may for example be considered as a mapping between stimulation settings and symptom suppression. By estimating the objective function, the Gaussian process can thus enable estimation of a value of the at least one stimulation parameter that optimally suppresses a disease biomarker. In the detailed example described below, an amplitude response curve (ARC) representing a change in a population synchrony ( $\rho$ ) (of a population of neurons) due to stimulation delivered at a particular phase ( $\psi$ ) is used as an objective function. By configuring the model to define strengths of contribution of measurement samples to the estimate based on the times at which the measurement samples were obtained, the model can take account of changes in the objective function over time.

The Gaussian process may be defined by a covariance function. As exemplified in the detailed example described below, the covariance function may comprise a temporal part. The covariance function may additionally comprise a time-invariant spatial part representing prior knowledge of a shape of the objective function.

In some arrangements, the temporal part and spatial part are combined into a spatio-temporal covariance function using the Hadamard product, as described below with reference to equation (9).

In some arrangements, the spatial part is implemented as described below with  
5 reference to equation (10).

The model (e.g., Gaussian process) may comprise a forgetting component. In cases where the model comprises a Gaussian process, the forgetting component may be represented by the temporal part of the covariance function. The forgetting component is configured to progressively decrease strengths of contribution of measurement samples as  
10 a function of increasing age of the measurement samples. A strength of the forgetting component (e.g., how quickly the forgetting component causes contributions from older measurement samples to be suppressed) may be represented by a forgetting factor (referred to as  $\epsilon$  in the detailed example below). The forgetting factor may, for example, have a value between 0 and 1, where 0 represents no samples being forgotten and values near to 1  
15 correspond to samples being forgotten so quickly that only a current sample is used to estimate the objective function.

In some arrangements, the forgetting component is implemented as described below with reference to equations (11) and (12). In some arrangements, the forgetting factor may be learned using past data (e.g., past samples).

20 The model (e.g., Gaussian process) may comprise a structured temporal component. In cases where the model comprises a Gaussian process, the structured temporal component may be represented by the temporal part of the covariance function. As exemplified in the detailed example described below, the structured temporal component may comprise one or more periodic components. Each periodic component is  
25 configured to contribute a variation in the strength of contribution of measurement samples to the estimate that is periodic as a function of time. Each periodic component may thus promote a greater degree of similarity between strengths of contributions of measurement samples taken at corresponding time points in different cycles. For example, if a periodic component has a period of one day, the periodic component will contribute to samples  
30 taken at the same time on different days having more similar weightings than samples taken at different times of day. Furthermore, when the model is estimating an optimal

value of the at last one stimulation parameter at a particular time point in the cycle of a periodic component, the periodic component will promote increased strengths of contribution to the estimate from samples taken at the same time point in previous cycles. This is illustrated for example in Figures 6(b) and 7(b), discussed below. In some  
5 arrangements, the periodic component is implemented using an exponential sine squared covariance function as exemplified in the detailed example below.

In some arrangements, the periodic component is implemented as described below with reference to equation (13).

In some arrangements, the combination of the forgetting component and the one or  
10 more periodic component is implemented as described below with reference to equation (14).

The one or more periodic components may arise for example due to factors such as biological rhythms, such as the circadian sleep-wake cycle or medication cycles. The one or more periodic components may have a predetermined period (e.g., one day, or a period  
15 between administration of treatments to a patient. Alternatively, or additionally, the controller 8 may be configured to learn a period of one or more of the periodic components. The structured temporal component may additionally or alternatively comprise an aperiodic component. As exemplified in the detailed example described below, the aperiodic component may comprise a progressive (e.g., monotonic) drift. The  
20 progressive drift may be caused, for example, by progression of a disease state of the biological system 12.

Embodiments of the disclosure may be implemented as methods. The methods may include computer-implemented methods of generating a stimulation signal for brain or nerve stimulation. The methods may comprise a method of performing brain or nerve  
25 stimulation. Such methods may include generating a simulation signal and applying stimulation to the biological system based on the generated stimulation signal. The methods may include steps within the example framework method depicted in Figure 2. A computer program and/or a computer program product (e.g., a non-transitory computer program product) may be provided for executing computer-implemented method steps.

30 For example, the method may comprise receiving time series biomarker data (e.g., step S1 in Figure 2). The biomarker data comprises measurement samples obtained by

measurements performed on a biological system at different respective times. The biomarker data represents one or more biomarkers indicative of a condition of the biological system that is affected by brain or nerve stimulation applied by a stimulation unit based on a stimulation signal. The stimulation signal is defined by at least one  
5 stimulation parameter. The method may comprise generating a stimulation signal (e.g., step S2 in Figure 2) for the stimulation unit based on the received biomarker data by using a model of the biological system to estimate an optimal value of the at least one stimulation parameter using the biomarker data. The model is configured to define strengths of contribution of measurement samples to the estimate based on the times at which the  
10 measurement samples were obtained. The generating of the stimulation signal may comprise generating a stimulation signal for which the at least one stimulation parameter has the estimated optimal value.

### Detailed Example

#### 15 Overview

Reference is made below to an algorithm based upon Bayesian optimization (referred to as “BayesOpt”) with modifications to enable time-adaptive therapy optimization (Time-Varying Bayesian Optimisation, referred to as “TV- BayesOpt”). The TV-BayesOpt algorithm is an example of how the controller 8 may implement the  
20 estimation of the optimal value of the at least one stimulation parameter using biomarker data, as described above.

Further details are given below about implementation choices for TV- BayesOpt, namely the *surrogate model* (i.e., a model used to estimate the shape of the true objective function) and the *acquisition function* (i.e., how to select the next sample from the  
25 parameter space). The surrogate model is an example of the model of the biological system discussed above. The acquisition function is an example of how the controller may use the biomarker data, for example to update the model and/or exploit the model to provide an estimation of an optimal value of the at least one stimulation parameter.

*Spatiotemporal covariance functions* used for a time-varying implementation of the  
30 BayesOpt algorithm (*TV-BayesOpt*) are described that incorporate prior knowledge regarding the form of the objective function along with its anticipated drift in time. These

covariance functions are examples of the covariance functions defining the Gaussian process implementation of the model of the biological system discussed above. A form is established of a dynamical model of neural synchrony and is used to evaluate the performance of the system 2 in delivering optimal stimulation to suppress a rhythmic biomarker as a proxy for symptom suppression.

### *Kuramoto Model*

The Kuramoto model was used to investigate the performance of the TV-BayesOpt algorithm. The Kuramoto Model provides a simple mathematical model to describe synchronization in biological systems. In this specific case the model represents a population of neurons as  $N$  coupled oscillators that oscillate with an intrinsic frequency  $\omega_{biomarker}$ . This intrinsic frequency reflects the rhythmicity of the disease biomarker, which in the case of Parkinson's disease ranges from 14-30 Hz, and for tremor ranges from 3-8 Hz. The strength of interaction between the oscillators is captured by a coupling term  $K_{biomarker}$ , where the influence of the  $j^{th}$  oscillator on the  $i^{th}$  oscillator is described in terms of phase progression (i.e. speeding up or slowing down) of the  $i^{th}$  oscillator depending upon their difference in phase.

$$\frac{d\theta_i}{dt} = \omega_{biomarker_i} + \frac{K_{biomarker}}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i) \quad (1)$$

The model has been extended to account for the effects of exogenous perturbations which is given by:

$$\frac{d\theta_i}{dt} = \omega_{biomarker_i} + \frac{K_{biomarker}}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i) + IX(t)Z(\theta_i) \quad (2)$$

where  $I$  is the intensity of stimulation,  $X(t)$  is a function whose value equals 1 at time  $t$  if stimulation is applied and is 0 otherwise, and  $Z(\theta_i)$  is the phase response curve (PRC), which describes the response of the  $i^{th}$  oscillator to stimulation (i.e. speeding up or slowing down) depending on the precise timing of the stimulation with respect to the oscillator's phase,  $\theta_i$ . Neurons have been classified into two distinct categories based on their PRCs that characterise their response to perturbation: type I PRCs either exclusively delay or advance spike firing, whereas type II PRCs can both advance or delay dependent



upon the specific phase at which stimulation is delivered. It has been demonstrated that neural oscillators in tremor exhibit type II PRCs. Therefore, in the present example, the PRC is described as  $Z(\theta_i) = -\sin(\theta_i)$ .

Under certain assumptions, the Kuramoto model can be solved to express the  
 5 behaviour of each oscillator in terms of the order parameters,  $\rho$  and  $\psi$ , describing the population's mean phase-coherence and mean phase respectively.

$$\frac{d\theta_i}{dt} = \omega_{\text{biomarker}_i} + K_{\text{biomarker}} \rho \sin(\psi - \theta_i) + IX(t)Z(\theta_i) \quad (3)$$

The population's mean phase-coherence,  $\rho$ , represents synchrony on the interval  $[0, 1]$  where  $\rho = 0$  and  $\rho = 1$  correspond to either complete desynchrony or synchrony, respectively. In this manner,  $\rho$  can be used as a control signal used a proxy for population  
 10 level synchrony similar to that currently used in adaptive deep brain stimulation (DBS) to capture instantaneous biomarker power, or level of synchronization, in specific frequency bands of a neural oscillation. Following from the introduction of oscillopathies, it is assumed that lower levels of  $\rho$  corresponds to reduced disease symptoms whereas higher levels of  $\rho$  reflects increased symptom severity. The performance of the TV-BayesOpt  
 15 algorithm can then be assessed by quantifying how the precise stimulation timing, i.e., the optimal  $\psi$  at which to apply stimulation, can minimize population synchrony (in terms of  $\rho$ ).

The objective function for the TV-BayesOpt algorithm is represented as the effect of phase-locked stimulation on synchrony ( $\rho$ ). This can be quantified with an amplitude  
 20 response curve (ARC) where the ARC represents the change in the population synchrony ( $\rho$ ) due to stimulation delivered at a particular phase ( $\psi$ ). The ARC therefore represents the underlying objective function (mapping between stimulation settings and symptom suppression) where certain phases of stimulation may increase or decrease population synchrony. The shape of population ARC can be varied over time by updating the PRC  
 25 ( $Z(\theta_i)$ ) for the oscillators to emulate gradual or periodic changes.

### *Bayesian Optimization*

To estimate the form of the ARC that maps from stimulation parameters to population synchrony (i.e., the objective function), a surrogate model based on a Gaussian

Process (GP) prior was implemented. The GP prior model can be thought of as a probability distribution over a set of non-linear regression functions (as opposed to a Gaussian distribution, which is a set over variables). This permits efficient estimation of a mean function, as well as its relative uncertainty across the parameter space, indicating the confidence in the evaluated objective function for each parameter value.

The objective function  $f(x)$  is fully defined by its mean  $\mu(x)$  and covariance  $K(x, x')$ , that form sufficient statistics such that  $f(x) \sim GP(\mu(x), K(x, x'))$ , where  $x$  in this context are the stimulation parameters to be explored. The covariance function describes the relationship between points in the parameter space, and is vital in determining the form (e.g., smoothness, or periodicity) of the estimated objective function. Thus  $K(x, x')$  describes the spatial covariance (i.e., the space between parameter  $x$  and its neighbour  $x'$ ). An example covariance function for use with arrangements of the present disclosure is described below.

The mean and covariance functions allow us to derive a joint distribution of the noisy outputs  $\mathbf{y}$  (previously sampled simulation parameters) and the estimated value of the objective function (mapping between stimulation settings and symptom suppression)  $f^*$  evaluated at a new sample point  $x^*$  in the parameter space. This may be defined as:

$$\begin{bmatrix} \mathbf{y} \\ f_* \end{bmatrix} \sim \left( \begin{bmatrix} \mu(\mathbf{x}) \\ \mu(x_*) \end{bmatrix}, \begin{bmatrix} \mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I} & \mathbf{K}(\mathbf{x}, x_*) \\ \mathbf{K}(x_*, \mathbf{x}) & K(x_*, x_*) \end{bmatrix} \right) \quad (4)$$

where  $\mathbf{y} \sim GP(\mu(\mathbf{x}), \mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I})$ ,  $\mathbf{x} = [x_1, \dots, x_n]$  and  $\mathbf{y} = [y_1, \dots, y_n]$ . As additional samples from the parameter space are evaluated, the prior is updated to form a posterior distribution to improve the model's approximation of the shape of the objective function,  $f(x)$ . Conjugacy between the prior and likelihood allows the posterior to be computed analytically by omitting the evidence term. The predictive distribution can therefore be written as

$$p(f_* | D_n, x_*) \sim N(\mu_*, \sigma_*^2) \quad (5)$$

Where  $D_n = \{\mathbf{x}, \mathbf{f}\}$  represents the previously sampled parameter values and their respective outcome values and the posterior distribution is defined as

$$\mu_* = \mathbf{K}(x_*, \mathbf{x}) [\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y} \quad (6.a)$$

$$\sigma_*^2 = K(x_*, x_*) - \mathbf{K}(x_*, \mathbf{x}) [\mathbf{K}(\mathbf{x}, \mathbf{x}) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{K}(\mathbf{x}, x_*) \quad (6.b)$$

The acquisition function is used to identify regions of the parameter space which should be tested during the optimization process. To determine the next sample to test during the optimization process, the acquisition function  $\alpha(x)$  uses the current estimate of the underlying surrogate model and calculates the expected utility of all samples in the parameter space. The next sample to be tested is subsequently selected as the sample point  $x^*$  which minimizes the acquisition function:

$$x_* = \underset{x_*}{\operatorname{argmin}} \alpha(x; D_n) \quad (7)$$

A Gaussian Process - Lower Confidence Bound (GP-LCB) acquisition function was chosen for the BayesOpt algorithm as it allows flexible tuning of the amount of parameter space exploration during the optimization process. The GP-LCB is defined as

$$\alpha_*(x_*) = \mu_*(x_*) - \kappa_n \sigma_*(x_*) \quad (8)$$

where the parameter  $\kappa_n$  rescales  $\alpha_*$  - the estimated uncertainty of the GP. To balance between exploration of under sampled regions of the parameter space, or exploit already sampled regions, the value of  $\kappa_n$  in the acquisition function is chosen according to the specific needs of the user where high  $\kappa_n$  values favour exploration, and low  $\kappa_n$  values favour exploitation, of the parameter space. The role of  $\kappa_n$  for guiding this exploration-exploitation trade-off when sampling from the parameter space between optimization steps is illustrated in Figure 3 (exploration) and Figure 4 (exploitation) in which examples are shown of either maximally explorative or exploitative acquisition functions.

To track a dynamically varying set of optimal therapy parameters, arrangements of the present disclosure eschewed the common assumption that the objective function is time invariant. When applied to time-varying systems the standard BayesOpt routine will exhibit diminished performance as it assumes all samples have equal contribution to the current estimate of the objective function. A TV-BayesOpt algorithm in contrast weighs the contribution of previously tested samples to the current estimate based on the time at which samples were taken in relation to the current estimation step. Sources of temporal variation which lead to changes in the optimal parameter(s) may be gradual (e.g., linear), such as during disease progression, or due to other factors such as biological rhythms, such as the circadian sleep-wake cycle or medication cycles (e.g., periodic). The goal of the TV-BayesOpt algorithm is thus to find the optimal parameter set for a given time, i.e., the

parameters which provide the best suppression of symptoms, and to subsequently track the development of this parameter set over time.

To meet these requirements, a spatio-temporal covariance function was constructed. The spatio-temporal covariance function comprises two parts which are referred to as the spatial covariance function (or spatial part) and the temporal covariance function (or temporal part), respectively. The spatial covariance function in this context is equivalent to the covariance function used for time- invariant BayesOpt and captures the prior knowledge of the shape of the objective function. On the other hand, the temporal covariance function captures the anticipated variation in the shape of the objective function over time, i.e., rescaling the relevance of previously sampled data in relation to the most recent sample. The resulting spatio-temporal covariance function,  $\widetilde{K}_t(x, x')$ , is formulated as the Hadamard product of the spatial,  $K(x, x')$ , and the temporal covariance functions,  $K(t, t')$ :

$$\widetilde{K}_t(x, x') = K(x, x') \odot K(t, t') \quad (9)$$

A multitude of covariance functions exist across the field of machine learning that can describe complex topologies. The performance of the BayesOpt algorithm is thus dependent on the selection of an appropriately shaped covariance function for the problem that is being investigated.

In the present example, the problem investigated concerns the optimization of a target phase for phase-locked stimulation to suppress a rhythmic biomarker of disease. Therefore, the parameter space (target phase values) is periodic, where  $-\pi$  and  $\pi$  radians represent the same point in the parameter space. Therefore, the spatial covariance function (i.e., the covariance between target stimulation phases in the parameter space) was therefore selected to also be periodic using the exponential sine squared covariance function:

$$K(x, x') = \exp\left(-\frac{2 \sin^2\left(\frac{\pi|x - x'|}{T_x}\right)}{l_x^2}\right) \quad (10)$$

where hyperparameters  $l_x$  and  $T_x$  are the length scale and period (incorporating our prior knowledge of the periodicity of the parameter space).

For biological systems there are several sources which may lead to variations in the shape of the objective function over time, including disease progression (leading to monotonic changes) and biological rhythms such as the circadian sleep-wake cycle (leading to periodic changes), or a combination of the two. Examples of temporal covariance functions that overcome these temporal variations are outlined in the following sections. Topologies of these covariance functions are summarized below in Figures 5-7.

To accommodate gradual variations in the shape of the objective function a ‘forgetting’ covariance function that reduces the influence of older samples on the current estimate can be implemented, as depicted in Figures 5(a), 6(a), and 7(a). Bogunovic et al. (2016) introduced a covariance function that was defined to meet these requirements (Bogunovic I, Scarlett J, Cevher V. Time-varying Gaussian process bandit optimization. *Artif Intell Stat.* 2016;314–23), where the covariance function is defined as

$$K(t, t') = (1 - \epsilon)^{\frac{|t-t'|}{2}} \quad (11)$$

where  $t$  is the time of the most current sample,  $t'$  is the time at which previous samples were taken and  $\epsilon$  is referred to as the forgetting factor. The forgetting factor,  $\epsilon$ , has a value between  $[0, 1)$  which determines how quickly the contribution of previously taken samples to the current estimate are reduced. A forgetting factor  $\epsilon = 0$  corresponds to no data being forgotten and all previous samples contributing equally to the estimate of the objective function, while  $\epsilon$  values close to 1 correspond to previous data samples being forgotten so fast that only the current sample is used to estimate the objective function. Moreover, as  $|t - t'| \rightarrow \infty$  the contribution of previous samples on the current estimate reduces to 0. This behaviour is graphically summarized in Figure 6(a).

An appropriate value of  $\epsilon$  can be learned using past data through maximum likelihood estimation or by empirically selecting  $\epsilon$  to reduce the contribution of previous samples based on a specified half-life:

$$\epsilon = \lambda = \frac{\ln 2}{t_{1/2}} \quad (12)$$

Where  $\lambda$  corresponds to the half-life decay factor and  $t_{1/2}$  defines half-life as a specified number of samples, i.e., a  $t_{1/2}$  value of 100 corresponds to the objective function value

associated with the 100th last sample taken contributing only 50 % of its original value to the current estimate of the surrogate model.

To accommodate periodic variations in the shape of the objective function a periodic covariance function (i.e., a temporal covariance function part representing a structured temporal component that is periodic) can be implemented to weight the contribution of previously sampled data based on the specific time points in the cycle at which the samples were taken, as depicted in Figures 5(b), 6(b), and 7(b). In this manner, the contribution of samples taken at time points near the same parts of the cycle are weighted more similarly to each other (i.e., such that the strengths of contribution of these measurement samples to the estimate are more similar to each other) than those taken at different time points in the cycle. This behaviour can be captured in a similar manner to the spatial covariance function by using the exponential sine squared covariance function,

$$K(t, t') = \exp \left( - \frac{2 \sin^2 \left( \frac{\pi |t - t'|}{T_t} \right)}{l_{t,p}^2} \right) \quad (13)$$

where  $T_t$  is the temporal period of the underlying biomarker oscillation (defined with the same units as  $t$ ) and  $l_{t,p}$ , the length scale, controls the amplitude of the temporal oscillation. In practice, the units of  $T_t$  will depend on the sampling protocol and may be defined in either sample counts or physical timescales such as hours or minutes. Furthermore,  $T_t$  can be estimated empirically by tracking the variation of symptom suppression under fixed stimulation parameters. Use of this covariance function captures temporal variations which repeat themselves exactly, as depicted in Figure 6(b).

For implementations in real-life scenarios where both gradual and periodic temporal variations may be superimposed due to both disease progression and biological rhythms, the forgetting component and the structured temporal component (e.g., periodic component) of the temporal covariance function can be combined, as depicted in Figures 5(c), 6(c), and 7(c). Doing so produces the temporal covariance function below which is the product of the forgetting covariance function (forgetting component) and the periodic covariance function (structured temporal component)

$$K(t, t') = (1 - \epsilon)^{\frac{|t-t'|}{2}} \exp\left(-\frac{2 \sin^2\left(\frac{\pi|t-t'|}{T_t}\right)}{l_{t,p}^2}\right) \quad (14)$$

where  $\epsilon$ ,  $l_{t,p}$  and  $T_t$  are the covariance function hyperparameters as described above in the previous sections. This temporal covariance function is used to subsequently reduce the contribution of older samples on the current estimate while also flexibly tracking periodic variations in the shape of the objective function over time, as depicted in Figure 6(c).

5

### *Simulation Details*

The behaviour of the TV-BayesOpt algorithm was investigated in the Kuramoto DBS model. The Kuramoto DBS model was initially simulated for a 200s period to allow the model to reach its steady state behaviour. After 200s, the model was subsequently  
 10 simulated for 3000 optimization steps where each optimization step corresponded to a 58s simulation period. During each optimization step, phase-locked stimulation at a target phase value was applied only during the last 8s of the 58s simulation period. The initial 50s of each optimization step was included to allow the model to return to its steady state post application of phase-locked stimulation in the last optimization step. The output of each  
 15 optimization step,  $y$ , was quantified as the normalized change in the mean order parameter during the stimulation period from its baseline value:

$$y = \Delta\rho = \frac{\overline{\rho_{stim}} - \overline{\rho_{baseline}}}{\overline{\rho_{baseline}}} \quad (15)$$

where  $\overline{\rho_{stim}}$  is the mean order parameter averaged over the 8 s stimulation period and  $\overline{\rho_{baseline}}$  is the baseline mean order parameter value and was calculated as the mean order parameter of the 25s period prior to stimulation application.

20 To initiate the TV-BayesOpt algorithm, data from the parameter space was needed. To this end, 12 equally spaced samples between  $[-\pi, \pi)$  were tested to coarsely characterize the population response to phase-locked stimulation at different stimulation phase values. Following this, selection of subsequent phase values to be tested at each optimization were selected by the TV-BayesOpt algorithm as described above.

25 Drifts in the objective function (mapping between stimulation settings and symptom suppression) shape were simulated by incrementally adding a phase offset value

to the Kuramoto DBS model PRC over the steps of the optimization process. Three drifts in the optimum were simulated to explore the TV-BayesOpt algorithm performance at locating and tracking an optimal phase value for desynchronizing the Kuramoto DBS model during phase- locked stimulation. The performance of the algorithm at each

5 timestep was quantified using the cumulative regret, defined as

$$CR = \frac{\sum_{k=1}^n (f(\psi_k^*) - f(\psi_{target_k}))}{n} \quad (16)$$

where  $\psi_k^*$  is the location of the true optimum value phase value for phase-locked stimulation at optimization step  $k$ ,  $\psi_{target_k}$  is the stimulation phase value tested at the  $k^{th}$  optimization step,  $f(*)$  is the measured objective function value, i.e.  $\Delta\rho$ , at the specified input value and  $n$  is the total number of optimization steps.

10 A gradual drift was simulated as a phase offset advancing from 0 to  $-\pi$  over 3000 optimization steps, a periodic drift was simulated as a sinusoidal phase offset advancing from 0 to  $-\pi$  and back to 0 over 100 optimization steps, and lastly the superimposed gradual and periodic drifts were simulated as the superposition of the gradual and periodic drifts.

15 Finally, sensitivity analyses were conducted to investigate the influence of the temporal covariance forgetting factor,  $\epsilon$ , and period,  $T_t$ , on the TV-BayesOpt algorithm's performance at tracking a gradual and periodic temporal drift and purely periodic temporal drift, respectively.

## 20 *Results*

The performance of the TV-BayesOpt algorithm was investigated in the Kuramoto DBS model for locating and tracking an optimum phase value to reduce population synchrony during phase-locked stimulation. The behaviour of the Kuramoto DBS model in response to phase-locked stimulation at different population phase values was first

25 investigated prior to simulation of gradual and periodic drifts in the location of the optimal stimulation phase. The performance of the algorithm in each scenario was investigated in comparison to the traditional time-invariant implementation of the BayesOpt algorithm.

After an initial 200s transient period the Kuramoto DBS model reached its steady state behaviour with the magnitude of the order parameter  $\rho$  (i.e., population synchrony



which is a direct correlate of symptom severity) settling to a value around 0.8, reflecting a highly synchronized state. Once the model reached its steady state behaviour, the population's response to phase-locked stimulation was investigated using a static PRC for the oscillator population that was set to  $-\sin(\theta)$ . Stimulation was applied to the population at 12 equally spaced mean population phase values,  $\psi_{target}$ , between  $(-\pi, \pi)$  (Figure 8). As shown in Figure 8(c), phase-locked stimulation of the population between  $\psi_{target} = (-\pi, -\frac{\pi}{2})$  radians and  $\psi_{target} = (\frac{\pi}{2}, \pi)$  radians resulted in desynchronization of the population and a reduction in the magnitude of the order parameter ( $\Delta\rho < 0$ ). In contrast, phase-locked stimulation applied between  $\psi_{target} = (-\frac{\pi}{2}, \frac{\pi}{2})$  radians resulted in increased population synchrony, which was reflected in an increase in the magnitude of the population order parameter ( $\Delta\rho > 0$ ). The optimal phase value for phase-locked stimulation,  $\psi^*$ , was identified as the value that resulted in the greatest reduction in population synchrony. This value was identified as  $\psi^* = \pi$  radians. Following this, the location of  $\psi^*$  was dynamically varied as described in the subsequent sections.

The location of the optimal stimulation phase,  $\psi^*$ , achieving maximum population desynchrony, was varied by incrementally adding an offset value,  $\Delta\theta$ , to the initial population PRF: i.e.,  $Z(\theta_j) = -\sin(\theta_j + \Delta\theta)$ . This is represented in Figure 9 where the shape of the population ARC and PRC are dynamically varied by inclusion of  $\Delta\theta$ . The location of  $\psi^*$  could thus be gradually varied over the course of the optimization process (Figure 9).

A gradual drift in the optimal stimulation phase value,  $\psi^*$ , needed to desynchronize the population of oscillators was tested using the TV-BayesOpt algorithm plus a temporal forgetting covariance function and a simulation in which the optimal phase drifts from  $\pi$  to 0 radians (Figure 10). Such a gradual drift may be anticipated in a context such as disease progression. The TV-BayesOpt algorithm resulted in lower cumulative regret over the optimization process, calculated as the area under the curve (AUC) of the cumulative regret plot, when compared to the time-invariant implementation of the BayesOpt algorithm (Figure 10).

A periodic drift in the optimal phase value,  $\psi^*$ , for phase-locked stimulation was emulated by periodically varying  $\psi^*$  from  $\pi$  to 0 radians over 100 optimization steps

(representing a 1-day period). Implementation of the TV-BayesOpt algorithm with a temporal periodic covariance function (where  $T_t = 100$ ) led to accurate tracking of the location of the drifting  $\psi^*$ , as depicted in Figure 11. Such a periodic change could be expected from regular medication intake, circadian cycles, and other repeating patterns which may influence therapy efficacy. The temporal periodic covariance function with a 1-day period led to better tracking of  $\psi^*$  in comparison to static BayesOpt, as evidenced by a lower cumulative regret AUC (Figure 11).

Superposition of gradual and periodic drifts in the optimal stimulation setting for phase-locked stimulation,  $\psi^*$ , of the Kuramoto oscillators required implementation of the temporal forgetting-periodic covariance function. Implementation of the TV-BayesOpt algorithm with this covariance function resulted in accurate tracking of the location of  $\psi^*$  and produced a lower cumulative regret AUC in comparison to the time-invariant BayesOpt algorithm (Figure 12).

A sensitivity analysis was conducted to investigate the influence of the temporal kernel's forgetting factor value,  $\epsilon$ , on the performance of the TV-BayesOpt algorithm (Figure 13). For each  $\epsilon$  value investigated, the AUC of the cumulative regret plot was calculated for tracking the gradual and periodic drift illustrated in Figure 12, with high values indicating inaccurate optimization. Variation of the TV-BayesOpt algorithm performance due to prior knowledge of the temporal drift was investigated by comparing the algorithm performance when solely implementing a forgetting covariance function in comparison to a forgetting-periodic covariance function; Figure 13 (upper and lower curves respectively). For  $\epsilon$  values less than 0.04 the forgetting covariance function resulted in greater cumulative regret AUC values than the forgetting-periodic covariance function. For  $\epsilon$  values above 0.04, both temporal covariance functions resulted in equivalent average regret AUC measures, where the algorithm provided the best performance with an  $\epsilon$  value of 0.22 and corresponded to a data half-life ( $t_{1/2}$ ) of 3.15 samples. Incorporating prior knowledge of periodicity in temporal variation by using the temporal periodic covariance function led to more stable performance of the TV-BayesOpt algorithm than when the forgetting factor alone was implemented. Consistent cumulative regret AUC values lower than the forgetting covariance function performance alone were observed for the algorithm when incorporating the temporal periodic covariance function.

$\epsilon$  values above 0.22 resulted in increased cumulative regret AUC values as  $\epsilon$  values in this range led to the algorithm forgetting data too quickly to accurately estimate the location of the optimal value.

Thus, in summary, good algorithm performance is found to be achieved by  
 5 implementation of the TV-BayesOpt algorithm with smooth forgetting alone, but this requires careful selection of an appropriate forgetting factor well-matched to the temporal oscillation being tracked, as shown in Figure 13. When too low a value of  $\epsilon$  is selected, the algorithm weighs the contribution of older samples too heavily which leads to poor tracking of the optimum stimulation phase over time, Figure 13. When the TV-BayesOpt  
 10 algorithm additionally incorporates a periodic temporal covariance function whose period is aligned to the temporal drift the algorithm performance is greatly improved for low  $\epsilon$  values. Robust algorithm performance is subsequently observed for all  $\epsilon$  values below 0.4. Above this value the algorithm performance begins to worsen due to previously taken data samples being forgotten too quickly.

15 A sensitivity analysis was conducted to investigate the TV-BayesOpt algorithm performance when the temporal drift anticipated by the algorithm,  $T_t$ , differed from the true period of the example periodic temporal variation highlighted in Figure 11. The performance of the algorithm was also compared to that of a scheduler which made scheduled adjustments to maintain stimulation at the optimal phase value over a 24-hour  
 20 period, Figure 14.

The scheduler and TV-BayesOpt with a periodic temporal covariance function were unable to track the location of the optimum stimulation phase when the temporal drift in the optimal stimulation phase was different from the anticipated period. This resulted in large AUC values for both the scheduler and algorithm when the offset to the true temporal  
 25 period was not 0.

When the temporal covariance function for the TV-BayesOpt algorithm was implemented with smooth forgetting in addition to a periodic kernel the algorithm was able to accommodate the offset in the difference between the true and anticipated temporal periods, as evidenced by the lower regret AUC values.

30 Thus, in summary, the observations discussed above demonstrate that when there is no difference between the actual and expected temporal periods, the algorithm performs

well with a periodic temporal covariance function alone and maintains tracking of the optimal stimulation phase. In this case, once the algorithm has locked to the location of the optimum stimulation phase its performance is equivalent to a scheduler that updates the location of the optimum stimulation phase value based on the anticipated temporal period (Figure 14). Worsened performance is observed by both the algorithm and scheduler however when there is a difference between the anticipated and actual temporal period (Figure 14). This worsened algorithm performance can be compensated for by use of the smooth forgetting covariance function which maintains algorithm performance at tracking the location of the optimum phase when there is a mismatch between the actual and anticipated temporal period by the TV-BayesOpt algorithm.

The above results highlight the relationship between smooth forgetting and temporal rhythmicity in the temporal covariance function. Incorporation of both smooth forgetting and knowledge of the rhythmic variations in therapy efficacy leads to better performance of the TV- BayesOpt algorithm than when either of these are utilized by the algorithm alone.

## CLAIMS

1. A system for generating a stimulation signal for brain or nerve stimulation, comprising:
- 5 a signal generation unit configured to generate a stimulation signal for performing brain or nerve stimulation on a biological system and to transmit the stimulation signal to a stimulation unit configured to apply the stimulation to the biological system, wherein the stimulation signal is defined by at least one stimulation parameter;
- a data receiving unit configured to receive time series biomarker data comprising
- 10 measurement samples obtained by measurements performed on the biological system at different respective times, the biomarker data representing one or more biomarkers indicative of a condition of the biological system that is affected by brain or nerve stimulation applied by the stimulation unit based on the stimulation signal; and
- a controller configured to control generation of the stimulation signal by the signal
- 15 generation unit based on the received biomarker data, wherein:
- the controller is configured to use a model of the biological system to estimate an optimal value of the at least one stimulation parameter using the biomarker data, wherein the model is configured to define strengths of contribution of measurement samples to the estimate based on the times at which the measurement samples were obtained.
- 20
2. The system of claim 1, wherein the model comprises a forgetting component configured to progressively decrease strengths of contribution of measurement samples as a function of increasing age of the measurement samples.
- 25
3. The system of claim 2, wherein the model further comprises a structured temporal component.
4. The system of claim 3, wherein the structured temporal component comprises one or more periodic components.
- 30
5. The system of claim 4, wherein each periodic component is configured to

contribute a variation in the strength of contribution of measurement samples to the estimate that is periodic as a function of time.

6. The system of claim 4 or 5, wherein one or more of the periodic components has a  
5 predetermined period, optionally corresponding to a periodicity of an underlying biomarker oscillation of the biological system.

7. The system of any of claims 4 to 6, wherein the controller is configured to learn a  
10 period of one or more of the periodic components.

8. The system of any of claims 3 to 7, wherein the structured temporal component further comprises an aperiodic component.

9. The system of claim 8, wherein the aperiodic component comprises a progressive  
15 drift, the progressive drift optionally corresponding to progression of a disease state of the biological system.

10. The system of any of claims 3 to 9, wherein the model comprises a Gaussian  
20 process defined by a covariance function.

11. The system of claim 10, wherein the Gaussian process is configured to estimate an objective function representing an expected variation of the biomarker data as a function of the at least one stimulation parameter of the stimulation signal.

25 12. The system of claim 10 or 11, wherein the covariance function comprises a temporal part, the temporal part representing the forgetting component and the structured temporal component.

30 13. The system of claim 12, wherein the structured temporal component comprises a periodic component and the temporal part of the covariance function,  $K(t, t')$ , comprises the following:

$$K(t, t') = (1 - \epsilon)^{\frac{|t-t'|}{2}} \exp\left(-\frac{2 \sin^2\left(\frac{\pi|t-t'|}{T_t}\right)}{l_{t,p}^2}\right)$$

where  $t$  is the time of a most current measurement sample,  $t'$  is a time at which a previous measurement sample was taken,  $\epsilon$  is a forgetting factor defining a strength of the forgetting component,  $l_{t,p}$  defines a length scale of the periodic component, and  $T_t$  is a temporal  
 5 period of the periodic component.

14. The system of any of claims 10 to 13, wherein the covariance function further comprises a time-invariant spatial part representing prior knowledge of a shape of the objective function, the temporal part and the spatial part optionally being combined using  
 10 the Hadamard product.

15. The system of any preceding claim, wherein:  
 the stimulation signal is configured to be phase-locked to a biomarker of the biological system and the at least one stimulation parameter comprises a stimulation phase  
 15 relative to the biomarker; and/or  
 the at least one stimulation parameter comprises a stimulation amplitude; and/or  
 the at least one stimulation parameter comprises a stimulation pulse width; and/or  
 the at least one stimulation parameter comprises a stimulation pulse frequency.

20 16. The system of any preceding claim, further comprising the stimulation unit.

17. The system of claim 16, wherein the stimulation unit comprises one or more electrodes for applying the stimulation signal to the biological system.

25 18. A computer-implemented method of generating a stimulation signal for brain or nerve stimulation, the method comprising:  
 receiving time series biomarker data comprising measurement samples obtained by measurements performed on a biological system at different respective times, the biomarker data representing one or more biomarkers indicative of a condition of the

biological system that is affected by brain or nerve stimulation applied by a stimulation unit based on a stimulation signal, wherein the stimulation signal is defined by at least one stimulation parameter; and

- generating a stimulation signal for the stimulation unit based on the received  
5 biomarker data by using a model of the biological system to estimate an optimal value of the at least one stimulation parameter using the biomarker data, wherein the model is configured to define strengths of contribution of measurement samples to the estimate based on the times at which the measurement samples were obtained.

- 10 19. The method of claim 18, wherein the generating of the stimulation signal comprises generating a stimulation signal for which the at least one stimulation parameter has the estimated optimal value.

20. A method of performing brain or nerve stimulation, comprising:  
15 generating a stimulation signal using the method of claim 20; and  
applying stimulation to the biological system based on the generated stimulation signal.

21. A computer program comprising instructions which, when the program is executed  
20 by a computer, cause the computer to carry out the method of claim 20.



1/15

Fig. 1




Fig. 2

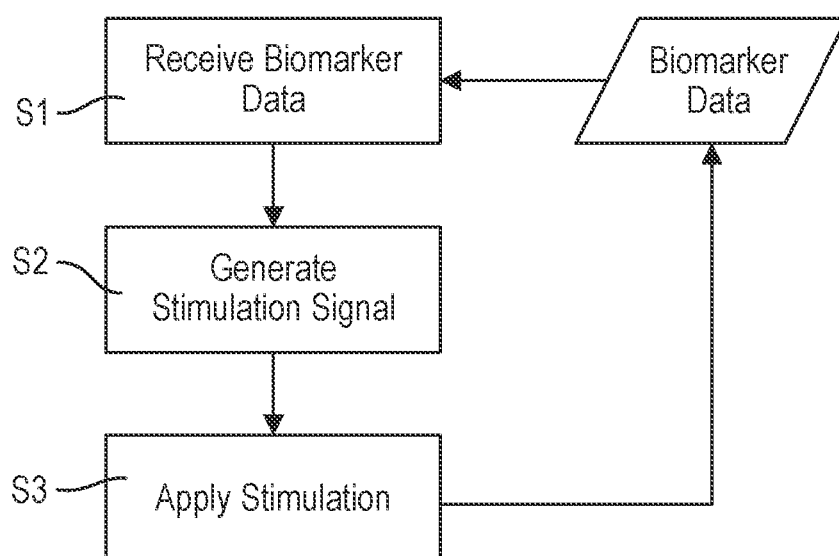


Fig. 3




Fig. 4




Fig. 5

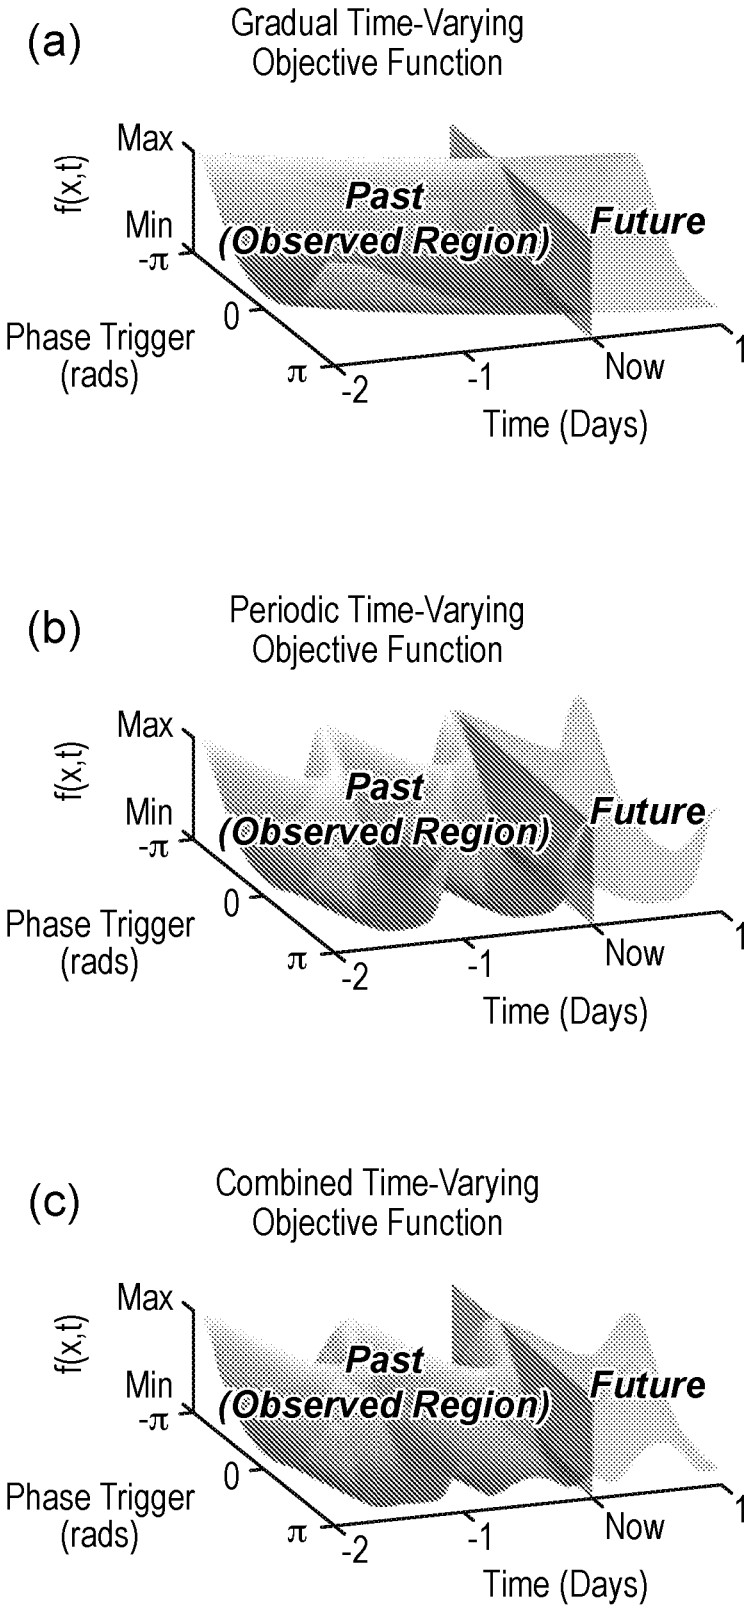




Fig. 6


(a) Forgetting Covariance Function



(b) Periodic Covariance Function




(c) Combined Covariance Function




6/15

Fig. 7


(a) Weighting of Samples in a Gradual Time-Varying Objective Function



(b) Weighting of Samples in a Periodic Time-Varying Objective Function

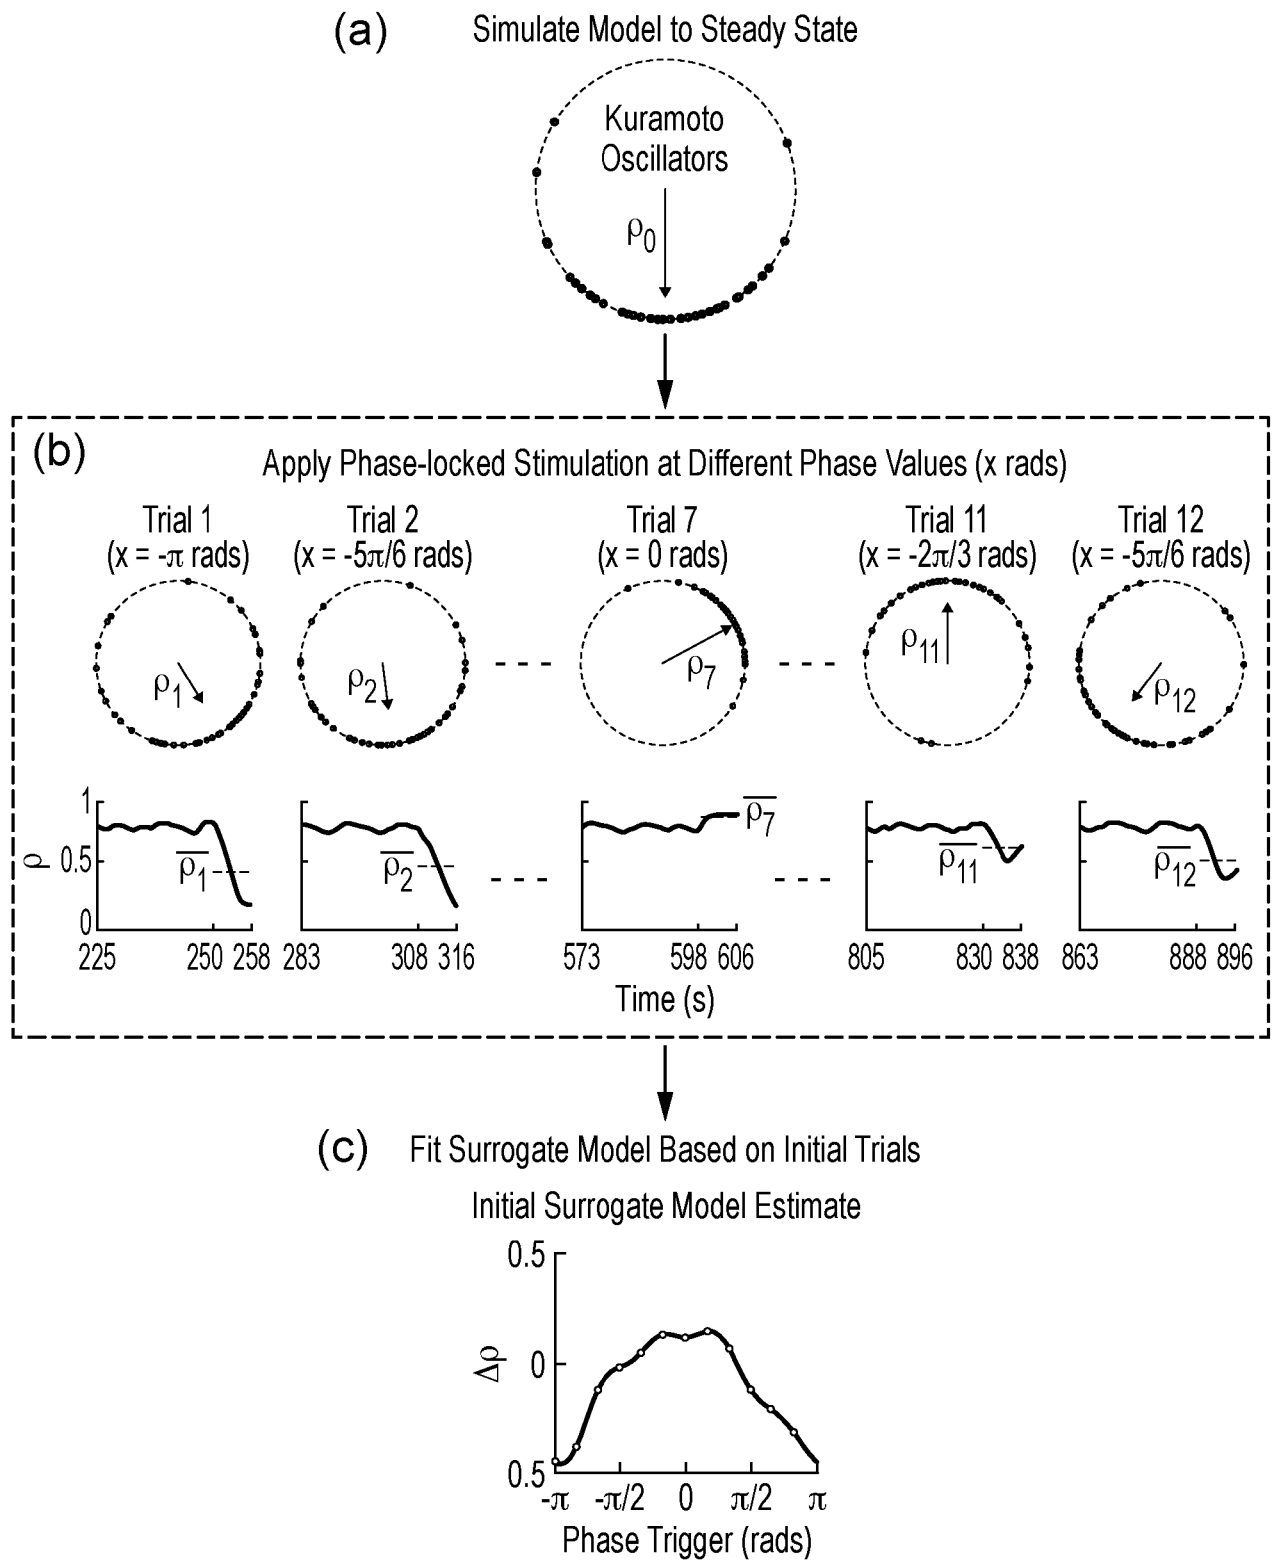



(c) Weighting of Samples in a Combined Time-Varying Objective Function



7/15

Fig. 8





## Time-Varying Offset Applied to population PRC

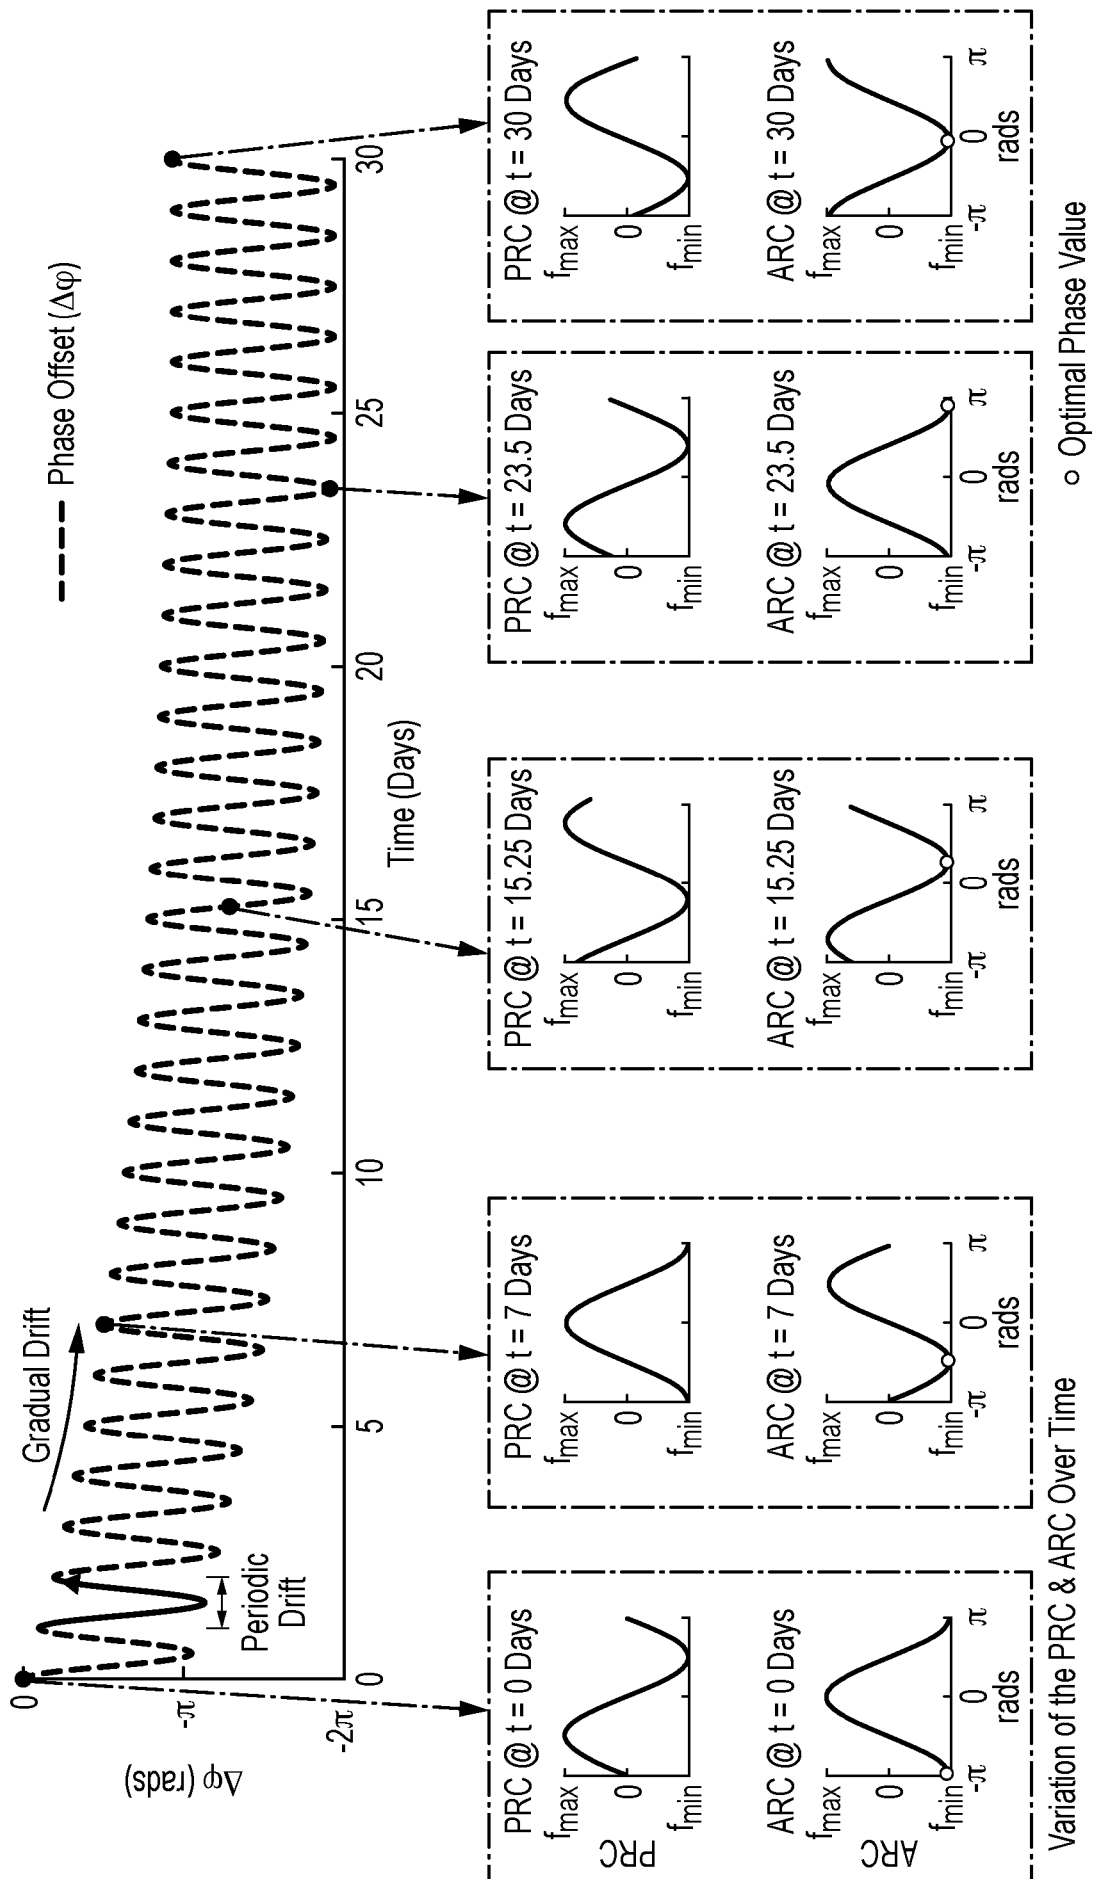




Fig. 10(a)




Fig. 10(b)




Fig. 10(c)




Fig. 10(d)





Fig. 10(e)







Fig. 11(c)




Fig. 11(d)




Fig. 11(e)




Fig. 12(a)

Algorithm Tracking Performance




Fig. 12(b)




Fig. 12(c)




Fig. 12(d)




Fig. 12(e)




Fig. 13





Fig. 14



## INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2024/051935

**A. CLASSIFICATION OF SUBJECT MATTER**

INV. A61N1/36

ADD. A61N1/05

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

**A61N**

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

**EPO-Internal, WPI Data****C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2019/140500 A1 (PNP SOLUCOES EM BIOENGENHARIA LTDA [BR]) 25 July 2019 (2019-07-25) the whole document -----	1 - 17
A	US 2020/230414 A1 (MADHAVAN RADHIKA [IN] ET AL) 23 July 2020 (2020-07-23) abstract paragraph [0036] - paragraph [0082] figures 1-13 -----	1 - 17
A	WO 2022/029445 A1 (UNIV OXFORD INNOVATION LTD [GB]) 10 February 2022 (2022-02-10) the whole document -----	1 - 17



Further documents are listed in the continuation of Box C.



See patent family annex.

\* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance;; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance;; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;" document member of the same patent family

Date of the actual completion of the international search

**17 September 2024**

Date of mailing of the international search report

**26/09/2024**

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2

NL - 2280 HV Rijswijk

Tel. (+31-70) 340-2040,

Fax: (+31-70) 340-3016

Authorized officer

**Artikis, T**



# INTERNATIONAL SEARCH REPORT

International application No.  
PCT/GB2024/051935

## Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☒ Claims Nos.: 18 - 20  
because they relate to subject matter not required to be searched by this Authority, namely:  
**Rule 39.1(iv) PCT - Method for treatment of the human or animal body by therapy due to the therapeutic steps of generating a stimulation signal and applying stimulation to a biological system based on said signal.**
2. ☒ Claims Nos.: 21  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:  
**see FURTHER INFORMATION sheet PCT/ISA/210**
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims;; it is covered by claims Nos.:

### Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

Continuation of Box II.1

Claims Nos.: 18-20

Rule 39.1(iv) PCT - Method for treatment of the human or animal body by therapy due to the therapeutic steps of generating a stimulation signal and applying stimulation to a biological system based on said signal.

-----

Continuation of Box II.2

Claims Nos.: 21

Claim 21 is drafted in such a way that it is not in compliance with the provisions of clarity of Article 6 PCT to such an extent that the subject-matter for which protection is sought cannot be established, because it refers to unsearched method claim 20 and lacks any means for interacting with the biological system.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No  
PCT/GB2024/051935

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2019140500 A1	25-07-2019	BR 102018000939 A2 WO 2019140500 A1	30-07-2019 25-07-2019
US 2020230414 A1	23-07-2020	US 2020230414 A1 WO 2020154138 A1	23-07-2020 30-07-2020
WO 2022029445 A1	10-02-2022	EP 4192573 A1 US 2023256248 A1 WO 2022029445 A1	14-06-2023 17-08-2023 10-02-2022