





PAPER

Online Bayesian optimization of vagus nerve stimulation

RECEIVED
13 March 2023REVISED
19 January 2024ACCEPTED FOR PUBLICATION
13 March 2024PUBLISHED
2 April 2024

Lorenz Wernisch^{1,*} , Tristan Edwards¹, Antonin Berthon¹, Olivier Tessier-Lariviere¹, Elvijs Sarkans¹, Myrta Stoukidi¹, Pascal Fortier-Poisson¹, Max Pinkney¹, Michael Thornton¹, Catherine Hanley¹, Susannah Lee¹, Joel Jennings¹, Ben Appleton¹, Phillip Garsed¹, Bret Patterson¹, Will Buttinger¹, Samuel Gonshaw¹, Matjaž Jakopec¹, Sudhakaran Shunmugam¹, Jorin Mamen¹, Aleksii Tukiainen¹, Guillaume Lajoie², Oliver Armitage¹  and Emil Hewage¹

¹ BIOS Health Ltd, Cambridge, United Kingdom² Université de Montréal and Mila-Quebec AI Institute, Montréal, Canada

* Author to whom any correspondence should be addressed.

E-mail: lwernisch@gmail.com**Keywords:** nerve stimulation, vagus nerve, neuromodulation therapy, Bayesian optimizationSupplementary material for this article is available [online](#)**Abstract**

Objective. In bioelectronic medicine, neuromodulation therapies induce neural signals to the brain or organs, modifying their function. Stimulation devices capable of triggering exogenous neural signals using electrical waveforms require a complex and multi-dimensional parameter space to control such waveforms. Determining the best combination of parameters (waveform optimization or dosing) for treating a particular patient's illness is therefore challenging. Comprehensive parameter searching for an optimal stimulation effect is often infeasible in a clinical setting due to the size of the parameter space. Restricting this space, however, may lead to suboptimal therapeutic results, reduced responder rates, and adverse effects. *Approach.* As an alternative to a full parameter search, we present a flexible machine learning, data acquisition, and processing framework for optimizing neural stimulation parameters, requiring as few steps as possible using Bayesian optimization. This optimization builds a model of the neural and physiological responses to stimulations, enabling it to optimize stimulation parameters and provide estimates of the accuracy of the response model. The vagus nerve (VN) innervates, among other thoracic and visceral organs, the heart, thus controlling heart rate (HR), making it an ideal candidate for demonstrating the effectiveness of our approach. *Main results.* The efficacy of our optimization approach was first evaluated on simulated neural responses, then applied to VN stimulation intraoperatively in porcine subjects. Optimization converged quickly on parameters achieving target HRs and optimizing neural B-fiber activations despite high intersubject variability. *Significance.* An optimized stimulation waveform was achieved in real time with far fewer stimulations than required by alternative optimization strategies, thus minimizing exposure to side effects. Uncertainty estimates helped avoiding stimulations outside a safe range. Our approach shows that a complex set of neural stimulation parameters can be optimized in real-time for a patient to achieve a personalized precision dosing.

1. Introduction

Bioelectronic medicine applies neuromodulation to modify neural activity and physiological function, providing a novel approach to treating a range of neurological and physiological diseases. In contrast to exploiting endogenous (or spontaneous) neural

activity, stimulation devices are able to trigger exogenous neural signals. For the stimulation of nerves, exogenous neural signals take the form of evoked compound action potentials (eCAPs), the result of a depolarization of a large number of fibers in the nerve triggered by an electrical stimulus (examples of eCAPs are shown in figure 4). The type and number

of fibers recruited depend on the wave form and dosing of the stimulus in a complex way. Through such signals, therapies utilizing stimulations can alter the function of organs, offering promising new avenues for medical treatment (Cracchiolo *et al* 2021). The exogenous signal is typically delivered by an implantable pulse generator via electrodes placed with nerve tissue.

The autonomic nervous system (ANS) is essential for maintaining homeostasis throughout the body and is implicated in the pathogenesis of a variety of disorders. The vagus nerve (VN) is a key target for neuromodulation of the ANS, as it innervates several thoracic and visceral organs, including the heart, lungs and gut. Accessing the VN through surgery is relatively straightforward (Reid 1990) and VN stimulation (VNS) has been suggested for the management of neurological diseases such as drug resistant epilepsy (Labiner and Ahern 2007a), depression (Nemeroff *et al* 2006), as well as respiratory (Carr and Undem 2003, Chang *et al* 2015), inflammatory (Pavlov and Tracey 2012) and cognitive conditions (Meyers *et al* 2018, Martin *et al* 2022), among others.

In this study, we explore the optimization of neural stimulation parameters to achieve targeted physiological changes or neural activations. Specifically, we focus on optimizing VNS parameters to elicit a precise change in heart rate (HR) or a defined level of B-fiber activation. A specific change in HR immediately after VNS is the basis of most VNS therapies for heart failure suggested to date: cyclic application of stimuli that achieve mild bradycardia have shown long term therapeutic effects in preclinical studies (Dusi and De Ferrari 2021). In particular, based on earlier studies (Ardell *et al* 2017a, Nearing *et al* 2020), stimulations that achieve a balance between tachy- and bradycardia for an overall net effect of zero change in HR are the focus of a recent clinical trial (Konstam *et al* 2019).

While the potential of VNS is widely recognized, important challenges remain in developing effective stimulation protocols. Notably, there is substantial patient-to-patient variation in the distribution and arrangement of fibers in the VN (Upadhye *et al* 2022, Jayaprakash *et al* 2023), as well as in the coupling of these fiber groups and stimulation electrodes (Qiao *et al* 2016). This necessitates a personalized optimization process. The optimization of VNS for HF therapies is primarily based on mitigating potential off-target effects prior to considering the efficacy of the resulting VNS setting (Labiner and Ahern 2007a, Nicolai *et al* 2020). Besides patient to patient variation in the effect of VNS parameters, changes in the same patient over time have been observed as well (for example, over a ten week titration period in Nearing *et al* (2020)). Extensive adjustments over a large number of stimulation parameters are time consuming and often induce discomfort or even more severe side effects. The development of rapid and flexible

optimization procedures which can quickly and efficiently identify stimulation parameters with a specific response is therefore of great interest.

Personalizing neural stimulation presents a multi-faceted challenge, necessitating specialized hardware and software solutions. The system utilized in this study encompasses a neural interface that captures high-quality, longitudinal neural and physiological data, as well as a cloud infrastructure tailored for real-time collection and processing of these signals. The focus of this work, however, is a method for optimizing stimulation parameters, in particular, Bayesian optimization (BO). We refer to the specific version of BO tailored for neural stimulation in this study as online BO of evoked signals (OBOES).

Parameter optimization in neuromodulation has been successfully applied across various domains. One notable instance is adaptive deep brain stimulation, which exemplifies the closed-loop dynamic control of brain activity via continuous adjustment of stimulation parameters (Arlotti *et al* 2018, Guidetti *et al* 2021). In contrast, our focus is on optimizing stimulations for a system that is, at least temporarily, non-dynamic. BO is a computational technique that seeks to identify optimal parameters with a minimum number of exploratory steps. BO algorithms have recently emerged as promising optimization tools for stimulations of the nervous system in several studies: for optimizing epidural spinal cord stimulation in rats (Desautels *et al* 2015), for intracortical stimulation in monkeys (Laferriere *et al* 2020), for distal limb movement in monkeys (Losanno *et al* 2021, Bonizzato *et al* 2023), treatment of epilepsy (Park *et al* 2020, Stieve *et al* 2023), and improved memory (Ashmaig *et al* 2018) in deep brain stimulation. Simulation studies, utilizing models based on neuromodulation data, demonstrate the potential of BO in comparison to other optimization methods (Grado *et al* 2018, Mao *et al* 2023, Aiello *et al* 2023).

The core principles of a BO approach involve identifying and optimizing an individual's response to stimulation inputs at a specific point in time, without relying on past responses from the same individual or from others. Given the absence of prior knowledge about the VNS response function under these conditions, exploration becomes essential. At the same time, steps need to be taken toward optimizing the response. Thus, a balance must be struck between exploration and optimization. Finally, a statistical assessment of the extent of exploration and of the confidence range on the discovered optimum is desirable. Knowledge of the probability that responses will likely fall outside a safety or comfort range helps managing the range of acceptable dosing.

We develop and tested OBOES by optimizing responses to stimulations controlling HR and eliciting optimized activations of certain nerve fiber types. To develop, compare and evaluate optimization

approaches *offline* (or *in silico*) before their real-time experimental deployment, we developed realistic simulation models of typical neural responses to VNS using experimental data. These simulation models are data-driven with minimal modeling assumptions about any underlying mechanisms, thus providing a realistic evaluation platform for assessing the effectiveness of optimization approaches for various VNS scenarios. It should be noted that the simulation models developed in this study were not designed to precisely replicate responses to neural stimulation. Instead, their purpose was to offer a reasonably realistic representation of these responses to facilitate the development and thorough evaluation of OBOES, ensuring it was robust and effective prior to its application in actual experimental settings. The effectiveness of a neural stimulation optimization using a Bayesian method approach in an *online* (or *in vivo*) setting was explored in a series of VNS experiments conducted on anesthetized swine. Since OBOES is able to optimize arbitrary response functions, the method is applicable to human patients as well, particularly given the functional similarities of the human and porcine VN (Settell et al 2020).

In addition to the optimization of changes in HR, we performed optimization of neural biomarkers, such as the activation of specific fiber types. To this end, we recorded exogenous nerve signals induced by VNS, in the form of eCAPs. It has been widely hypothesized that B-fiber activation is associated with bradycardia (Qing et al 2018a). Fiber activation in the exogenous signal can be optimized with minimal physiological effects using stimulations with low frequency and short duration. Optimizing fiber activation instead of directly targeting a change in HR is therefore less stressful for the subject and less time consuming. As an example of the applicability of OBOES to optimize specific neural signals in an online experiment, we maximized B-fiber activation while minimizing stimulation current.

In summary, we have demonstrated the practical applicability of a machine learning (ML) optimization algorithm for VNS in a series of preclinical trials. This system achieves, in real time, target changes in HR or maximization of B-fiber activation while minimizing stimulation current. Consequently, the system facilitates the optimization of multi-dimensional VNS parameters, thus enabling a personalized precision dosing with a minimum number of stimulations causing side effects.

2. Methods

We describe the surgical setup and technical infrastructure for recording neurograms of eCAPs in the VN and HR changes with VNS. To assess the performance of our optimization approach on neural data we built a series of simulators. Since the 'true'

response function is known for simulators, optimization algorithms can be compared and assessed and the efficiency of an optimization can be evaluated *offline*. Simulators for eCAPs were constructed based on VNS responses recorded from our subjects, termed the *porcine dataset*. A further set of simulators is based on data available in the literature, termed the *rodent dataset*. Finally, we performed *online* optimizations of HRs and neural fiber activations intra-surgically and describe a statistical method for assessing the uncertainty estimates of OBOES when the true response function is unknown.

2.1. Anesthesia and surgical preparation

Experiments were conducted on five female Landrace pigs (labeled A7 to A11) under the approval of the local Institute of Animal Care and Use Committee, and approved in accordance with the animal use policy of BIOS Health Ltd. All animals, weighing 43.2 ± 3.5 kg, were sedated using an intramuscular injection of ketamine (25 mg kg^{-1}), atropine (0.04 mg kg^{-1}) and acepromazine (1.1 mg kg^{-1}), then intubated and anesthetized using propofol (3 ml slow intravenous bolus for induction and $0.4 \text{ mg kg}^{-1} \text{ min}^{-1}$ for maintenance). The depth of anesthesia was monitored and adjusted based on the occipital reflex, jaw tone and hemodynamic indices. A mechanical respirator was available, although animals were left to breathe spontaneously unless end tidal carbon dioxide concentration exceeded 70 mmHg. A two-lead electrocardiography (ECG) was used in a lead I configuration. Data were digitized at a rate of 500 Hz using the Axon Digidata 1550B plus HumSilencer (Molecular Devices, CA, USA).

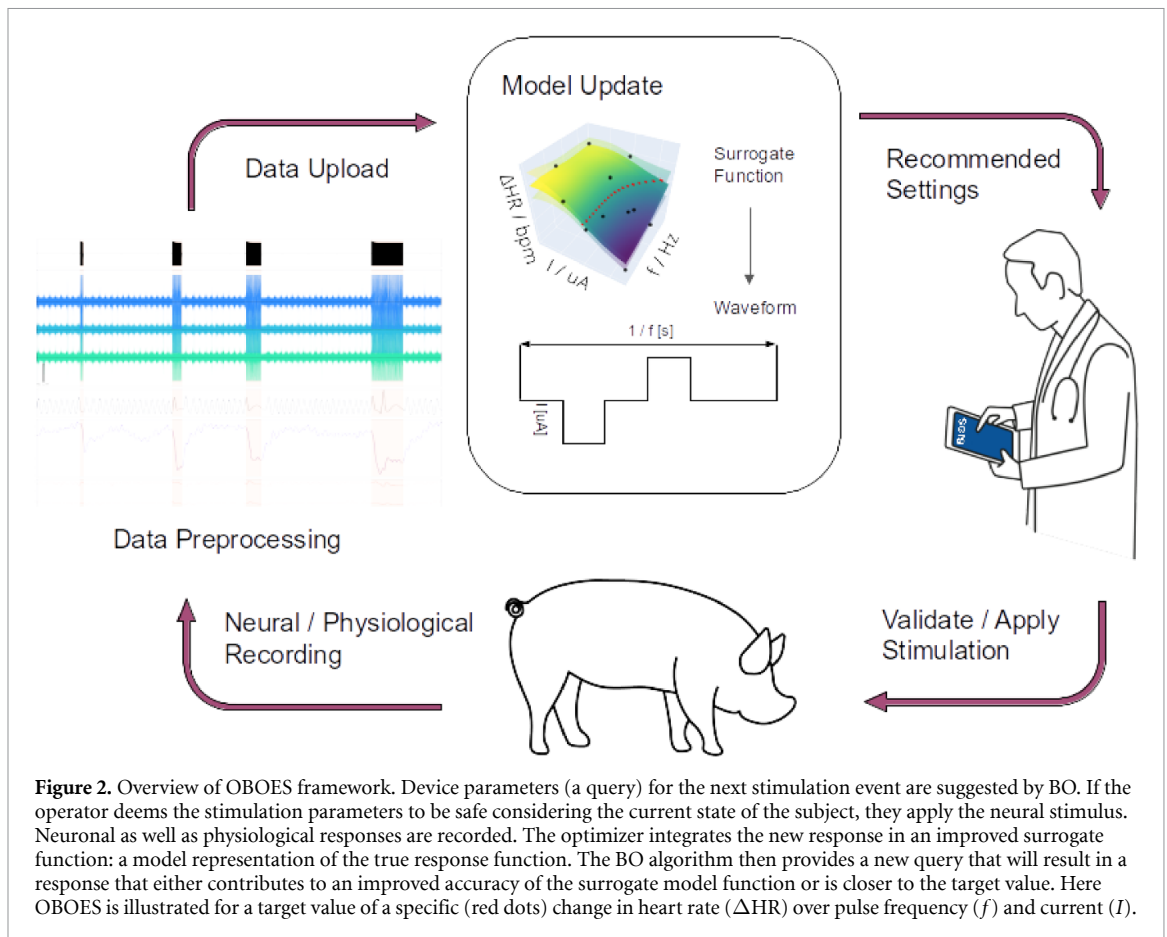
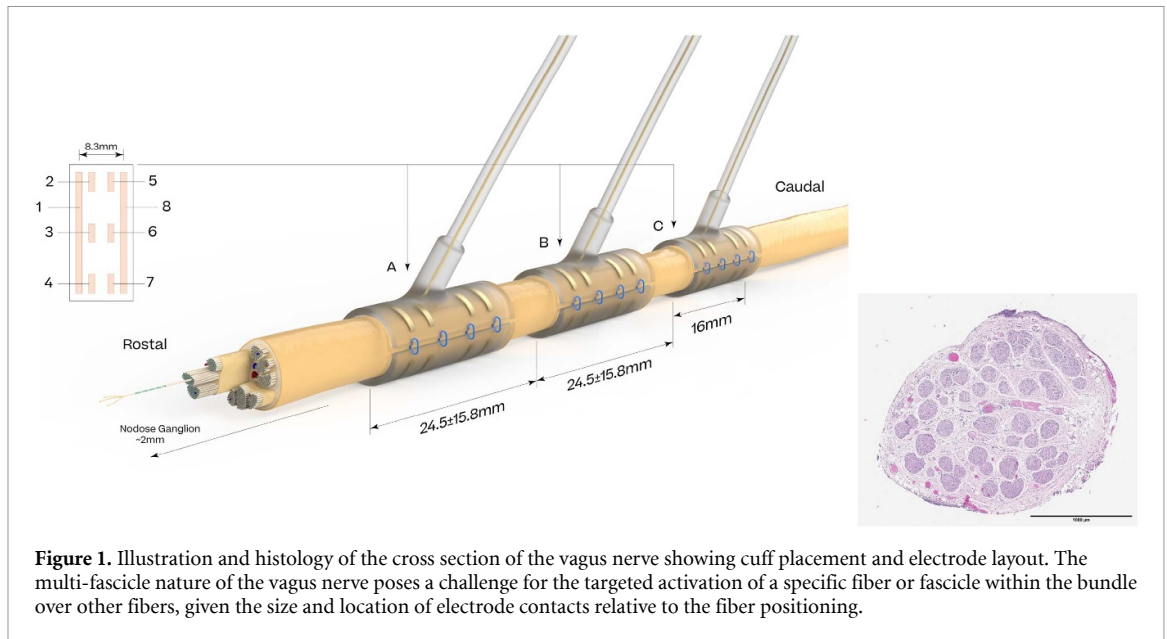
Animals were positioned supine, with both forelimbs and head extended to expose the ventral aspect of the neck. A 10 cm incision was marked 2 cm right of the midline. The nodose ganglion marks the most cranial point of our interaction with the vagus. Eight centimeters of the nerve was stripped and cleaned caudally of the nodose to fit three cuffs.

The surgical procedures each ran for 8 h at which point the animals were euthanized with an intravenous administration of T-61 (0.3 ml kg^{-1}).

2.2. Electrodes and equipment

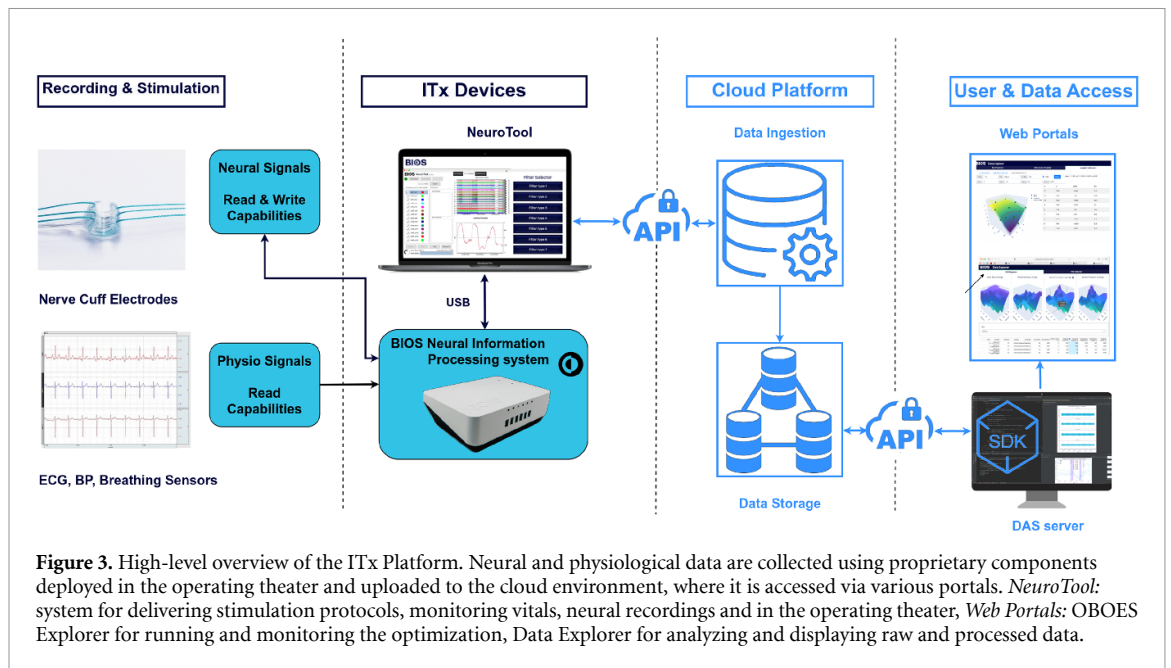
The multi-contact cuffs, built by Microprobes (NC, USA), each contained eight electrodes arranged in longitudinal bipolar pairs (see figure 1). All neural recordings and stimulation data presented in this paper were performed via the maximally spaced bipolar pair (electrodes 1 and 8 in figure 1), with a center-to-center separation of 8.3 mm. The cuff arrangement gave an average inter-cuff distance of 24.5 ± 15.8 mm between the stimulation cuff and the most cranial recording cuff across the cohort.

All cuffs were connected to a custom data acquisition system engineered in-house. Briefly,



an RHD2216 differential amplifier and RHS2116 stimulator/amplifier digital electrophysiology chips (Intan Technologies) are controlled through custom HDL on an Artix 7 (Xilinx, CA, USA) field-programmable gate array (FPGA) sampling up to 16 channels at 30 kHz. The system is capable of applying stimulation pulses with varying current (10–2500 μA), frequency (1–1500 Hz), individual

pulse width (1–1000 μs), and train duration (1–10 s). Biphasic symmetric constant-current pulses were delivered through a single channel in reference to a ground lead positioned proximal to the stimulation site. A spare pair of record channels were used as an additional ECG recording to ease synchronization of the neural and physiological data and to obtain HR data during online optimization.



2.3. Research and development system architecture

We developed a software suite to accelerate our research and development (R&D) process with the goal of coupling data generation with cloud-based data management and processing resources to support ML research in neuroscience. A schematic of the platform is shown in figure 3.

The data acquisition components of the system reside in the operating theater: the data acquisition system (neural interface) and user interface (UI, *NeuroTool*). The data acquisition system performs the VNS and VN signal recording, uploads data to the cloud services and has an onboard GPU for edge ML inference. Stimulation control models are deployed here for applications where rapid inference and data security are paramount. Data can be temporarily streamed to the UI for validation and checking of impedances. All stimulation parameters can be programmed, adjusted and triggered by an onsite electrophysiologist.

Data are uploaded to a data ingestion service in batches. From here, the neural responses of interest can be immediately visualized and examined on *Data Explorer* or on *OBOES Explorer* (see figure 3; snapshots of the explorer screens are shown in the supplementary material). Monitoring of responses informs fiber activation thresholds to guide other stages in the protocol. Data are stored on a remote server and can be downloaded for post surgical analysis and algorithm development.

2.4. Rodent dataset

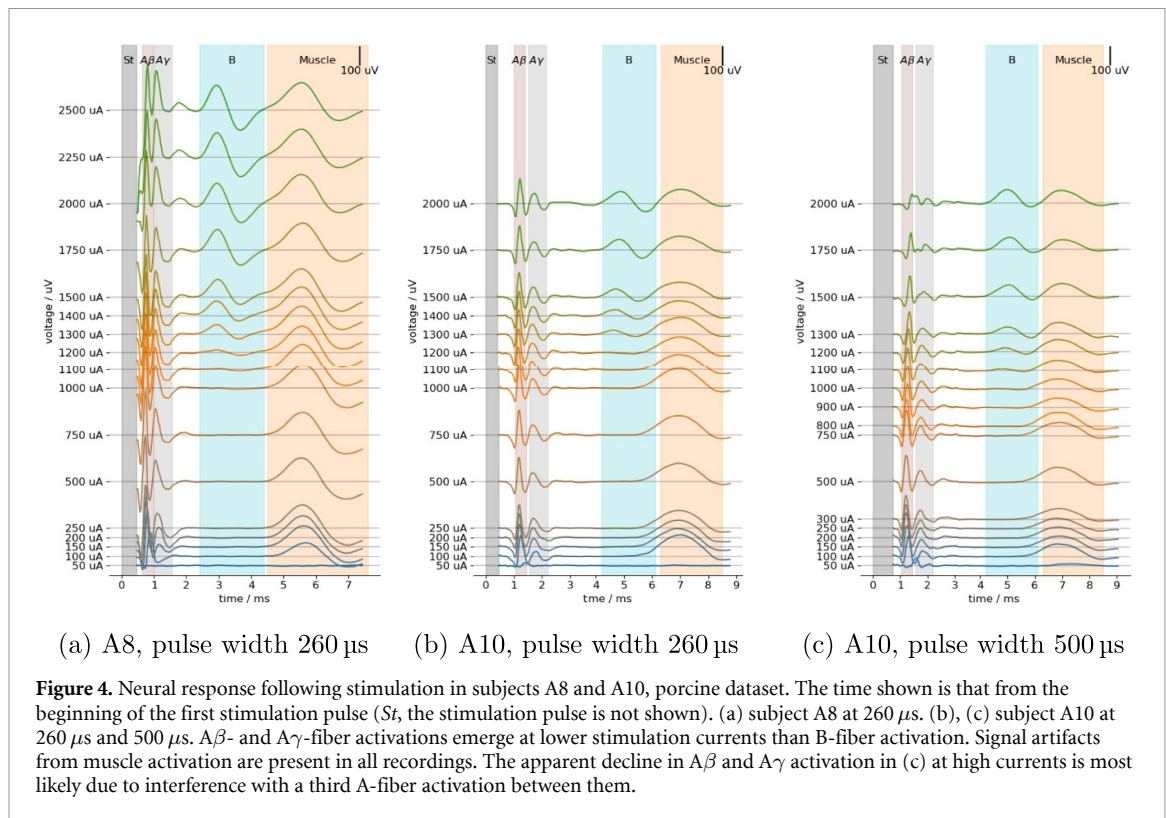
For the development of the OBOES framework, we devised a simulation platform for testing, refining and evaluating optimization approaches on simulated data. Such simulations should show characteristics

as close as possible to realistic neural or physiological responses to stimulations. We utilized data on neurograms of eCAPs by Ward *et al* (2015), available through the National Institutes of Health SPARC programme (Ward *et al* 2021). The dataset provides the responses to a grid of VNS parameter values varying over seven pulse currents and four pulse widths. Each stimulation was applied cathode-first with an alternating monophasic constant current. Twelve rodents received the complete sweep of 28 VNS parameter combinations at a fixed pulse frequency of 10 Hz. Here we refer to the rodents as subjects P1 to P12. The maximum eCAP values provided in the study form the basis of our simulator construction.

2.5. Porcine dataset

We extended the simulation platform using data from our own experiments. Exogenous signals in the form of eCAP responses to stimulation were collected for four porcine subjects A7 to A11. The pulse width was set to either 130 μs , 260 μs , or 500 μs and the pulse train duration to 5 s. A range of frequencies (2, 5, 10, 15, 20 Hz) and currents (50–2500 μA , in increments of 100 μA) were applied. We found an alternating, monophasic, constant pulse waveform to be the most suitable to elicit clean eCAP responses without amplifier artifacts.

Evoked neural and cardiovascular responses to VNS were recorded and vitals were allowed to return to baseline before starting the next cycle (with a minimum of 30 s between each cycle). Charge limits were set at the point where stimuli caused noticeable strong side effects, such as coughs or extended bradypnea. Neurograms were processed as described next.



2.6. Processing of neural data

Stimulations are administered as a train of consecutive pulses at a specific frequency. We observed no difference in the shape of eCAPs elicited by individual pulses, provided the pulses maintained identical width and current amplitude. So we only considered the eCAP response to the initial pulse for subsequent analysis. The neural recordings sometimes exhibited stimulus artifacts post-pulse, attributable to some coupling between stimulating and recording electrode and amplifier saturation. We rectified this by fitting robust polynomial regressions of low degree to the time series (Seabold and Perktold 2010) and subtracting the regression values. This is followed by convolution smoothing using a Kaiser window of increasing size, in order to retain fine details at the beginning of the signal for faster, less dispersed eCAPs, while providing enough smoothing toward the end of the signal for slower, more dispersed eCAPs. As an example, the evolution of the fiber engagement for two subjects, recorded from the most cranial cuff, are displayed in figure 4. Additional figures of eCAPs for all subjects can be found in the supplementary material.

Activations of three fiber types are visible in the neurograms of figure 4, denoted A β -, A γ -, B-fibers followed by a laryngeal muscle artifact. The muscle artifact was validated through caudal vagotomy as in Nicolai *et al* (2020). We noticed that the artifact appears instantaneously on both cuffs without any conduction delay in contrast to the other eCAPs. fiber activation thresholds differ from subject to subject.

This is likely due to differences in bioelectronic coupling at the cuff-nerve interface and anatomical variation in fiber distribution (Losanno *et al* 2021).

For optimizing B-fiber activation, activation was calculated from neurograms by identifying the B-fiber maximum and minimum peak and taking their difference.

2.7. Gaussian processes

Gaussian processes (GPs) (Rasmussen 2006) provide a flexible, nonlinear approximation of an unknown response function. The only way to gather information about this response function is to probe it with specific input values and to obtain corresponding response values. This is often referred to as a *black box* function, which is only accessible via such input queries. With only a finite number of inputs and corresponding response values, limited information about the underlying function is available, and the approximation might not be very accurate. If subsequently more data become available, it is possible to integrate them in a principled probabilistic fashion to improve the GP approximation. Here we use GP approximations for two purposes: (a) constructing simulation models for method development and testing, described in section 2.8, and (b) serving as a surrogate function in BO, as outlined in section 2.9.

Technical details of how a GP approximation is set up can be found in the supplementary materials. Essentially, a GP provides an approximation to a response function g that maps a d -dimensional input $x \in \mathbb{R}^d$ to a scalar response $y, y = g(x)$. A key feature

of a GP is the ability to provide an estimate of the precision of its approximation in the form of uncertainty regions above and below its approximation. The accuracy of the approximation improves with an increasing number of data points of input-response pairs. The GP is represented by a mean function $m(x)$ for any input x and an uncertainty estimate $v(x)$ at this point. The smaller $v(x)$ the more accurate the approximation. In this study GP approximations are shown as surfaces (of mean response) with upper and lower confidence bounds indicated as transparent surfaces when the uncertainty around the GP mean response is of interest. They consist of the surfaces for $m(x) \pm \sqrt{v(x)}$ enclosing the central 68.2% quantile of a normal distribution. For examples showing of all three surfaces see figure 8.

Fitting a GP to data (x, y) means choosing suitable kernel parameters that represent the variance of the signal, the multi dimensional length scales (smoothness), and the noise variance. Parameters can either be set manually or estimated by maximum *a posteriori* probability with or without a prior such as a gamma distribution (some details on priors used in this study can be found in the supplementary material).

2.8. Exemplary simulators for the rodent and porcine dataset

In order to develop and assess optimization approaches offline, we developed simulators based on the rodent as well as on our porcine dataset. Our primary intent with simulations was to ensure the readiness of our real-time optimization algorithm before deploying it in experiments on live animals. We viewed simulations as a critical debugging and validation step where: (i) the complete algorithm and pipeline could be tested end-to-end. (ii) runtimes could be assessed to ensure they were not excessively long. (iii) GP priors, such as noise and length scales, and optimization parameters like the balance factor between exploration and optimization or the number and location of starting points (as discussed in a following section), were tuned appropriately. (iv) Simulations gave us a ballpark estimate of the number of BO steps needed before our results could outperform other search methods, which was instrumental in adjusting our experimental schedule during surgery sessions. (v) We were able to introduce variability in form of controlled levels of randomness and noise into the simulated data.

While we initiated our simulations with standard functions commonly used in the BO literature, we soon discerned that the real-world response functions could exhibit specific features not present in these standard functions. This led us to derive functional relationships directly from actual data, which has two main facets: (i) the base response should authentically mirror the characteristics—like shape, range, number of local minima and maxima—typically found in real response functions based on physiological or

eCAP features; (ii) in practical experiments, the measured response invariably contains random fluctuations and measurement noise. GP approximations are aptly suited to address both these points.

Based on actual experimental input-response data, we employ GP regression to fit a (mean) response surface to the data. GP regression is well suited to deal with irregular, high-dimensional input points and noisy responses. Once a response surface is fitted, for the purpose of a simulation, we are free to designate the fitted surface as the ‘true’ response, serving as a benchmark for evaluating optimization algorithms. Notice, however, that such simulated data are typically too smooth, since disturbances or measurement noise are not included. In real-world scenarios we typically are unable to obtain clean values of the true responses without distortion by random exogenous factors and measurement noise. To reproduce such characteristics of actual data, our simulator adds Gaussian random noise to the smooth response at a noise-to-signal ratio of 0.2 for each output request. This reflects a noise intensity equivalent to about 20% of the signal—a parameter observed across the majority of our data sets.

For the rodent dataset we constructed GP simulators for maximal eCAP activation for each of the 12 subjects P1 to P12. A GP approximation of maximal eCAP action in subject P10, for example, is shown in figure 6(c). The original data provide 28 data points (shown in blue in figure 6(c)). A GP approximation provides an output for any arbitrary input within a given search space.

For the porcine dataset of subjects A7 to A10 we constructed GP simulators that include the time dimension as well to represent the neural signal as time series, that is, a three dimensional input point of current, width, and time is mapped to the corresponding electrode voltage value at that time (for an example see figure 5). For further details of GP parameter settings see the supplementary material. For an input of stimulation parameters this type of simulator returns a time-series vector, a neurogram of eCAPs. In order to obtain a scalar response representing B-fiber activation, the activation value was computed as difference of the maximum in the first half to the minimum in the second half of the B-fiber range defined as 2.47–4.3 ms, 2.3–4.97 ms, 2.97–5.97 ms, and 3.97–6.3 ms for A7 to A10.

2.9. BO

BO (Archetti and Candelieri 2019) is a popular choice for the optimization of objectives over unknown functions. BO is a sequential search strategy optimizing the input to an unknown response function balancing optimization with exploration (Mockus (2012), for an accessible introduction see also BorealisAI (2020)). Typically the response function to be optimized over is unknown. However, it can be probed at any arbitrary input point (possibly with certain

boundaries). In the context of a BO the mechanism producing responses to inputs is referred to as the *oracle* or *black box* function. In what follows, we will use BO in two contexts: when the oracle is provided *offline* by a simulator (a GP fitted on a full data set), and when the oracle is provided *online* by a real-time measurement of the physiological or neural response following VNS in an alive subject. For details, see sections 2.11 and 2.12, respectively.

Since the optimizer has no access to the true response function and can only query it at a few input points through the oracles, it builds a *surrogate function* which it maintains and improves internally to guide its search. Sequential querying of the oracle is used to improve the surrogate function. The surrogate function therefore needs to be flexible enough to adapt to new data. It should also provide some estimation of uncertainty in order to balance optimization with exploration. This is required in order to detect regions of high uncertainty that might require more data to narrow the uncertainty gap.

2.10. BO with GPs

GPs are a popular choice for representing surrogate functions in BOs. As discussed in section 2.7, GPs are probabilistic models that are able to capture functional relationships between multidimensional inputs and a nonlinear responses. Starting from a typically flat representation of the function before any data are available, a GP provides an increasingly accurate estimation of the unknown response function by incorporating an growing number of data points.

A BO run consists of a series of requests in the form of queries to the oracle. To find the next query input to forward to the oracle, a BO maintains an *acquisition function* in addition to its surrogate function. The next query input is an input that optimizes the acquisition function. The acquisition function needs to support finding an optimum of the (unknown) objective function based on the BO's surrogate function. In addition, it also needs to support exploration of new regions for an improved estimation of the surrogate function. To achieve these complementary aims an acquisition function is typically some combination of the mean of the surrogate function with its uncertainty. Once the acquisition function is optimized and the corresponding input is forwarded as query to the oracle, the oracle returns the response to that input. The acquisition function is designed so that this response either provides an improved optimum, or reduces the uncertainty about the response function in an under-explored region. The new data point contributes to an improved estimation of the objective function by the surrogate function.

The upper (lower) confidence bound UCB (LCB) is a popular acquisition function if the goal is to maximize (minimize) the objective. As in section 2.7 a GP is given in form of its mean function $m(x)$

and the uncertainty variance $v(x)$ or standard deviation $s(x) = \sqrt{v(x)}$ at input x . LCB and UCB are then defined as $lcb(x) = m(x) - \lambda s(x)$ and $ucb(x) = m(x) + \lambda s(x)$ (λ a scaling factor that trades off exploration with exploitation; we call it the *confidence bound factor*).

Since we often not necessarily aim to maximize or minimize a response, but to get close to a specific set point, we propose a novel acquisition function. The *targeted confidence bound* method or TCB, is a variation of the UCB and LCB acquisition function and defined as $tcb(x) = |m(x) - t| - \lambda s(x)$. An optimal query point is found as $x^* = \operatorname{argmin}_x tcb(x)$. A justification for the use of this acquisition function to find an optimal point achieving a target set point is provided in the supplementary material.

2.11. Offline BO: assessment using simulations

Before deploying OBOES in a live experiment, we assessed the feasibility, practicability, and the convergence rate to be expected for typical neural response functions. We used the simulators from section 2.8 to provide challenging objective functions from realistic neural data sets. With a simulator, we have access to the true objective function and can therefore provide the 'ground truth' optimum. This optimum can be used as a yardstick to assess how quickly an optimization algorithm converges to it in terms of numbers of required queries.

We expect the number of queries required by an efficient optimization algorithm for reaching an optimum to be small. Using a simulator, an evaluation of BO in terms of the required number of iterations was done as follows. First we obtained the true optimal value from the simulator. Then, starting with two initial queries at the extrema of lowest and highest values (for all input dimensions) within the search space, we tracked the progress of BO iterations from query to query. At each iteration we obtained the BO's suggestion for the next query point (which optimizes its acquisition function) as input to the oracle in the next iteration. For the purpose of assessment, however, we also obtained the BO's best current guess of an input that might elicit an optimal response from the oracle (which optimizes its surrogate function). Since the best current guess is obtained by optimizing the surrogate function, the estimated optimal input typically differs from the query input, which is obtained by optimizing the acquisition function. For the assessment, the algorithm's estimated optimal input is evaluated by the simulator. The value returned for that input is compared to the true optimum (which was obtained from the simulator in the beginning). The relative error is the absolute difference between guessed and true optimum divided by the true optimum.

The most straightforward alternative to a BO search is a *grid search* with a fixed number of input points laid out in a grid. Notice that optimization

using a grid requires a large amount of oracle evaluations before an optimum can be called, since all grid points need to be evaluated. The relative error of this optimum serves as a threshold any competitive optimization method should be able to undercut in fewer iterations than the size of the grid. A *random search* evaluates random input points sequentially through the oracle and keeps track of the best result along the sequence. For examples of a comparison of all three search strategies see figure 7. We expect that the relative error of a BO is smaller than that of a random search for the same number of iterations.

2.12. Online BO: modulating HR change and B-fiber activation

Two types of online experiments were conducted using OBOES. The first experiment aims to optimize VNS to meet a target modulated HR change. The second aims to optimize a specific vagal fiber recruitment, as measured by recorded eCAPs.

Online experiments for HR were conducted in subjects A9, A10, A11 and for B-fiber optimization in subject A12. A safe search space for pulse frequency and current was established based on previous stimulation experiments and the judgment of an onsite electrophysiologist. We initialized the BO algorithm with two fixed stimulation settings: one at lowest frequency and current setting, one at the highest within the search space. The optimization cycle proceeds as shown in figure 2. To achieve a target HR, a target value for ΔHR was chosen: -5 bpm for A9 and A11, and -10 bpm for A10. After initialization, the BO model suggests a query input, consisting of a current ($I/\mu\text{A}^{-1}$) and a frequency (f/Hz^{-1}). A stimulation with these query settings is applied. The cardiovascular and vagus response for this stimulation are uploaded via the Data Ingestion Service serving OBOES Explorer (figure 3). The HRs on either side of the stimuli are estimated from ECG data (for details see supplementary material) and their difference ΔHR returned to the BO algorithm as response to the query.

Similarly, OBOES was set up for maximization of B-fiber activation as recorded by cuff electrodes. Neurograms were recorded during stimulation and processed as described in section 2.6. B-fiber activation was defined as the difference between maximum and minimum value of the recording in the range of 3.5–5.2 ms. In order to mitigate side effects due to strong stimulation currents, the objective function was defined as the ratio of B-fiber activation (in mV) and stimulation current (in μA). This objective function encourages high levels of B-fiber activation and simultaneously low levels of stimulation current.

There is no immediate ground truth for assessing an online BO. However, a targeted BO aims for a specific target response (for example see red points in figure 6(b)). Instead of a single input that leads

to a maximum or minimum response the BO typically provides a range of inputs along the iso-response curve with responses close to the target value. We can therefore choose from a range of input values to assess the BO. Ideally we would like to demonstrate the accuracy of the BO's surrogate function by establishing that the true response value falls within the uncertainty range of one standard deviation above and below the mean of the surrogate function GP of the BO (see section 2.7). However, the experimental response to VNS is affected by noise, and a single response value is not necessarily representative of the true response. We therefore acquired three responses $r_1(x), r_2(x), r_3(x)$ for the same input x and used them to estimate the true response value $\rho(x)$ at that input by Bayesian inference. Instead of a single value, such inference results in a probability distribution for $\rho(x)$ (for details see supplementary material). We consider the accuracy of the surrogate function as sufficient if 95% of this probability mass lies within the BO's uncertainty range. This uncertainty range is then shown to be a reliable measure of how close the BO's surrogate function is to the real response, at least around the iso-response curve. We repeated this calculation for three different input points along this curve. For an illustration of this approach and an example of the posterior distribution see figure 9(b).

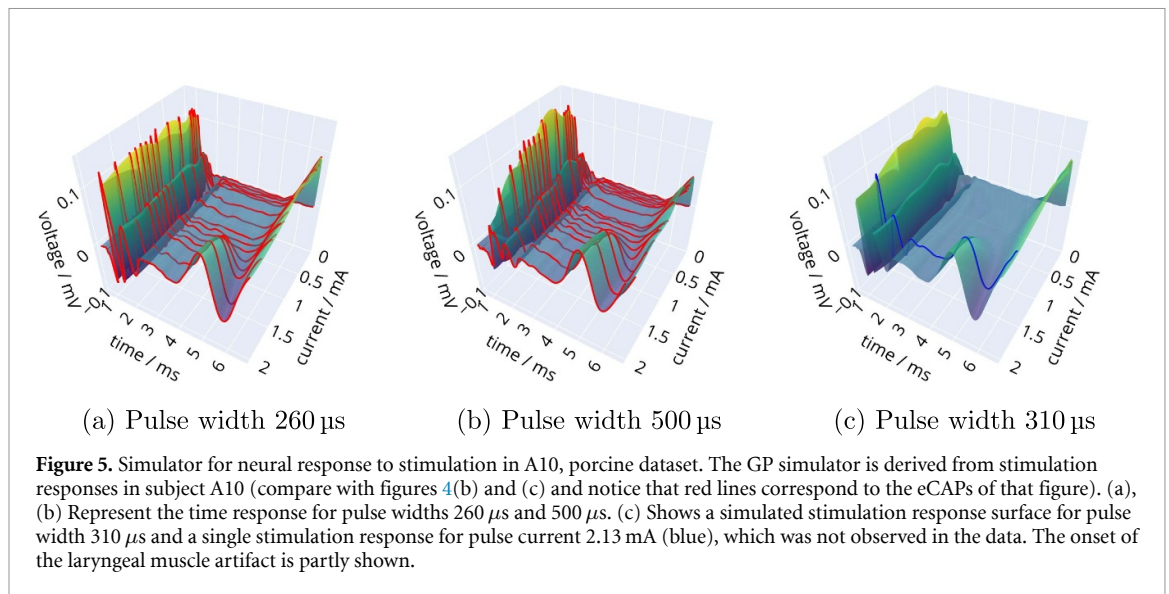
For maximizing the B-fiber response over current, we obtain a single optimal input instead of an iso-response curve of possible inputs. So instead of testing the accuracy of the surrogate function and the uncertainty estimate at three different inputs along an iso-response curve, we just tested the accuracy of the surrogate function at the optimal input.

3. Results

In order to assess OBOES offline, realistic simulators were constructed from neural data as reported in the following section. Subsequently, the results of evaluating the approach offline are discussed, followed by sections detailing the results from applying OBOES online in an intra-operative setting with alive animals, controlling HR and B-fiber activation.

3.1. eCAPs simulators for porcine dataset

For assessing the optimization of neural fiber activations through VNS we built simulators emulating eCAP responses. eCAPs for four different subjects A7 to A10 were obtained for this study. Figure 4 shows eCAPs following VNS for subject A8 and A10 for pulse widths of 260 μs and 500 μs . Figure 5 shows the result of fitting a GP to a three dimensional input space of current, width, and time to provide simulated voltage reading at this point for A10. Figures 5(a) and (b) show the fit to the two different levels of pulse widths available in the data set. However, notice that the simulator can provide eCAPs for arbitrary inputs: figure 5(c) shows a simulated eCAP response for pulse



width of 310 μs and current of 2.13 mA, neither of which was covered by the original data. Figures of eCAP simulators for A7, A8, and A9 are provided in the supplementary material.

3.2. Offline assessment of BO using simulations

Parameter optimization for achieving a specific target response was assessed by applying targeted BO to B-fiber activation as derived from eCAP simulators. One example of a B-fiber activation response is shown in figure 6, which is obtained by measuring B-fiber activation in neurograms provided by the eCAPs simulator for A10 (figure 5) (similarly for A7–A9 in the supplemental material). A targeted BO was set up to find inputs that yield responses close to a target value of 0.04 mV for A10 (0.06, 0.15, 0.1 mV for A7 to A9). BO was performed with a prior factor of $\kappa = 0.8$ and confidence bound factor of $\lambda = 2.0$. The approximation achieved by surrogate function of the BO after 11 evaluations (including two initial input points at the extreme of the search space in magenta) is shown in figure 6(b). The iso-response curve is indicated by red markers. Notice how closely the BO's surrogate approximation in figure 6(b) follows the true response in figure 6(a) after only a few stimulation results. A small amount of uncertainty is indicated by upper and lower surfaces close to the surrogate approximation.

Similarly, GP simulators were constructed for maximal eCAP activation in 12 subjects, P1 to P12, from the rodent dataset. Figure 6(c) shows a simulator for subject P10, while figure 6(d) shows the result of optimizing pulse current and width to achieve maximal overall eCAP activation (examples for all 12 subjects can be found in the supplementary material). Each BO was started by two initialization points (in magenta) at the low and the high value corners of the search spaces. The maximum suggested by the BO is marked by a point in red. Uncertainty estimates

around the reconstructed response are indicated by transparent surfaces. They are narrower in the region of interest around the maximum response where BO tends to provide more queries. BO was performed with a prior factor $\kappa = 0.8$ and confidence bound factor $\lambda = 4.0$. Again notice the similarity between the BO's reconstructed surrogate function in figure 6(d) compared to the true response in figure 6(c), as well as the BO's ability to find the correct maximum for a surface with several local maxima.

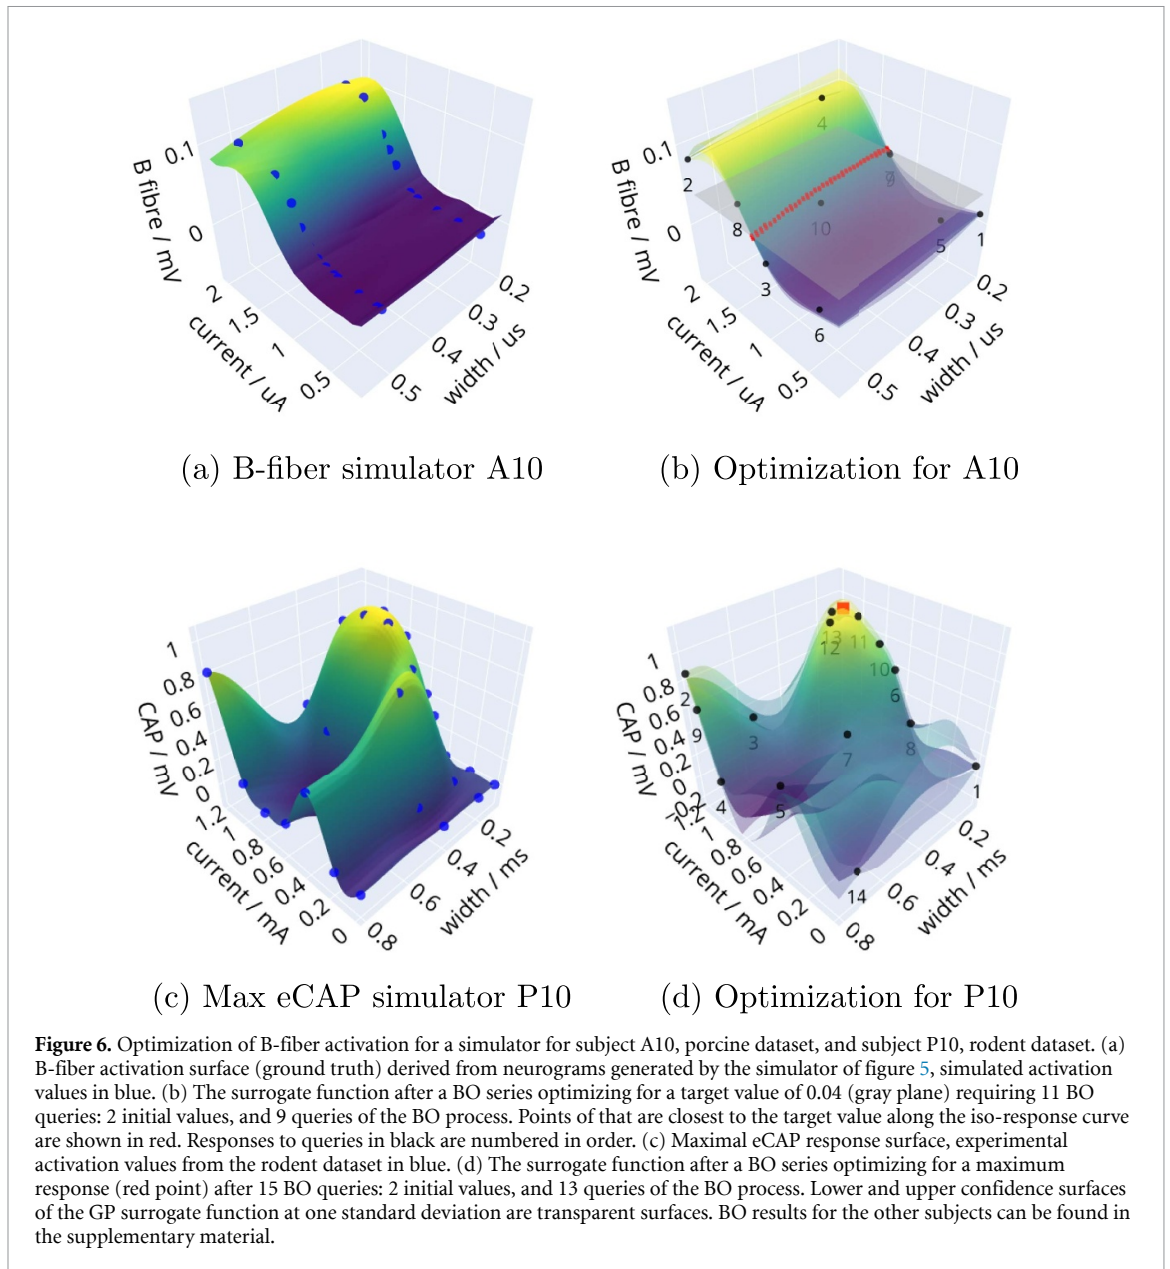
In order to evaluate the performance of a BO we compared it to the optimization results from a grid as well as a random search. The results are shown in figure 7(a) for the porcine and in figure 7(b) for the rodent dataset. The BO result (blue markers) is compared to the performance of a random search (gray markers). Also shown are the threshold values provided by the optima in grid searches. The BO matches or outperforms all grid searches after 13 evaluations. The exceptional performance of the grid searches and BO for P3, P4, and P5 in figure 7(b) is explained by the fact that the optima are in or close to one of the corners of the search space and therefore easily captured in a grid search as well as the initial BO evaluation.

Overall, apart from the greater efficiency in terms of the number of stimulations required, BO also provides a more robust search result compared to grid searches, which are sensitive to the grid layout relative to the optimum input.

3.3. Online BO of HR change

Simulation results for BO were encouraging and it was decided to apply OBOES online during surgery on porcine subjects to achieve targeted changes in HR through manipulation of VNS parameters.

A typical OBOES series with 11 evaluations is shown in figure 8 for subject A10 targeting a ΔHR of -10 bpm. After two initial evaluations OBOES



suggests further VNS parameters. The response to the stimulation is recorded, measured, and forwarded to the algorithm. The progression in updating OBOES's surrogate function is shown. Notice a narrowing of the upper and lower uncertainty surfaces (transparent) of the GP with a growing number of evaluations. Similar figures for subject A11 can be found in the supplementary material.

The final result for A10 is shown in figure 9 (and for A11 in supplementary material). In order to assess the quality of the surrogate surface three stimulations for each of three different inputs were performed (nine in total) as indicated in figure 9(a) by blue markers. Figure 9(b) shows the deviations from the predicted value for each of the three inputs. The posterior probability distribution of the estimated true experimental response is indicated. The expectation is that little probability mass p falls outside the

GP's confidence bounds, which is indeed the case ($p < 0.001$). In other words, the GP's upper and lower confidence surfaces provide reliable guidelines on how much uncertainty stemming from measurement noise and limited knowledge of the true response function affects predicting the effects of optimal stimulation. Since confidence bounds reduce with an increasing number of OBOES steps, they potentially provide a guideline when to stop the optimization procedure, namely, when the prediction uncertainty falls below a specific level of acceptance. One of the test points is slightly offset from the target value. This is due to a small error in the stimulation input. However, it fortuitously demonstrates that the confidence bounds typically get wider further away from the target, since they fall in a less explored region, while still providing valid confidence limits.

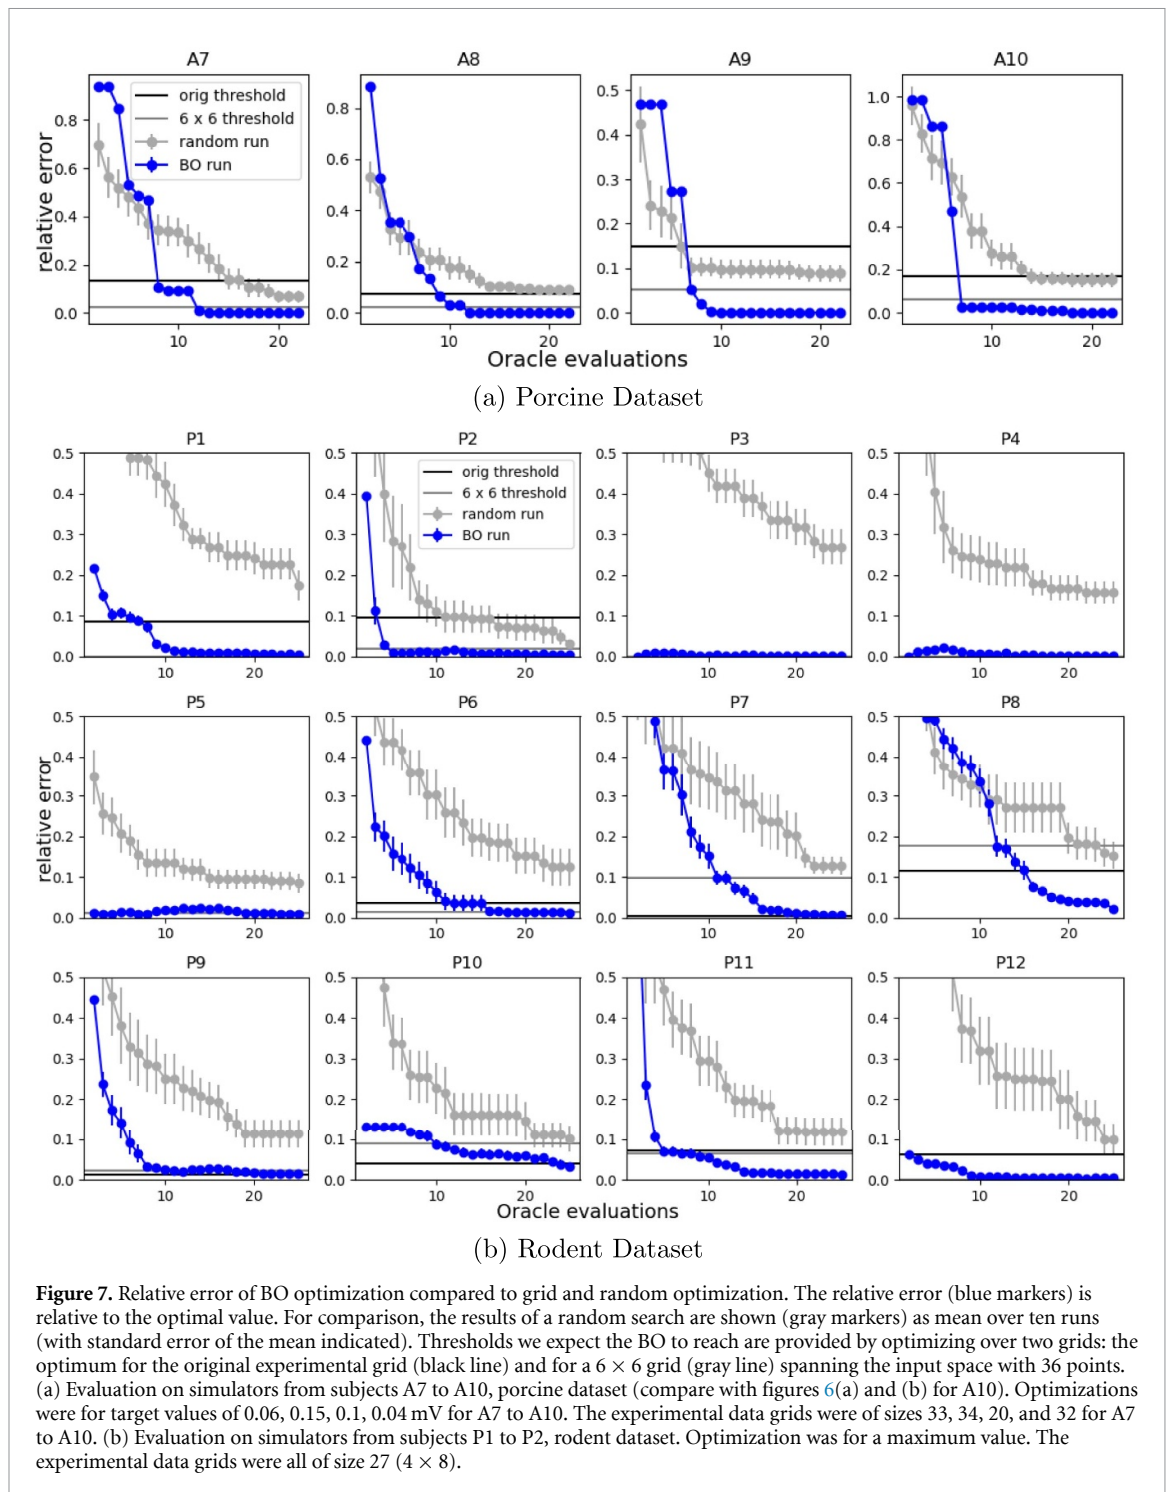


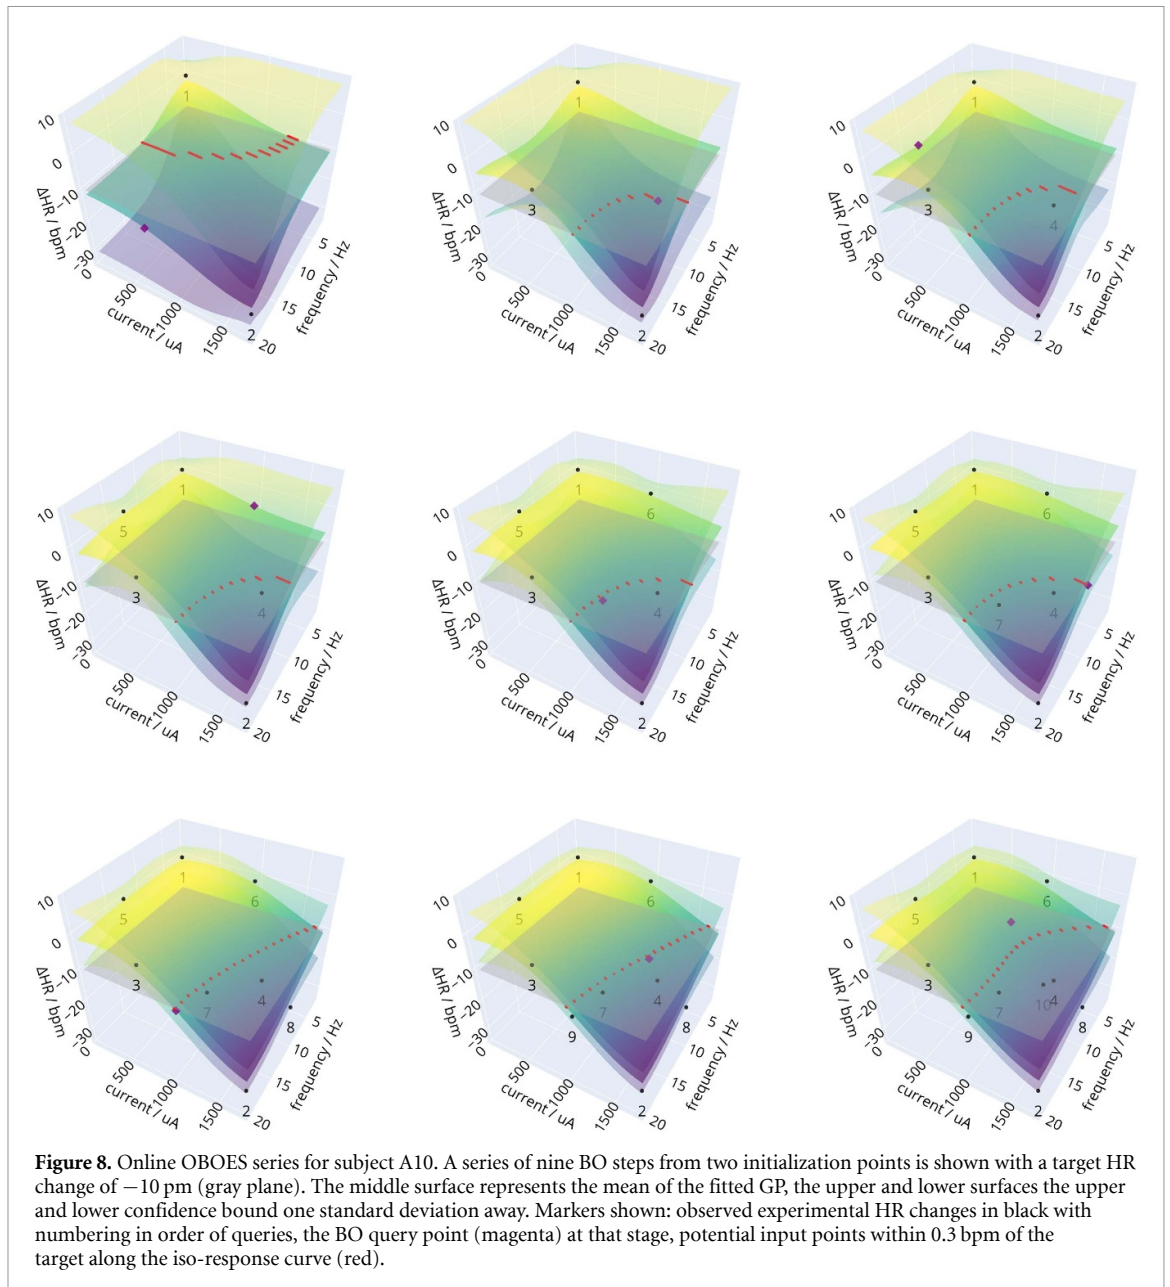
Figure 7. Relative error of BO optimization compared to grid and random optimization. The relative error (blue markers) is relative to the optimal value. For comparison, the results of a random search are shown (gray markers) as mean over ten runs (with standard error of the mean indicated). Thresholds we expect the BO to reach are provided by optimizing over two grids: the optimum for the original experimental grid (black line) and for a 6×6 grid (gray line) spanning the input space with 36 points. (a) Evaluation on simulators from subjects A7 to A10, porcine dataset (compare with figures 6(a) and (b) for A10). Optimizations were for target values of 0.06, 0.15, 0.1, 0.04 mV for A7 to A10. The experimental data grids were of sizes 33, 34, 20, and 32 for A7 to A10. (b) Evaluation on simulators from subjects P1 to P2, rodent dataset. Optimization was for a maximum value. The experimental data grids were all of size 27 (4×8).

3.4. Online BO of B fiber activation

The parasympathetic effect of B-fiber activation on HR is widely accepted (Qing *et al* 2018a). Figure 10 shows a comparison of the effect of stimulation current with that of B-fiber activation on Δ HR in subject A10. This suggests that optimizing B-fiber activation in place of HR is a viable alternative to optimizing HR directly.

To avoid side effects the total stimulation charge needs to be controlled. In the OBOES experiment of figure 11 we aimed to maximize B-fiber activation while minimizing the stimulation current (sequential

figures illustrating this optimization in supplementary materials). A way to achieve that is by maximizing the ratio between activation and current. As expected the maximum is achieved around 1500 μ A, away from the upper limit that was put on the current by the electrophysiologist. The pulse width was not considered a major cause of potential side effects and was not incorporated as a constraint into the objective function of the BO (i.e. the ratio between B-fiber activation and pulse current). The objective function therefore increases monotonically with the pulse width and the optimum is reached at the maximum of its



allowed range, at $500 \mu\text{s}$. Figure 11(b) demonstrates that the GP confidence range provides a reliable estimate of the uncertainty: the probability p that the true response, as estimated from three test points, lies outside the range is negligible ($p < 0.001$). However, although within the uncertainty range, some underestimation bias by the GP of the response is noticeable.

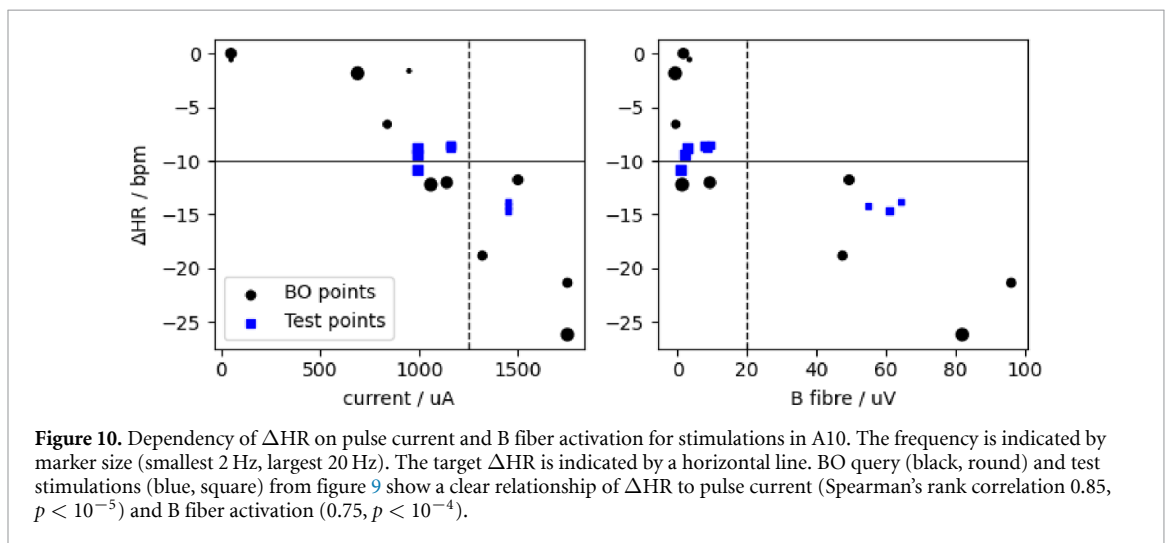
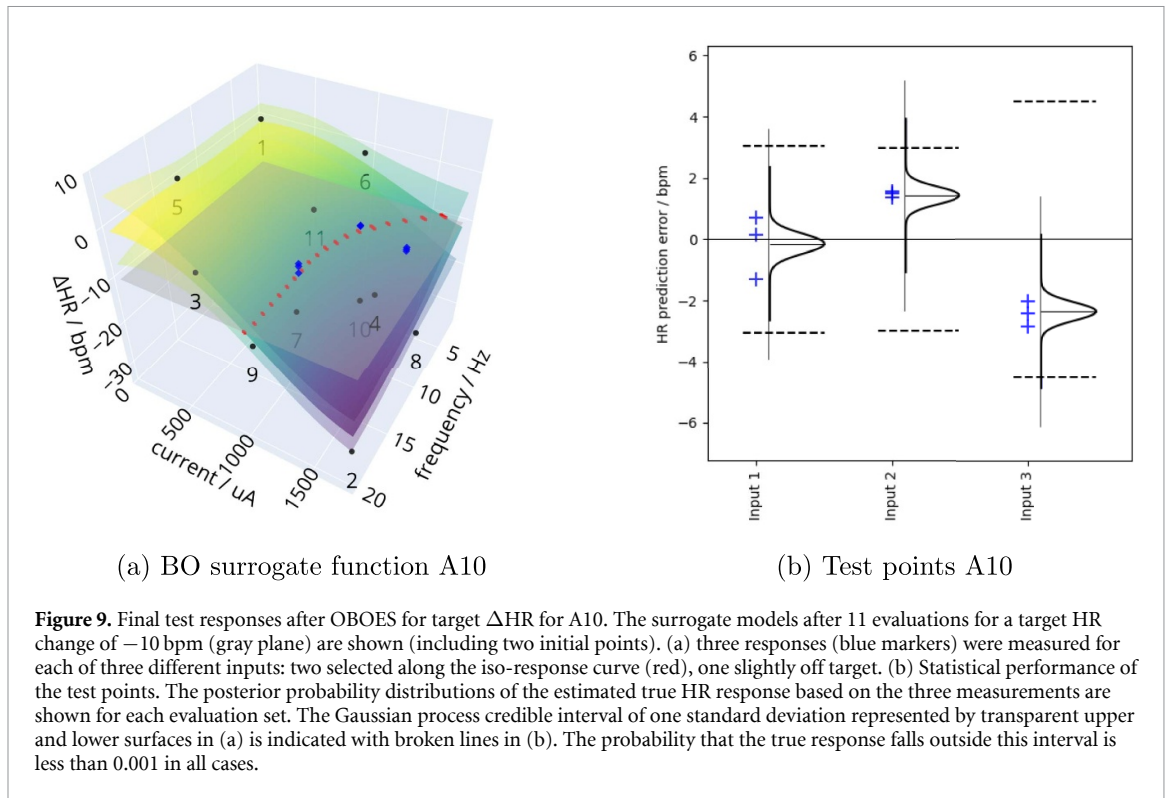
4. Discussion

To our knowledge the OBOES framework presented in this study is the first example of achieving targeted changes in HR and maximizing specific fiber activations efficiently through joint optimization of VNS parameters in alive animals. We also present a comprehensive cloud based R&D system enabling real time automated data collection, analysis

and optimization. Furthermore, we present a realistic simulation framework modeling responses to VNS. We demonstrate the usage of such simulation models for the development and testing of parameter optimization procedures for VNS.

4.1. Performance of OBOES

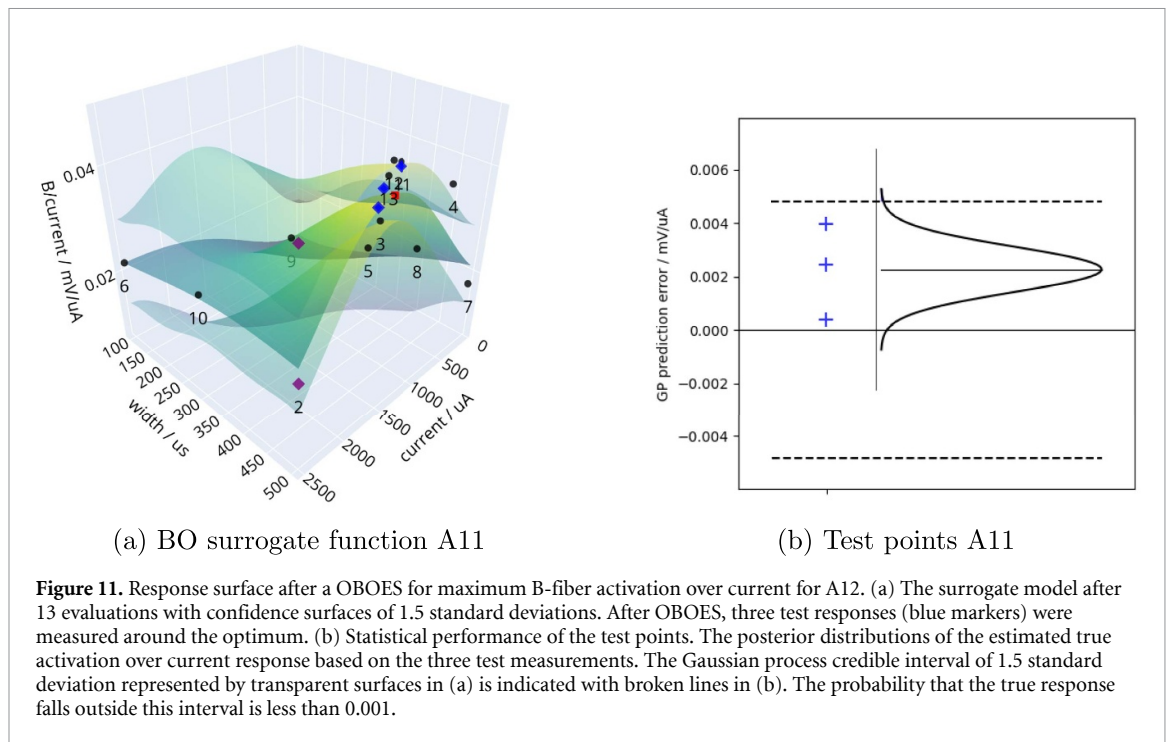
Current state-of-the-art clinical VNS systems use cardiovascular responses (Ardell *et al* 2017a) or tolerance-based dosing strategies (Premchand *et al* 2014) to guide VNS parameter choices. This is typically done on a reduced set of parameters, for example, through dose titration using only stimulation current. If multiple parameters are used, they are explored in a rigid order (LivaNova 2021). The neural stimulation optimization using a Bayesian method, BO as suggested in this study, is more flexible and efficient, particularly for high-dimensional parameter spaces.



BO is a well-developed and established ML technique (Kushner 1964) geared toward optimizing an unknown function via iterative queries. These queries, defined by the algorithm, are designed to both explore and optimize the function. Numerous empirical studies have employed BO to fine-tune neuromodulation parameters (Desautels *et al* 2015, Ashmaig *et al* 2018, Laferriere *et al* 2020, Park *et al* 2020, Losanno *et al* 2021, Stieve *et al* 2023). Several simulation studies on neuromodulation models, that enable a precise comparison with some ground truth, show its promising performance when optimizing parameters (Grado *et al* 2018, Mao *et al* 2023, Aiello *et al* 2023). Evaluations using simulated data have consistently showcased BO's superior performance

compared to random and grid searches—a point we highlight in this study as well—or greedy strategies (Bonizzato *et al* 2023). We emphasize that this study's primary objective was not to pioneer a completely new approach for optimizing nerve stimulations. Instead, we expand the range of previous applications of BO in neuromodulation by showcasing its efficacy when applied to VNS for controlling physiological states and eCAP activations in simulations as well as real-time experiments.

GPs are recognized for their resilience against the curse of dimensionality—the increasing demand for larger numbers of sample points in higher dimensions—by imposing smoothness assumptions on the response function (Spigler *et al* 2020).



Consequently, BO leveraging GPs becomes an ideal choice for optimizing functions in high-dimensional input spaces. Since the current study we began to apply OBOES to five dimensional parameter spaces including discrete spaces such as electrode locations, which will be reported on in future studies. Due to its speed of optimization, OBOES is better suited to cope with variation in bioelectronic coupling of the electrode–nerve interface at implantation or changes at the interface over time, for example, with inflammation (Qiao *et al* 2016). Given the frequent need to adapt parameter settings in response to changes at the nerve-cuff interface, there is a compelling motivation for an adjustment approach, such as OBOES, that minimizes the number of required stimulations.

Our choice to develop and test optimization algorithms through simulations, derived from real data, proved invaluable. Not only did it offer a sandbox environment for refining our algorithm, but it also highlighted potential challenges and pitfalls we might encounter in real-time experimental settings. For increased realism we derived response functions for the simulations from actual data. This was for the purpose of enhancing realism rather than replicating the exact true response function. Achieving a wholly accurate representation would demand rigorous vetting and testing of assumptions, a process beyond our primary study’s scope. Instead, we sought to capture salient features of the true response, such as its overall shape and noise characteristics. Capturing these features via GP modeling offered a balance between mirroring genuine conditions and maintaining computational feasibility. They ensured that OBOES underwent testing under conditions that resonated closely,

albeit not perfectly, with the challenges of real-world applications.

In the simulation studies, OBOES performed well compared to grid searches and random searches in terms of the number of steps to get close to an optimum as well as the quality of the approximation (figure 7 of section 3.2). The aim of this study, however, was not necessarily to demonstrate superiority over alternative optimization approaches. There are undoubtedly many approaches that perform comparably in terms of requiring only a small number of steps to reach an optimum, BO being one of them. We opted for a BO approach because it (i) enables a principled probabilistic way to integrate prior knowledge, e.g. about the smoothness of the response function, (ii) constructs an explicit model of the unknown response function in the form of a surrogate function, (iii) is less susceptible to increasing complexity with increasing dimensions, the ‘curse of dimensionality’ and (iv) provides information about the amount of model uncertainty. This last property allows one, for example, to assess the plausibility of the GP’s internal surrogate model and to make a rational assessment of the risk when relying on its predictions in a clinical setting.

A BO approach also enables great flexibility in the choice or design of the acquisition function, which can be adapted to specific requirements. For example, in this study we developed a variant, TCB, of an acquisition function that allows a BO to home in on an iso-response curve of inputs that result in the same defined target value as responses. TCB turned out to be very robust, in simulations as well as in online settings, leading to fast convergence to an iso-response set of inputs. An additional benefit of this

approach is that the iso-response curve allows one to pick preferred parameter sets with ancillary benefits, for example, mitigating side effects or promoting device battery life, while maintaining the same biological response.

In this study, GP surrogate functions for OBOES were constructed for each subject individually. We experimented with GP models based on coregionalization kernels (Álvarez *et al* 2012), that allowed us to share information between several subjects. Such models showed promise when applied to simulated data by reducing the number of query evaluations even further. However, the gain in efficiency was only noticeable for simulated subjects with similar responses. In this study as well as for further experiments on life subjects—as discussed in greater detail in the companion manuscript Berthon *et al* (2023)—the responses from different subjects to the same stimulation inputs typically diverged considerably and the added complexity of applying a model based on multilevel modeling or coregionalization did not seem to outweigh the small loss in efficiency compared to a subject specific GP model.

4.2. Online implementations for medical devices

The ITx Platform of section 2.3 shows two services where a ML model can be deployed: either on the acquisition hardware, an edge deployment, or in the ITx Platform portal, a cloud deployment. We opted for the second choice, since we aimed for a proof-of-principle demonstration. A cloud deployment allowed us to manually adjust settings and algorithms and to control stimulus parameters if intervention was deemed necessary. The online BO experiments, however, were surprisingly stable and were obtained with little manual intervention. The use of GP parameter priors stabilized their optimization, particularly in the early phases of the BO where only few data points are available. We have since deployed OBOES directly on the acquisition hardware to run fully automated with the option to monitor and intervene manually if necessary (results not shown here).

Safety considerations are paramount for medical devices, and those leveraging ML techniques are no exception. BO is a probabilistic method that balances exploration with optimization of an unknown response function with as few data points as possible. It is possible to tweak the balance between optimization and exploration through the confidence bound factor λ of section 2.10. Here we set the factor to a constant that performed well in simulations. Typically the greatest uncertainty tends to occur at the more extreme values of the input space. There exists a risk that high input values for stimulations may lead to intolerable side effects. In our current implementation of OBOES, we address this by predefining constraints on the input space in collaboration with an electrophysiologist overseeing the algorithm's

application. As long as the patient's responses stay within acceptable limits, the size of the input space can be increased gradually, allowing higher charge stimulations to be explored by the model. An alternative approach, which we will explore in future, could exploit the GPs' natural handling of uncertainty to pursue dose titration in a safe way (Berkenkamp *et al* 2016): analysis of the surrogate function can in principle suggest safe extensions of the search space constrained by a risk model that balances expansion of the search space with safety considerations.

4.3. Neural biomarkers for physiological effects

Anesthetics are known to severely dampen the baroreflex and other hemodynamic indices (Ahmed *et al* 2021). We observed that to maintain the anesthetic plane, propofol administration occasionally required adjustment. During grid searches in online experiments, this sometimes drastically changed the magnitude of bradycardia for fixed stimulation parameters. In contrast, nerve recruitment itself is much more robust to anesthesia or the severity of heart disease. Furthermore, neurograms can be explored rapidly and with few physiological side effects by keeping pulse numbers low (by lowering frequency and train duration). For an intraoperative procedure, optimizing neural biomarkers is thus preferable over optimizing physiological ones. We investigate the relationship between eCAPs and physiological responses in detail in a companion study (Berthon *et al* 2023). We have shown that OBOES is suitable for optimizing B-fiber activation, a prime example for an indirect optimization of HR using a neural biomarker. We also hypothesize that neural biomarkers, based on their enhanced reliability as measures of mechanism engagement and bioelectronic coupling, make good candidates for true closed-loop neuromodulation therapies. Of course, utilizing neural biomarkers in this way depends on the ability to model the relationship between neural and physiological responses to neural stimulation in a precise way, which is the subject of Berthon *et al* (2023).

4.4. Clinical use of intraoperative optimization

VNS parameter optimization for neural or physiological targets can be performed any time after the implantation of the device. We have successfully performed such optimizations in an awake, freely moving animal (as will be reported in a forthcoming study).

However, intraoperative optimization during clinical implantation, while monitoring fiber engagement, could provide immediate objective feedback on (1) the successful installation of the stimulation lead, (2) whether the intended fiber group is being activated and (3) a range of stimulation parameters (on the iso-response line) that can engage the target mechanism. After implantation, device titration can be guided by the range of stimulation parameters

known to generate on-target therapeutic effects, as established during implantation. This could drastically reduce the number of clinical interactions needed following device implantation and the overall time required to achieve optimal postoperative dosing.

4.5. Limitations and future directions

A limitation of the current study is that OBOES was performed on anesthetized subjects and caution is necessary when extrapolating the findings to awake individuals. We are actively exploring the efficacy of OBOES on awake subjects and anticipate presenting these findings in subsequent publications. One challenge in studying awake subjects is acquiring consistent physiological measurements exclusively linked to neural stimulation; this is due to increased interference from other external factors like movement. A GP framework, in theory, is well suited to tackle such challenges, for example, by incorporating additional external data, such as readings from movement sensors. However, this integration may complicate measurement tools and computational processes. Future experiments are essential to strike a balance between accuracy on one hand and complexity of data acquisition and processing on the other.

In this study, as well as in a companion study (Berthon *et al* 2023), we noted considerable variability in responses to stimulations, both between subjects and within the same subject over time. While the OBOES approach we have detailed assumes a consistent physiological response throughout an optimization run (relying only on data gathered during that run), the versatility of the GP framework allows for the inclusion of data from other subjects or past optimization sessions. Preliminary tests with simulated data indicate that an approach based on coregionalization kernels (Álvarez *et al* 2012) is able to integrate data from multiple subjects. This integration can, in turn, reduce the number of BO steps needed for a new subject, if prior data are available for subjects with similar response patterns. This integration can, in turn, reduce the number of BO steps needed, especially when database subjects exhibit similar response patterns. Subjects presented in this study, however, diverged too much in their responses for multiple subject integration to improve BO estimation sufficiently, when considering the additional complexity of the approach.

While GP regression offers remarkable flexibility in modeling diverse response functions, it does operate under certain assumptions (Rasmussen 2006). Of particular practical significance are the assumptions about the function's smoothness and the nature and magnitude of noise. Smoothness relates to the expected extent of response variability in the face of changing stimulation parameters: do small changes in parameters have a large or small effect. GP regression enables modeling of rougher or smoother responses by a choice of kernel type and the length scale

parameters. In contrast, noise relates to the variation observed in response to identical stimulation inputs: unavoidable variation when repeating a simulation. In our simulations and practical applications, the finer details of the type of smoothness and the exact probability distributions of noise had minimal impact on the overall outcomes. Among a range of kernel types we explored, squared exponential kernels consistently produced stable and satisfactory results. Moreover, assuming Gaussian noise not only proved effective but also facilitated efficient estimation. Even though hyperparameters for length scales and noise can be estimated from data, we observed that putting some broad prior constraints on them, as described in the [methods](#) section, increased stability.

In this study, our emphasis was on two-dimensional input spaces: combining pulse width and current for B-fiber optimization, as well as frequency and current for HR optimization. This approach was primarily chosen for its simplicity and to allow for a graphical representation of the progress of OBOES. We are presently exploring the possibility of expanding these parameter spaces. This includes parameters such as train duration, multiple electrode locations, pulse shapes (defined by a finite set of continuous parameters), or the clustering of pulses into multiple bursts, among others. As highlighted in Berthon *et al* (2023), the output dimension could be expanded to monitor potential side effects, such as bradypnea or laryngeal spasms. This results in a multitask optimization problem. How to deal with multiple tasks, either by combining them into one target function, or by maintaining several surrogate functions, will be the subject of future studies.

5. Conclusion

Optimization of VNS parameters for optimizing a target response, particularly in an intraoperative setting, is challenging. We demonstrated that BO is well-suited to support this task. For offline exploration of a wide range of neural biomarkers, accurate simulation models based on Gaussian processes can be derived from experimental data. For online optimization, BO enables an efficient and accurate reconstruction and robust optimization of the response function with comparatively few test stimulations.

Data availability statement

A simulation framework for applying OBOES including data used in this study for demonstration purpose is available through the o^2S^2 PARC platform³. The data that support the findings of this study will be openly available in downloadable format following a delay on the same platform. The data

³ <https://osparc.io/study/861f8d70-a2d4-11ed-8c93-02420ab00f9>.

cannot be made publicly available upon publication because they contain commercially sensitive information. The data that support the findings of this study are available upon reasonable request from the authors. <https://osparc.io/study/861f8d70-a2d4-11ed-8c93-02420a0b00f9>.

Funding

This research was, in part, funded by the National Institutes of Health (NIH) under other transaction award number OT2OD030536. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the NIH. BIOS acknowledges support from the MEDTEQ+ program, including contributions from Healthy Brains for Healthy Lives, CFREF and Mitacs.

Authors contribution

L W designed the OBOES algorithms and analyzed the data. L W and T E prepared the first draft of the manuscript. T E, P F P, O A, M S designed the animal experiments. P F P and J M guided the animal experiments. L W, A B, O T L designed, implemented and operated the BO interface of the ITx platform. B P, W B, M J designed and implemented the cloud based data processing and storage pipeline. B A, W B, S G, M J designed and implemented the NeuroTool interface. S S, S G, P G, M P, M S, B A, W B, M J designed and built the BIOS Neural Interface. E S, M T, C H, S L, J J, A T, G L contributed to reviewing and rewriting of the manuscript. E H, A T, O A, T E conceived of the study. O A supervised the study.

Conflict of interest

All authors, except G L, A T, and J J, are (or were at the time of their contribution) employees of BIOS Health Ltd and declare that BIOS Health has filed US and international patent applications relating to a system, apparatus and method for utilizing neural biomarker response in clinical decision-making. The concepts of this work is contained in the UK Patent GB 2214547.8. A T and J J declare a consulting role with BIOS Health at the time the research was conducted. G L declares a present consulting role with BIOS Health on topics related to the present publication, but that was not in place at the time of the development of this manuscript.

ORCID iDs

Lorenz Wernisch  <https://orcid.org/0000-0003-3113-7942>

Oliver Armitage  <https://orcid.org/0000-0001-8790-7710>

References

- Ahmed U, Chang Y-C, Lopez M F, Wong J, Datta-Chaudhuri T, Rieth L, Al-Abed Y and Zanos S 2021 Implant- and anesthesia-related factors affecting cardiopulmonary threshold intensities for vagus nerve stimulation *J. Neural Eng.* **18** 046075
- Aiello G, Valle G and Raspopovic S 2023 Recalibration of neuromodulation parameters in neural implants with adaptive Bayesian optimization *J. Neural Eng.* **20** 026037
- Álvarez M A, Rosasco L and Lawrence N D 2012 Kernels for vector-valued functions: a review *Found. Trends Mach. Learn.* **4** 195–266
- Archetti F and Candelieri A 2019 *Bayesian Optimization and Data Science* (Springer)
- Ardell J L, Nier H, Hammer M, Southerland E M, Ardell C L, Beaumont E, KenKnight B H and Armour J A 2017a Defining the neural fulcrum for chronic vagus nerve stimulation: implications for integrated cardiac control *J. Physiol.* **595** 6887–903
- Arlotti M et al 2018 Eight-hours adaptive deep brain stimulation in patients with Parkinson disease *Neurology* **90** e971–6
- Ashmaig O, Connolly M, Gross R E and Mahmoudi B 2018 Bayesian optimization of asynchronous distributed microelectrode theta stimulation and spatial memory 2018 *40th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBC)* vol 2018 pp 2683–6
- Berkenkamp F, Krause A and Schoellig A P 2016 Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics *CoRR* (arXiv:1602.04450)
- Berthon A et al 2023 Using neural biomarkers to personalize dosing of vagus nerve stimulation *bioRxiv Preprint* <https://doi.org/10.1101/2023.08.30.555487> (posted online 1 September 2023, accessed 5 September 2023)
- Bonizzato M et al 2023 Autonomous optimization of neuroprosthetic stimulation parameters that drive the motor cortex and spinal cord outputs in rats and monkeys *Cell Rep. Med.* **4** 101008
- BorealisAI 2020 Bayesian optimization (available at: www.borealisai.com/en/blog/tutorial-8-bayesian-optimization/)
- Carr M J and Udem B J 2003 Bronchopulmonary afferent nerves *Respirology* **8** 291–301
- Chang R B, Strohlic D E, Williams E K, Umans B D and Liberles S D 2015 Vagal sensory neuron subtypes that differentially control breathing *Cell* **161** 622–33
- Cracchiolo M, Ottaviani M M, Panarese A, Strauss I, Vallone F, Mazzoni A and Micera S 2021 Bioelectronic medicine for the autonomic nervous system: clinical applications and perspectives *J. Neural Eng.* **18** 041002
- Desautels T A, Choe J, Gad P, Nandra M S, Roy R R, Zhong H, Tai Y-C, Edgerton V R and Burdick J W 2015 An active learning algorithm for control of epidural electrostimulation *IEEE Trans. Biomed. Eng.* **62** 2443–55
- Dusi V and De Ferrari G M 2021 Vagal stimulation in heart failure *Herz* **46** 541–9
- Grado L L, Johnson M D and Netoff T I 2018 Bayesian adaptive dual control of deep brain stimulation in a computational model of Parkinson's disease *PLOS Comput. Biol.* **14** e1006606
- Guidetti M et al 2021 Clinical perspectives of adaptive deep brain stimulation *Brain Stimul. Basic Transl. Clin. Res. Neuromodul.* **14** 1238–47
- Jayaprakash N et al 2023 Organ- and function-specific anatomical organization of vagal fibers supports fascicular vagus nerve stimulation *Brain Stimul.* **16** 484–506
- Konstam M A et al 2019 Impact of autonomic regulation therapy in patients with heart failure *Circ. Heart Fail.* **12** e005879
- Kushner H J 1964 A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise *J. Basic Eng.* **86** 97–106

- Labiner D M and Ahern G L 2007a Vagus nerve stimulation therapy in depression and epilepsy: therapeutic parameter settings *Acta Neurol. Scand.* **115** 23–33
- Laferrriere S, Bonizzato M, Cote S L, Dancause N and Lajoie G 2020 Hierarchical Bayesian optimization of spatiotemporal neurostimulations for targeted motor outputs *IEEE Trans. Neural Syst. Rehabil. Eng.* **28** 1452–60
- LivaNova 2021 LivaNova (available at: www.livanova.com)
- Losanno E, Badi M, Wurth S, Borgognon S, Courtine G, Capogrosso M, Rouiller E M and Micera S 2021 Bayesian optimization of peripheral intraneural stimulation protocols to evoke distal limb movements *J. Neural Eng.* **18** 066046
- Mao X, Chang Y-C, Zanos S and Lajoie G 2023 Rapidly inferring personalized neurostimulation parameters with meta-learning: a case study of individualized fiber recruitment in vagus nerve stimulation *bioRxiv Preprint* <https://doi.org/10.1101/2022.09.06.506839> (posted online 4 November 2022, accessed 14 November 2022)
- Martin K A, Papadoyannis E S, Schiavo J K, Fadaei S S, Temiz N Z, McGinley M J, McCormick D A and Froemke R C 2022 Vagus nerve stimulation recruits the central cholinergic system to enhance perceptual learning *bioRxiv Preprint* <https://doi.org/10.1101/2022.01.28.478197> (Posted online 28 January 2022, accessed 1 September 2023)
- Meyers E C, Solorzano B R, James J, Ganzer P D, Lai E S, Rennaker R L, Kilgard M P and Hays S A 2018 Vagus nerve stimulation enhances stable plasticity and generalization of stroke recovery *Stroke* **49** 710–7
- Mockus J 2012 *Bayesian Approach to Global Optimization: Theory and Applications* (Kluwer Academic)
- Nearing B D, Anand I S, Libbus I, DiCarlo L A, KenKnight B H and Verrier R L 2020 Vagus nerve stimulation provides multiyear improvements in autonomic function and cardiac electrical stability in the ANTHEM-HF study *J. Card. Fail.* **27** 208–16
- Nemeroff C B, Mayberg H S, Krahl S E, McNamara J, Frazer A, Henry T R, George M S, Charney D S and Brannan S K 2006 VNS therapy in treatment-resistant depression: clinical evidence and putative neurobiological mechanisms *Neuropsychopharmacology* **31** 1345–55
- Nicolai E N et al 2020 Sources of off-target effects of vagus nerve stimulation using the helical clinical lead in domestic pigs *J. Neural Eng.* **17** 046017
- Park S-E, Connolly M J, Exarchos I, Fernandez A, Ghetiya M, Gutekunst C-A and Gross R E 2020 Optimizing neuromodulation based on surrogate neural states for seizure suppression in a rat temporal lobe epilepsy model *J. Neural Eng.* **17** 046009
- Pavlov V A and Tracey K J 2012 The vagus nerve and the inflammatory reflex—linking immunity and metabolism *Nat. Rev. Endocrinol.* **8** 743–54
- Premchand R K et al 2014 Autonomic regulation therapy via left or right cervical vagus nerve stimulation in patients with chronic heart failure: results of the ANTHEM-HF trial *J. Card. Fail.* **20** 808–16
- Qiao S, Stieglitz T and Yoshida K 2016 Estimation of the electrode-fiber bioelectrical coupling from extracellularly recorded single fiber action potentials *IEEE Trans. Neural Syst. Rehabil. Eng.* **24** 951–60
- Qing K Y, Wasilczuk K M, Ward M P, Phillips E H, Vlachos P P, Goergen C J and Irazoqui P P 2018a B fibers are the best predictors of cardiac activity during vagus nerve stimulation *Bioelectron. Med.* **4** 5
- Rasmussen C E 2006 *Gaussian Processes for Machine Learning* (MIT Press)
- Reid S A 1990 Surgical technique for implantation of the neurocybernetic prosthesis *Epilepsia* **31** S38–S39
- Seabold S and Perktold J 2010 Statsmodels: econometric and statistical modeling with Python *9th Python in Science Conf.*
- Settell M L et al 2020 Functional vagotomy in the cervical vagus nerve of the domestic pig: implications for the study of vagus nerve stimulation *J. Neural Eng.* **17** 026022
- Spigler S, Geiger M and Wyart M 2020 Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm *J. Stat. Mech.* **124001**
- Stieve B J, Richner T J, Krook-Magnuson C, Netoff T I and Krook-Magnuson E 2023 Optimization of closed-loop electrical stimulation enables robust cerebellar-directed seizure control *Brain* **146** 91–108
- Upadhye A R et al 2022 Fascicles split or merge every ~560 microns within the human cervical vagus nerve *J. Neural Eng.* **19** 054001
- Ward M P, Nowak T, Phillip R, Tan Z and Powley T 2021 Influence of left vagal stimulus pulse parameters on vagal and gastric activity in rat (available at: <https://sparc.science/datasets/9/version/4>)
- Ward M P, Qing K Y, Otto K J, Worth R M, John S W M and Irazoqui P P 2015 A flexible platform for biofeedback-driven control and personalization of electrical nerve stimulation therapy *IEEE Trans. Neural Syst. Rehabil. Eng.* **23** 475–84