

Highly efficient modeling and optimization of neural fiber responses to electrical stimulation

Received: 16 February 2024

Accepted: 13 August 2024

Published online: 31 August 2024

 Check for updatesMinhaj A. Hussain¹, Warren M. Grill^{1,2,3,4} & Nicole A. Pelot¹ 

Peripheral neuromodulation has emerged as a powerful modality for controlling physiological functions and treating a variety of medical conditions including chronic pain and organ dysfunction. The underlying complexity of the nonlinear responses to electrical stimulation make it challenging to design precise and effective neuromodulation protocols. Computational models have thus become indispensable in advancing our understanding and control of neural responses to electrical stimulation. However, existing approaches suffer from computational bottlenecks, rendering them unsuitable for real-time applications, large-scale parameter sweeps, or sophisticated optimization. In this work, we introduce an approach for massively parallel estimation and optimization of neural fiber responses to electrical stimulation using machine learning techniques. By leveraging advances in high-performance computing and parallel programming, we present a surrogate fiber model that generates spatiotemporal responses to a wide variety of cuff-based electrical peripheral nerve stimulation protocols. We used our surrogate fiber model to design stimulation parameters for selective stimulation of pig and human vagus nerves. Our approach yields a several-orders-of-magnitude improvement in computational efficiency while retaining generality and high predictive accuracy, demonstrating its robustness and potential to enhance the design and optimization of peripheral neuromodulation therapies.

The success of therapeutic electrical stimulation to treat disorders—spanning inflammatory, cardiovascular, cognitive, metabolic, and pain conditions, among others—hinges on appropriate modulation of targeted neurons. Neural responses to stimulation are highly nonlinear, and they are influenced by the delivered electrical signal, physical electrode-tissue relationships, and neuronal biophysics. Designing effective therapies thus requires analysis of a vast parameter space, including waveform shape, amplitude, frequency, pattern, active contact configuration, as well as individual and inter-species differences. Computational models of neural anatomy and biophysics enable rigorous optimization of application-specific parameters¹ (e.g.,

for selective stimulation of peripheral nerves^{2–5}), as well as investigation of the mechanisms of complex non-linear phenomena (e.g., conduction block^{6–8}). However, these approaches are limited by exceedingly high computational costs due to the non-linear properties of neurons. Therefore, we designed a highly efficient model of nerve fibers, and we used this “surrogate” model for rapid estimation of stimulation parameters to achieve selective activation.

Recent computational tools support neural simulations on GPUs^{9–11} and provide significant speedup over traditional CPU-based methods. However, extracellular stimulation cannot be modeled using these tools: NEURON—the “industry-standard” neural simulation

¹Department of Biomedical Engineering, Duke University, Durham, NC 27708, USA. ²Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA. ³Department of Neurobiology, Duke University, Durham, NC 27708, USA. ⁴Department of Neurosurgery, Duke University, Durham, NC 27708, USA. ✉e-mail: nikki.pelot@duke.edu

environment—is CPU-based, and it is the only platform that supports simulating the effects of extracellular voltages (as produced by extracellular stimulation) and the complex ultrastructure of state-of-the-art fiber models¹².

Prior efforts to reduce computational demand designed surrogate models that enabled higher throughput of simulations^{13–15}. However, these surrogates were restricted to a predetermined class of waveform and thus did not permit full exploration of the parameter space, including temporal features of stimulation. Further, they only considered fiber activation from rest, and thus could not predict other neural responses such as conduction block, action potential collision, subthreshold modulation, or dependence on prior history of excitation. These surrogates also only predicted activation thresholds, without spatiotemporal responses of the gating parameters and transmembrane voltage, which are critical for mechanistic studies.

Our objective is to estimate rapidly the full spatiotemporal response of large populations of nerve fibers to electrical stimulation. Previous studies across many fields applied machine learning to model physical systems by solving specific initial- or boundary-value problems^{16,17}, and recently, neural networks to solve such differential equations were conditioned to account directly for the governing physical laws¹⁸. These methods provide significant computational benefits, such as mitigating the need for fine spatial and/or temporal discretization while yielding high-resolution solutions; however, they offer limited speed gains for our purposes due to the need for optimization that is specific to the particular set of initial and/or boundary conditions, i.e., specific extracellular stimulation conditions. Other recent studies showed significant computational acceleration using generic function approximators to emulate mechanistic biological models^{19,20}. However, these methods do not accommodate enforcing boundary conditions that vary continuously in both space and time, as is necessary to represent electrical stimulation.

In this work, we present a framework, AxonML²¹, to implement, parameterize, efficiently execute, and perform optimization on GPU-based models of peripheral nerve fibers. Under the AxonML framework, we developed a high-throughput generalized computational model of myelinated peripheral fibers—hereafter labeled as the surrogate myelinated fiber, S-MF, pronounced “smurf”—that accurately predicts responses to electrical stimulation orders-of-magnitude more quickly than conventional methods, thus enabling efficient optimization of stimulation parameters for neural interface design. S-MF reproduces the spatiotemporal dynamics of the McIntyre-Richardson-Grill (MRG) fiber¹², a well-established and well-validated nonlinear model of a myelinated mammalian fiber; hereafter, we refer to the reference implementation of the MRG fiber as the “NEURON” model. The MRG model is the current gold standard for predicting fiber responses to electrical stimulation²²; it was developed using detailed electrophysiological data¹², and it has been validated experimentally to design new energy-efficient modes of stimulation^{23,24}, to predict responses to high-frequency signals for nerve block²⁵, and to predict species-appropriate stimulation thresholds of *in vivo* experiments²⁶. It has also been used to accurately predict responses to unconventional methods of stimulation, such as transcatheter amplitude-modulated signals for the treatment of overactive bladder²⁷. Commercially, the MRG model has been used in the development of a number of FDA-approved platforms to predict the volume of tissue activated, and it is the basis for Boston Scientific’s digital interface for setting clinical deep brain stimulation parameters in patients with Parkinson’s disease (Vercise Neural Navigator with STIMVIEW XT)²⁸.

S-MF is executed on a GPU for exceptional computational efficiency and uses a simplified myelinated cable geometry with nonlinear ionic conductances that are reparameterized to reproduce accurately the responses of the NEURON model. AxonML can also be used to implement other fiber models; we show that it generalizes effectively to produce a surrogate of an established model of an

unmyelinated C-fiber²⁹. We quantified the performance of S-MF across a wide range of parameters: monopolar and multipolar stimulation, fiber diameters, nerve morphologies, stimulus waveforms, states of intrinsic neural activity, and neuromodulation protocols (including excitation and kilohertz frequency block). We compared the performance of S-MF to simulations in NEURON. We used S-MF to design stimulation parameters that achieve selective stimulation in pig and human vagus nerves using gradient-free and gradient-based optimization approaches. We achieved 2,000 to 130,000× speedup over single-core simulations in NEURON, while maintaining accuracy.

Results

We solved morphologically realistic models of human ($n=1$) and pig ($n=1$) vagus nerve stimulation (VNS) using the finite element method to simulate distributions of electric potential within the nerves. We applied these potentials as time-varying boundary conditions to McIntyre-Richardson-Grill (MRG)¹² models of myelinated fibers (Fig. 1f) in NEURON to generate a dataset of responses to linear combinations of rectangular waveforms with randomized amplitudes, delays, and pulse widths, delivered with the six-contact ImThera cuff electrode (Fig. 1a). We used this dataset to train by backpropagation and gradient descent a surrogate model of the MRG fiber, implemented as a simplified cable model (Fig. 1g) with trainable parameters.

We used the trained surrogate model (S-MF) to simulate VNS, a pivotal neuromodulation method due to its versatility—including treatment of drug-resistant epilepsy³⁰, treatment-resistant depression³¹, stroke sequelae³², and heart failure³³—and favorable safety profile³⁴. We tested S-MF for accuracy in predicting activation thresholds to stimulus pulses with various waveforms, as well as accuracy in predicting other nonlinear responses to stimulation in pig ($n=6$) and human ($n=6$) vagus nerve models instrumented with either an ImThera (six-contact) or helical (bipolar, circumneural) cuff electrode (Fig. 1a,b).

We used S-MF to optimize stimulation parameters for selective activation of fascicles in models of pig and human VNS, using both rectangular biphasic and arbitrary waveforms. We implemented two optimization methods, gradient-free and gradient-based, and we compared performance across tasks and optimization methods using both the NEURON model and S-MF.

Surrogate model accurately and rapidly predicts activation thresholds

We quantified the accuracy of S-MF’s activation thresholds. Activation thresholds are the most common output metric in neuromodulation modeling studies. The MRG fiber model has been used to explain neuroanatomy-dependent differences in activation thresholds for responses related to heart rate changes and neck muscle contractions in animal experiments of VNS³⁵. We placed surrogate fibers of various diameters in pig ($n=6$) and human ($n=6$) vagus nerve models instrumented with one of two cuffs geometries: ImThera, with six circular contacts (Fig. 1a), or LivaNova helical, with two helical contacts spanning -270° around the circumference of the nerve (Fig. 1b). We delivered stimulus currents with waveforms of various shapes (see Methods “Surrogate model: Testing—activation thresholds”) and pulse widths, in monopolar configuration with ImThera and bipolar configuration with the helical cuff. We determined to within 1% accuracy the minimum current amplitude required to evoke at least one propagating action potential (AP) in each S-MF for all stimulus configurations ($n=17,280$ for ImThera and $n=25,200$ for helical). We compared the S-MF responses to MRG fibers implemented in NEURON, stimulated under the same conditions; see Methods “Surrogate model: Testing—activation thresholds” for full details.

S-MF accurately predicted NEURON thresholds across fiber diameters, nerve morphologies, electrode geometries, and waveforms ($R^2=0.999$; Fig. 2a), with mean absolute percentage error (MAPE) of $<2.5\%$ across all tested fiber diameters (6–14 μm), with lower error for

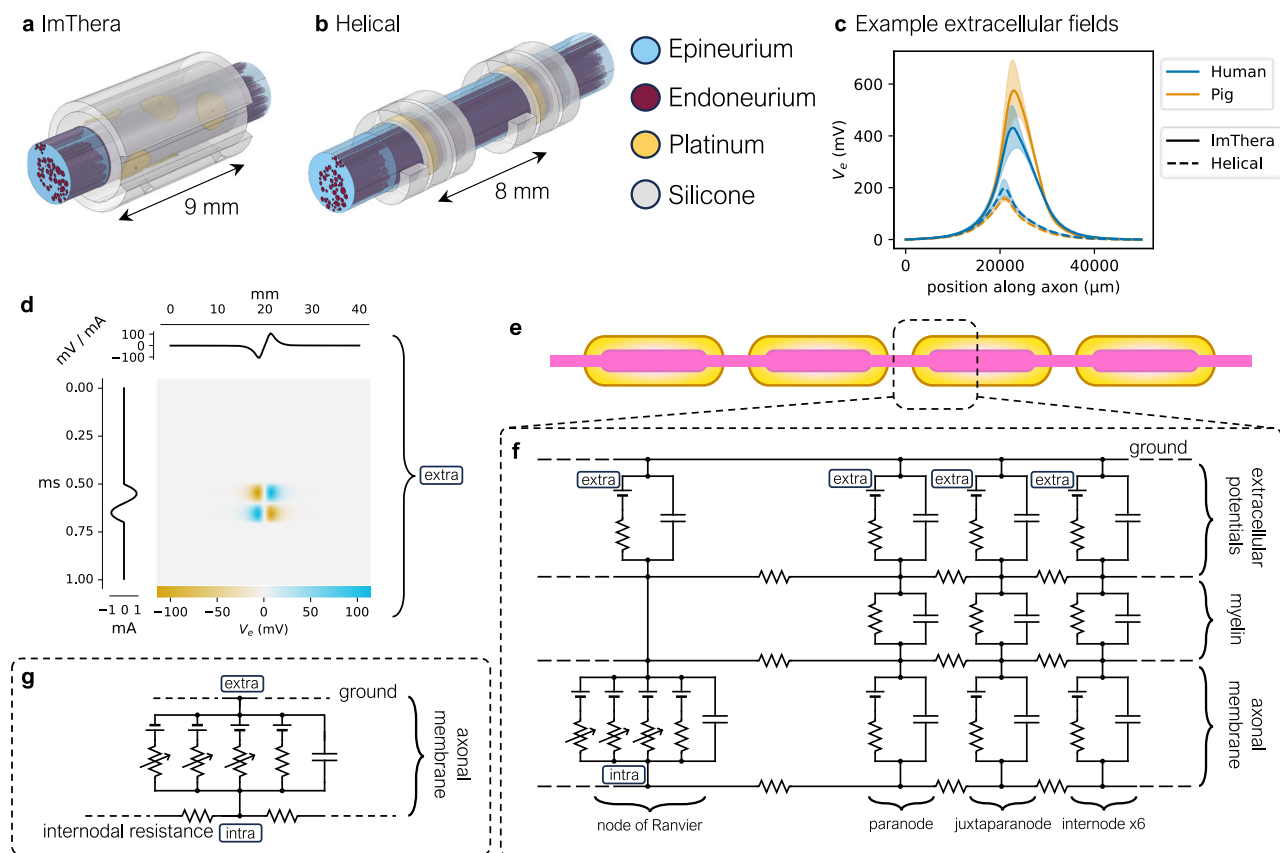


Fig. 1 | Components of computational models of electrical stimulation of peripheral nerves. ImThera (a) and helical cuff (b) geometries (with example pig vagus nerve morphology). The perineurium (side boundaries of each fascicle), muscle (tissue surrounding the nerve and cuff), and saline (fill between the nerve and cuff) are not shown. c Mean (lines) and standard deviation (shaded regions) potential distributions along the centroid of 10 example fascicles for pig (orange; Supplementary Note 1 P1) and human (blue; Supplementary Note 1 H1) vagus nerves in response to a 1 mA monopolar stimulus delivered via the ImThera (solid lines) or helical (dashed lines) cuff. d 2D array representation of spatiotemporal boundary

condition imposed by extracellular stimulation, generated by an outer product between the potential distribution along the fiber (V_A^{-1}) and the stimulus waveform (mA). e Cartoon of the longitudinal cross section of a myelinated fiber, with the axon in pink and myelin in yellow. f Double-cable structure of the MRG¹² NEURON model, with example circuit representations of the node of Ranvier, paranode, juxtapanode, and internode compartments. Inputs may include extracellular [extra] and intracellular [intra] current stimulation. g Single-cable structure of the surrogate model (S-MF) that retains only the nodes of Ranvier. As with the NEURON model, S-MF permits both extracellular [extra] and intracellular [intra] current stimulation.

larger fibers ($r = -0.21$, $p < 0.005$; Fig. 2b). Threshold errors spanned -11.0 to 7.3% (with $>95\%$ of errors within $\pm 5\%$), and they were comparable between human and pig nerve morphologies, as well as between ImThera and helical cuff geometries, although the training data used only the ImThera cuff (Fig. 2c). Threshold accuracy was consistent across all six tested waveforms, although the training data included only monophasic rectangular pulses (Fig. 2d).

For comparison, we implemented a previously published threshold estimator¹⁴ (Supplementary Note 5). The Peterson surrogate overestimated thresholds with MAPE = 31% across all tested fiber diameters and pulse widths, and in some cases, threshold error exceeded 150% (Fig. 2e).

S-MF massively reduced computation time (Fig. 2f). Calculating all 17,280 ImThera thresholds required 92 min 18 s in NEURON parallelized across 375 CPU cores but only 26 s for S-MF on a single GPU, i.e., $\sim 80,000\times$ speedup over single-core NEURON. Calculating all 25,200 helical cuff thresholds required 155 min 12 s in NEURON but only 37 s for S-MF, i.e., $\sim 95,000\times$ speedup over single-core NEURON.

Surrogate model exhibits emergent nonlinear spatiotemporal phenomena

The responses of nerve fibers to stimulation depend on states of intrinsic activity and prior history of excitation, and local excitation

does not necessarily imply action potential propagation to downstream targets. These complex interactions give rise to phenomena such as conduction block. The MRG fiber model has been used to predict high-frequency block^{36–38} and to explain non-monotonic effects of frequency on block thresholds from *in vivo* studies³⁹. S-MF accurately reproduced a variety of highly nonlinear phenomena that were not represented in the training set and were not captured by threshold estimators^{13–15,40}. See Methods “Surrogate model: Testing - other non-linear responses to stimulation” for full methodological details.

We first examined suprathreshold effects of stimulation with a single pulse. We selected one fiber location in a pig nerve model instrumented with the ImThera cuff and determined the activation threshold for cathodic monopolar stimulation with a rectangular pulse delivered from a random contact. We then stimulated with increasing suprathreshold amplitudes. S-MF reproduced diverse responses at increasing amplitudes, from excitation to unidirectional propagation, to bidirectional failure of action potential propagation, to re-excitation with a slightly delayed action potential relative to initial low-threshold excitation (Fig. 3a).

We then examined S-MF responses to kilohertz frequency stimulation, as is used to induce reversible conduction block. We selected a pig nerve model instrumented with the ImThera cuff and populated it with fibers of different diameters (5.7, 8.7, and 14.0 μm). We generated

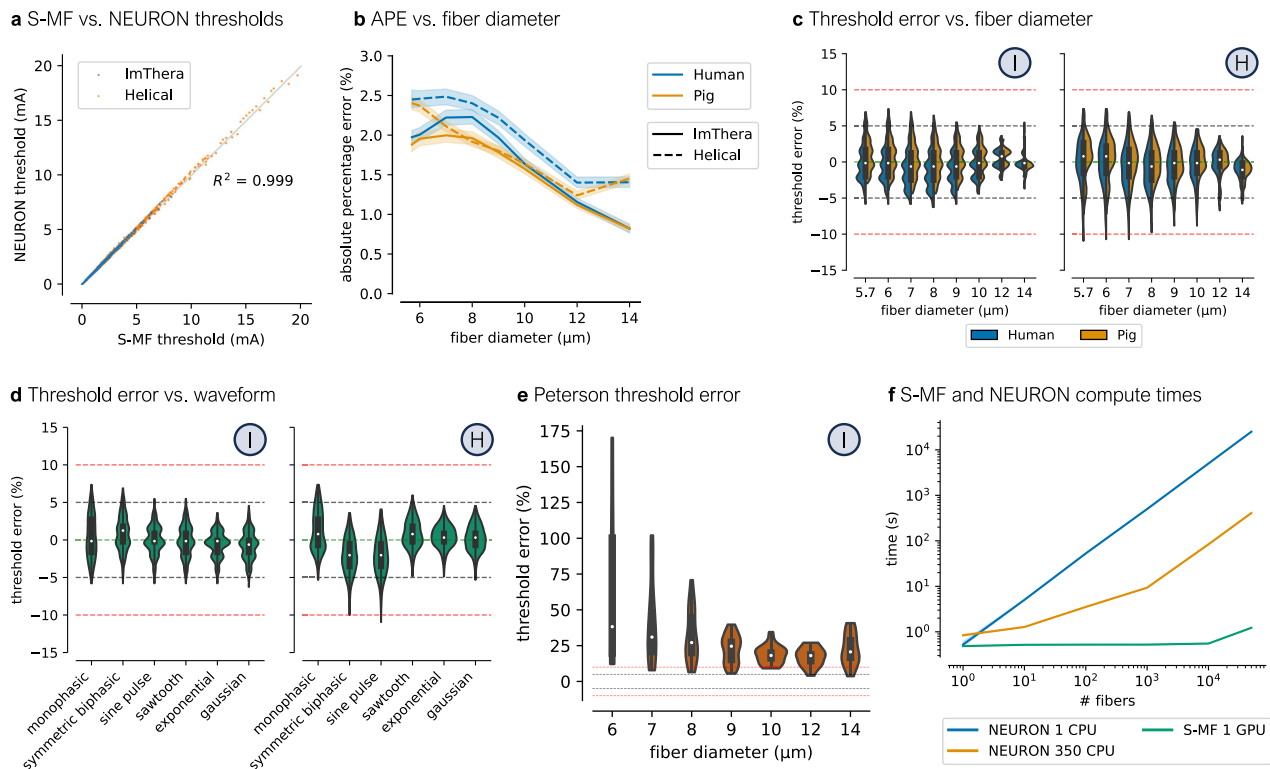


Fig. 2 | Surrogate model (S-MF) accurately predicts activation thresholds and dramatically reduces compute time. All tests were conducted in nerves other than those used to train S-MF. Data collected with monopolar ImThera and bipolar helical cuff electrodes are indicated with (I) and (H), respectively. **a** Thresholds calculated with S-MF vs. NEURON model across fiber diameters ($n = 8$), waveforms ($n = 6$), pulse widths ($n = 5$), species ($n = 2$), nerve morphologies ($n = 6$ per species), fiber locations ($n = 5$ to 10 per nerve), and cuff geometries ($n = 2$). We simulated thresholds using 375 CPU cores (Intel(R) Xeon(R) Gold 6154 CPU @ 3.00 GHz) for NEURON, and 1 CPU and 1 GPU (NVIDIA RTX 24GB A5000) for S-MF. The gray line denotes a 1:1 relationship. **b** Absolute percentage error (APE) in S-MF thresholds relative to the NEURON model across fiber diameters. Lines indicate the mean APE ($n = 1080$ per fiber diameter for ImThera cuff, $n = 1800$ per fiber diameter for helical cuff on pig nerves, $n = 1350$ per fiber diameter for helical cuff on human nerves), where the mean was calculated across waveforms, fiber locations, and nerve morphologies. Shaded regions indicate 95% confidence intervals. **c** Distributions of

percent threshold error versus fiber diameter across waveforms ($n = 6$), pulse widths ($n = 5$), nerve morphologies ($n = 12$), and fiber locations ($n = 5$ to 10 per nerve) for a given cuff. **d** Distributions of percent threshold error versus stimulus waveform across fiber diameters ($n = 8$), pulse widths ($n = 5$), species ($n = 2$), nerve morphologies ($n = 6$ per species), and fiber locations ($n = 5$ to 10 per nerve) for a given cuff. **e** Distributions of percent threshold error for the Peterson surrogate versus fiber diameter across pulse widths in a pig vagus nerve ($n = 1$) instrumented with the ImThera cuff. **c–e** Boxplots are displayed using the Tukey method (center line, median; box limits, upper and lower quartiles; whiskers, maximum / minimum point within $1.5 \times$ interquartile range of nearest hinge). **f** Compute times for different numbers of fibers (53 nodes of Ranvier, 5 ms simulation, 0.005 ms time step), simulated using NEURON or S-MF on different hardware. Compute times for single-core NEURON were extrapolated assuming linear scaling beyond 1000 fibers. Source data are provided as a Source Data file.

100 Hz intrinsic activity using intracellular current pulses at one end of each fiber. We then selected one electrode contact and delivered sinusoidal extracellular signals at a range of frequencies and amplitudes and recorded the number of action potentials arriving over a 100 ms period at the other end of each fiber opposite the site of intracellular stimulation. Kiloherz frequency extracellular signals generate complex interactions with ongoing neural activity across signal frequencies, amplitudes, and fiber diameters, including excitation, partial block, complete block, re-excitation, and re-block, and these responses were captured by S-MF (Fig. 3b).

Next, we examined S-MF's ability to represent interactions between propagating action potentials. We generated action potentials at opposite ends of fibers of various diameters (5.7, 8.7, 10.0, and 14.0 μm) using intracellular current injection. We recorded membrane potential along the length of the fibers over time. S-MF accurately reproduced action potential annihilation by collision for multiple fiber diameters (Fig. 3c).

Finally, we examined S-MF's ability to represent interactions between intrinsic firing and stimulation at "conventional" (sub-kHz) frequencies. We selected a human nerve model instrumented with the helical cuff and populated the nerve with fibers of different diameters

(5.7, 8.7, and 14.0 μm) at the centroids of randomly selected fascicles. We generated random intrinsic activity at different mean rates (10, 50, 100 Hz) in each fiber through intracellular current injection at one end. We simultaneously delivered bipolar extracellular stimulation with 0.1 ms biphasic pulses delivered at different frequencies (30, 50, and 100 Hz) at amplitudes ranging from 50% below to 50% above the mean threshold for each fiber diameter. We recorded the timing of spikes arriving at the end of each fiber opposite from the site of current injection and compared the spike times with and without extracellular stimulation. S-MF accurately predicted spiking desynchronization (i.e., changes in intrinsic spiking activity caused by extracellular stimulation) across mean intrinsic firing rates, frequencies of extracellular stimulation, stimulation amplitudes, fiber diameters, and electrode-fiber distances (Fig. 3d). For kiloherz-frequency and spike synchronization modeling tasks, S-MF produced up to 90,000 \times speedup over single-core NEURON (Supplementary Note 8.1).

Surrogate model greatly accelerates optimization for spatial selectivity

Given an appropriate model of a nerve's anatomy and biophysics, optimization methods can define stimulation parameters to achieve a

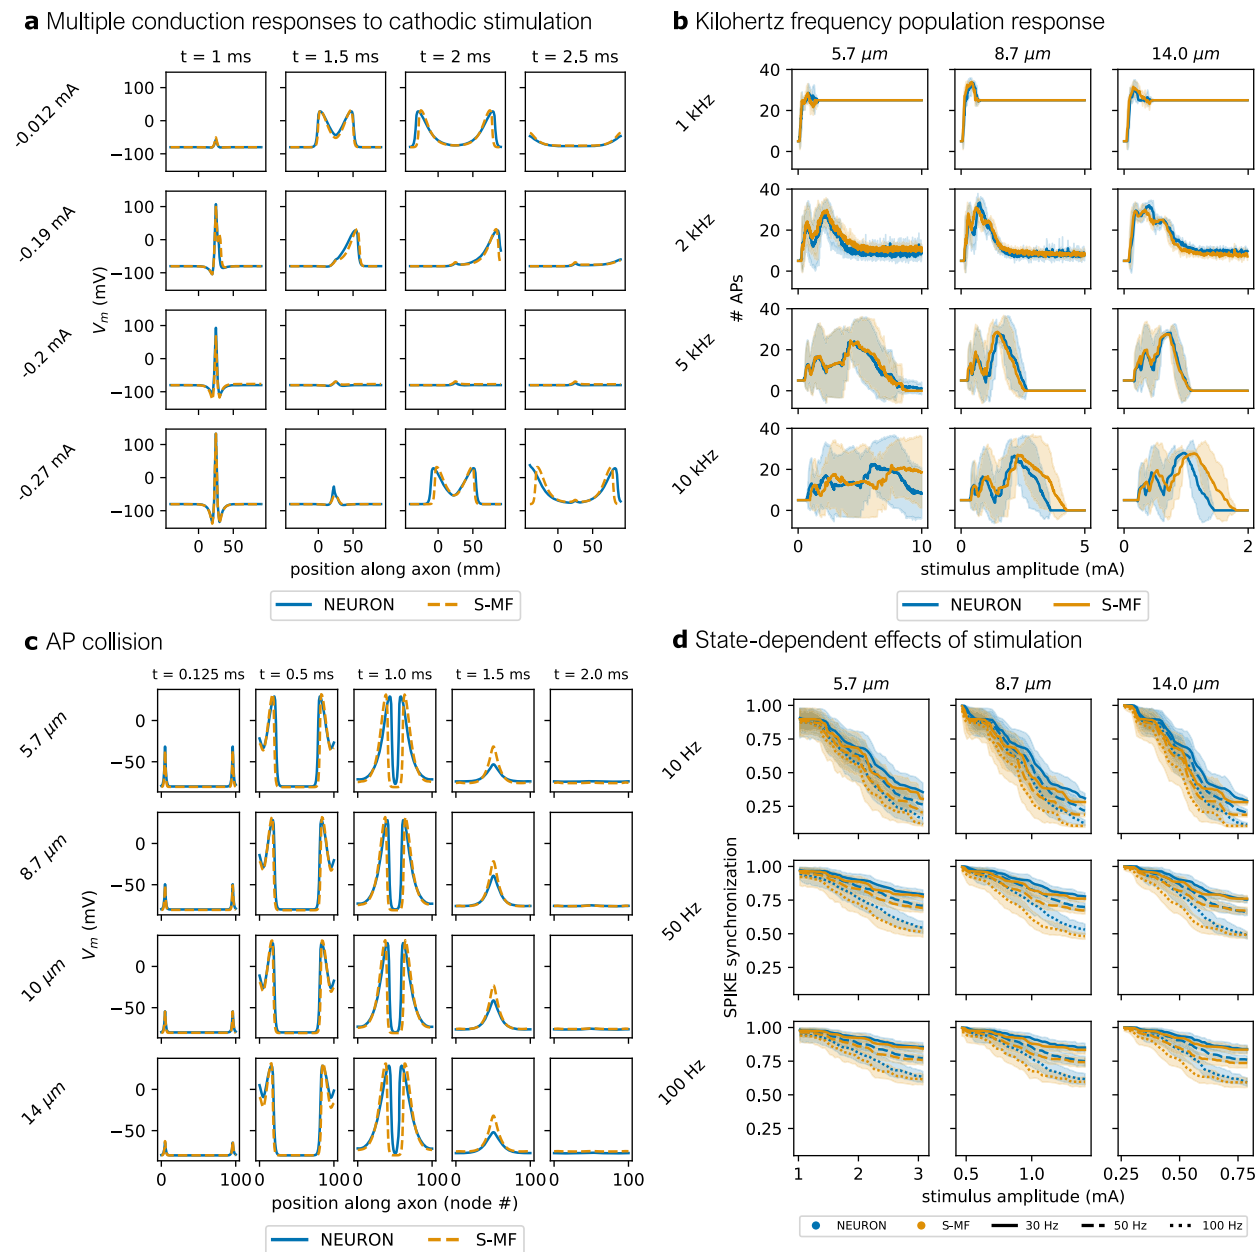


Fig. 3 | Surrogate model (S-MF) predicts a range of nonlinear responses to stimulation. All tests were conducted in nerves other than those used to train S-MF. **a** S-MF (dashed orange) reproduced NEURON (solid blue) responses to a 0.75 ms monopolar cathodic extracellular pulse delivered to pig vagus nerve (Supplementary Note 1 P3) with the ImThera cuff across different stimulation amplitudes (rows). Columns indicate the timepoints during each simulation at which transmembrane potential (V_m) is plotted. **b** S-MF (orange) reproduces the NEURON (blue) responses to kilohertz frequency extracellular stimulation with intrinsic activity. Fibers were simulated in a pig vagus nerve (Supplementary Note 1 P2) instrumented with the ImThera cuff. Intrinsic activity was generated at 100 Hz with intracellular current pulses at one end of each fiber and action potentials (APs) were counted at the opposite end. Rows correspond to the frequency of the extracellular sinusoidal signal and columns correspond to fiber diameters. Solid lines are the mean number of APs recorded across 7 sampled fascicles, and shaded regions are 95% confidence intervals. **c** S-MF (dashed orange) reproduces NEURON

(solid blue) AP interactions. Action potentials initiated by intracellular current pulses (2 nA, 0.1 ms) at opposite ends of a fiber propagate bidirectionally at appropriate speeds and mutually annihilate when they meet. Rows correspond to different fiber diameters, and columns correspond to the timepoints at which V_m along the nerve is plotted. **d** S-MF (orange) reproduces state-dependent fiber responses to extracellular stimulation predicted by NEURON (blue). Fibers were placed in a human nerve (Supplementary Note 1 H2) instrumented with the helical cuff. Rows correspond to mean intrinsic firing rates, and columns correspond to fiber diameters. Lines correspond to mean SPIKE synchronization⁶⁴ (a quantity between 0 and 1 that measures the proportion of coincident spikes between two spike trains) across all modeled fibers of a given diameter and mean intrinsic firing rate ($n = 35$, 7 fiber locations \times 5 distinct patterns of intrinsic firing) at a given amplitude and frequency (line style) of extracellular stimulation (0.1 ms symmetric biphasic rectangular pulses). Shaded regions correspond to 95% confidence intervals. Source data are provided as a Source Data file.

targeted neural response, such as activation of specific fibers while avoiding activation of other fibers. We examined this selective stimulation task under various constraints.

We instrumented each vagus nerve ($n = 6$ human and $n = 6$ pig) with a six-contact ImThera cuff and modeled a 5.7 μm diameter fiber at

the centroid of each fascicle. Vagal motor fibers pass by the nodose ganglion, where the cell bodies of visceral afferent sensory fibers are located. In pig models, we used published histology⁴¹ to assign fibers as target or off-target based on this sensory / motor distinction. Analogous immunohistochemical data are not yet available for the human

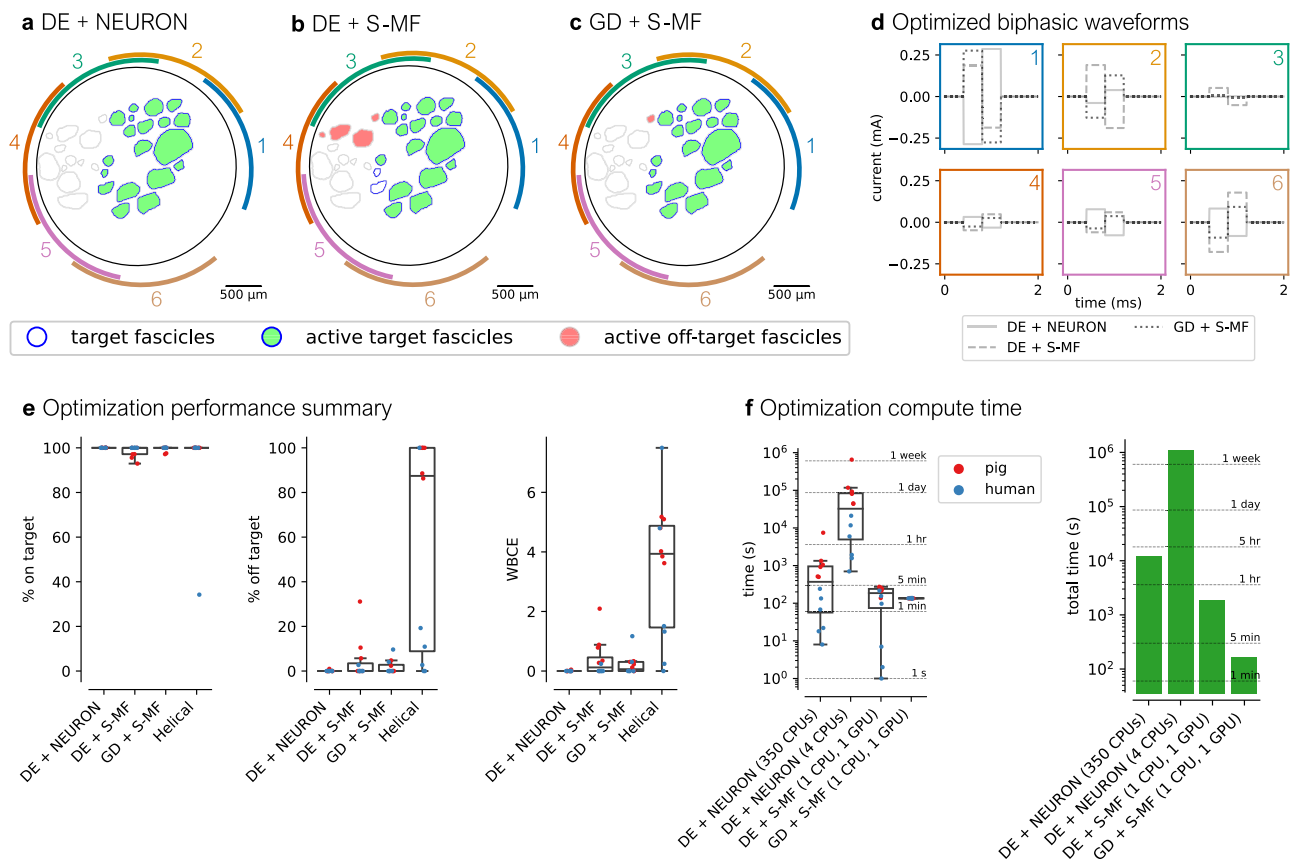


Fig. 4 | Optimization of current amplitudes (6 degrees of freedom) for spatially selective activation of 5.7 μm fibers using symmetric biphasic rectangular waveforms. **a–c** Example selectivity performance for a pig vagus nerve (Supplementary Note 1 P4) using gradient-free optimization (differential evolution; DE) and NEURON (**a**), gradient-free optimization and the surrogate model (S-MF) (**b**), and gradient descent (GD) with S-MF (**c**). **d** Final optimized stimulation waveforms corresponding to panels **a–c**. Frame colors and numbers correspond to contact colors and numbers in panels **a–c**. **e** Summary of optimization performance across all 12 nerves (6 pig vagus nerves in red and 6 human vagus nerves in blue) for all optimization methods: % target fascicular area activated, % off-target fascicular area activated, and weighted binary cross entropy (WBCE, Eq. (18)). We assumed that all

fibers of a specific diameter within a given fascicle had the same threshold⁷⁸; we thus converted percent fibers activated to percent fascicular area activated. For all optimized stimulus parameters generated by algorithms using S-MF (DE + S-MF, GD + S-MF), we evaluated the ‘true’ fiber activations using the NEURON model and report performance metrics for those activations (Supplementary Note 7). **f** Compute time for each optimization method on applicable hardware, per nerve (left, $n = 12$) and total (right). Compute time for 4 CPUs was estimated from 350 CPUs assuming linear scaling. **e, f** Boxplots are displayed using the Tukey method (center line, median; box limits, upper and lower quartiles; whiskers, maximum/minimum point within $1.5 \times$ interquartile range of nearest hinge). Source data are provided as a Source Data file.

vagus nerve; therefore, we assumed a similar vagotomy to pigs, and we randomly assigned one-half of the nerve as target and the other half as off-target.

We applied our optimization algorithms to two tasks: amplitudes of symmetric rectangular current pulses and arbitrary waveforms. First, we limited stimulation to symmetric biphasic rectangular waveforms with a fixed 400 μs pulse width, yielding a six-dimensional optimization problem for each nerve, i.e., the optimization algorithm was tasked with identifying the amplitude of the rectangular waveform for each of the six electrode contacts. Second, each contact was permitted to deliver simultaneously a 1 ms charge-balanced signal with an arbitrary waveform, yielding a 1200-dimensional optimization problem for each nerve, i.e., the optimization algorithm was tasked with identifying the current delivered at 200 consecutive time points for each of the six electrode contacts.

The computational demands to conduct optimization using CPU-based biophysical models can be prohibitive for many users and applications. Further, typical black-box optimization methods may not perform well for high-dimensional problems (e.g., waveform-agnostic optimization). We developed two optimization strategies (gradient-free and gradient-based) that used S-MF in lieu of a standard

biophysical model to address both the computational demands and performance on problems with many degrees of freedom. All reported selectivity results for optimized stimulus parameters were evaluated in NEURON. Supplementary Table 3 provides a summary of optimization performance with respect to waveform constraint (biphasic or arbitrary), optimization algorithm (gradient-free or gradient-based), and fiber model (NEURON or S-MF).

We optimized the amplitudes of rectangular pulses delivered via the six electrode contacts to achieve spatially selective fiber activation (Fig. 4a–d). Comparable selectivity was achieved across nerve morphologies, optimization methods, and fiber models (Fig. 4e). Optimization with S-MF yielded consistently high-quality solutions, with 100% median target activation and 0% median off target activation for both gradient-free optimization (differential evolution; DE) and gradient descent (GD). Relative to single-CPU execution time for DE with the NEURON model, S-MF achieved $\sim 2,300\times$ speedup when using DE (single GPU) and $\sim 26,000\times$ when using GD (single GPU) (Fig. 4f; Supplementary Note 8.2)

We used the weighted binary cross entropy (WBCE)—a weighted distance between two sets of binary probability vectors—to assess selectivity (Fig. 4e), where larger WBCE corresponds to poorer

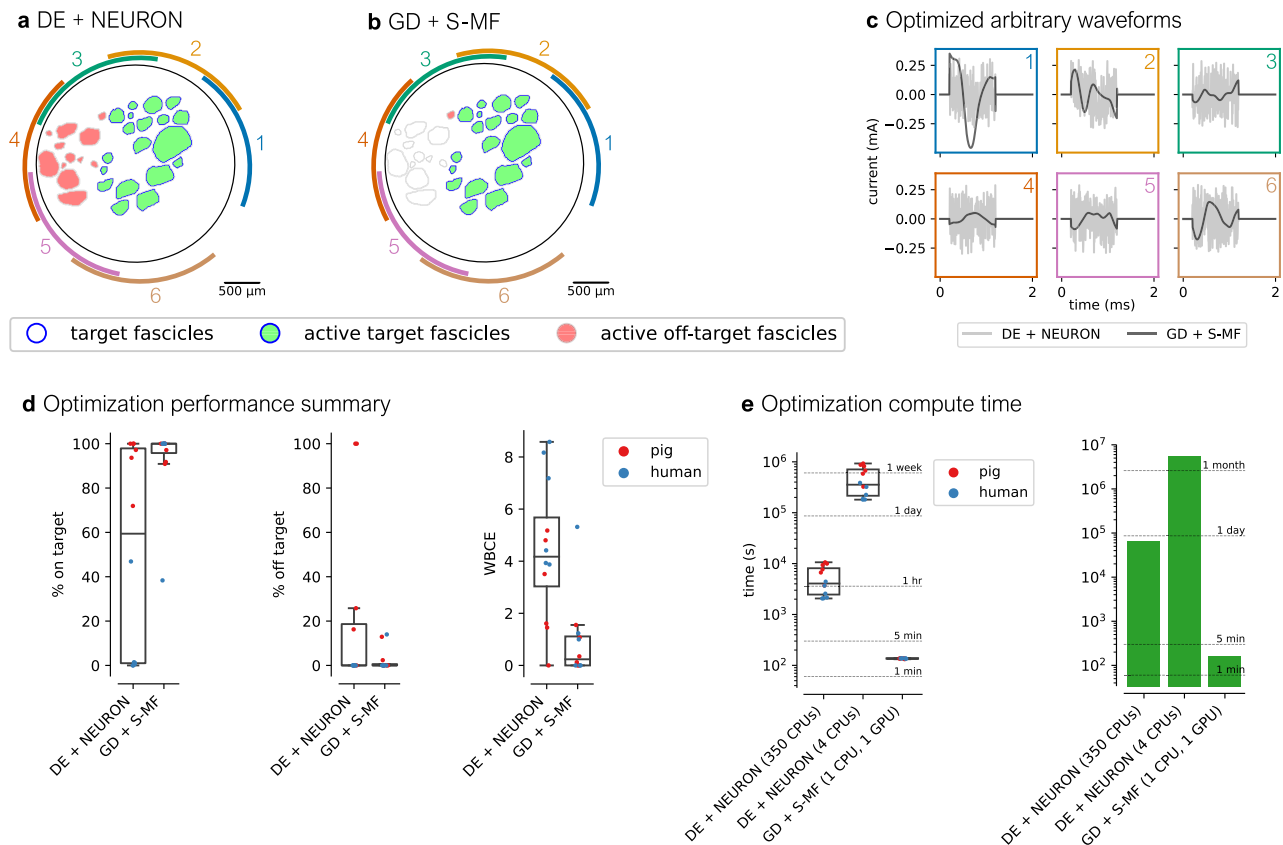


Fig. 5 | Optimization of current amplitudes (1200 degrees of freedom) for spatially selective activation of 5.7 μm fibers using arbitrary waveforms. **a–b** Example selectivity performance (Supplementary Note 1 P4) using gradient-free optimization (differential evolution; DE) and NEURON (**a**), and gradient descent (GD) with the surrogate model (S-MF) (**b**). **c** Final optimized stimulation waveforms corresponding to panels **a**, **b**. The total pulse duration for every contact for all optimizations was constrained to 1 ms (between 0.2 and 1.2 ms). Frame colors and numbers correspond to contact colors and numbers in **a**, **b**. **d** Summary of optimization performance using arbitrary waveforms across all 12 nerves (6 pig vagus nerves in red and 6 human vagus nerves in blue) for both optimization methods: % target fascicles activated, % off target fascicles activated, and weighted

binary cross entropy (Eq. (18)). We assumed that all fibers of a specific diameter within a given fascicle had the same threshold⁷⁸; we thus converted percent fibers activated to percent fascicular area activated. For all optimized stimulus parameters generated by algorithms using S-MF (GD + S-MF), we evaluated the ‘true’ fiber activations using the NEURON model and report performance metrics for those activations (Supplementary Note 7). **e** Compute time for each optimization method on different hardware, per nerve (left, $n = 12$) and total (right). Compute time for 4 CPUs was estimated from 350 CPUs assuming linear scaling. **d**, **e** Boxplots are displayed using the Tukey method (center line, median; box limits, upper and lower quartiles; whiskers, maximum / minimum point within 1.5 \times interquartile range of nearest hinge). Source data are provided as a Source Data file.

selectivity. We detected no significant differences in WBCE for solutions generated using either DE with S-MF (Wilcoxon signed-rank test; $p = 0.03 > 0.0125$) or GD with S-MF (Wilcoxon signed-rank test; $p = 0.02 > 0.01$) compared to DE with NEURON. While optimization with the six-contact ImThera cuff achieved spatial selectivity, the standard clinical cuff has helical contacts spanning most of the nerve circumference; optimization with the helical cuff activated target fibers (100% median activation), but with strong concomitant activation of off-target fibers (87% median activation).

We optimized the shape of arbitrary, charge-balanced waveforms delivered through each of the six contacts to achieve spatially selective fiber activation (Fig. 5). The DE and NEURON-based optimization was not successful in identifying selective stimulus parameters, consistent with the limitations of randomization-based evolutionary algorithms in searching high-dimensional spaces⁴². In contrast, gradient-based optimization consistently achieved selectivity across all nerves, with 100% median target activation and 0% median off-target activation, but selectivity was worse than with rectangular pulses (Wilcoxon signed-rank test of WBCE between biphasic waveforms generated using DE + NEURON and arbitrary waveforms generated using GD + S-MF yielded $p = 0.003 < 0.006$) likely due to additional degrees of

freedom increasing the probability of generating waveforms for which S-MF prediction is less accurate (Supplementary Note 7). GD + S-MF yielded a total speedup of $\sim 135,000\times$ over single-core DE + NEURON (Supplementary Note 8.3). Additionally, the gradient descent algorithm consistently generated smooth waveforms, in contrast to the stochastic gradient-free procedure that consistently returned noisy, random-looking waveforms (Fig. 5c).

Discussion

Electrical stimulation of the nervous system has been used for a broad range of applications, including VNS for epilepsy, depression, obesity, type 2 diabetes, heart failure, and rheumatoid arthritis; deep brain stimulation for Parkinson’s disease, essential tremor, and epilepsy; spinal cord stimulation for chronic pain; sacral nerve stimulation for incontinence; and carotid sinus nerve stimulation for hypertension. These therapies have a current estimated total global market of $\sim \$10$ billion. However, there currently exists only one simulation platform (NEURON) that directly represents the effects of electric fields on neurons, with some significant associated computational challenges; therefore, there is a need for high-quality estimation and optimization tools.

We introduced a framework, AxonML, for defining surrogate models of nerve fibers that dramatically reduces computational demands while preserving key biophysical attributes by (1) simplifying the fiber geometry, (2) developing a GPU-compatible implementation, and (3) tuning the model parameters to maintain accuracy relative to the original biophysical model. The surrogate preserves the dynamic behavior at the nodes of Ranvier where action potentials are initiated and propagated by nonlinear voltage-gated ion channels, including the full spatiotemporal response of the gating variables, transmembrane currents, and transmembrane potential. We defined a surrogate model to mimic the MRG fiber model; we chose the MRG fiber model because it is the “industry standard”, it is highly relevant to the development of peripheral nerve neuromodulation therapies, it balances complexity and accuracy, it was developed based on extensive experimental data from mammalian myelinated fibers, and it reproduces *in vivo* thresholds of rat, pig, and human VNS^{26,35}. AxonML can be applied to implement other neuronal models; for example, we also used the framework to create a surrogate of the Sundt²⁹ C-fiber model, an unmyelinated fiber model-of-choice for computational studies²² (Supplementary Note 10).

Our surrogate model (S-MF) scales efficiently on multithreaded hardware, like the RTX A5000 24GB GPU, and trains effectively on large spatiotemporal datasets. This GPU compatibility is an important advantage given that existing platforms like NEURON lack GPU support for solving differential algebraic equations inherent to extracellular potential boundary conditions. GPU deployment yielded a 4 to 5 orders-of-magnitude speedup over single-core NEURON, where the speedup depends on the problem size and parallelism factors (i.e., the number of fiber simulations, the length of each fiber, and the proportion of simulations that can be performed in parallel). Specifically, GPU execution enables the concurrent solution of several thousand fiber models, achieving concurrency both at the population level, where multiple fibers are processed simultaneously, and at the individual level, allowing parallel computation of state variable updates for all nodes within a single fiber. This approach significantly enhances scalability (Fig. 2f): for example, as shown in Fig. 4f, optimizing stimulation parameters only required ~2 min for either 1 nerve (whether ~10 fibers for a human nerve or ~50 fibers for a pig nerve) or 12 nerves. GPUs consistently outperform CPUs in both cost-per-gigaflop and performance-per-watt metrics^{43,44}; securing equivalent computational performance with CPUs via NEURON would be financially burdensome and would necessitate dedicated management of resources (for both maintenance and job allocation). Conversely, GPUs comparable to the RTX A5000 are commercially available and can be readily integrated into standard desktop computers, thus circumventing the need for elaborate, expensive infrastructure to tackle large-scale simulation and optimization problems.

Training a surrogate efficiently and effectively requires consideration of dataset generation, dataset consumption, and gradient descent configuration. Generating training datasets requires solving costly finite element models of the peripheral nerve and several thousand biophysical fiber models, which is made practical by access to a high-performance computing platform. Storing the training data requires moderate hard drive space (~100 GB), and training requires efficient streaming of data from the hard drive, depending on the available RAM. We found that training with the Adam⁴⁵ or RMSprop⁴⁶ optimizer generated surrogates of comparable quality. Training was also sensitive to the length of the chunks of data over which backpropagation was performed: 50 timesteps were suitable in our case, but this may vary depending on the system being approximated. These considerations, however, do not apply at the time of inference; we provide our trained surrogate publicly for others to use without having to train their own surrogate model.

S-MF greatly expands the types of waveforms and stimulation protocols that can be predicted as compared to threshold-

activation-based estimators, thus providing broader insight into how information transmitted along fibers is transformed by electrical stimulation. Prior efforts were exclusively designed to predict activation thresholds from rest for a specific class of waveform^{13–15,40}. In contrast, S-MF can represent states of intrinsic activity, effects of intracellular current injection, and history of excitation, while being agnostic to waveform shape, even though the training data featured only fiber responses to simple rectangular pulses delivered extracellularly via a cuff electrode. Despite this limited training set, S-MF accurately reproduced emergent non-linear responses, including responses to high-frequency signals, direct current (DC) block, and interactions between intrinsic activity and stimulation. Further, S-MF achieved rapid and accurate optimization across complex design spaces, where traditional computational approaches are intractable.

We introduce a gradient descent-based stimulus parameter optimization algorithm using S-MF, in addition to an implementation of a gradient-free population-based differential evolution algorithm. Both optimization strategies using S-MF preserved accuracy while significantly reducing run time compared to traditional metaheuristic methods coupled with a NEURON model. The gradient descent approach scaled especially well as it is straightforward to parallelize on a single GPU. In gradient-free methods, each optimization step requires evaluating a set of candidate solutions (“population”) for every problem, cumulatively requiring a total number of fiber simulations that equals the product of the population size and the number of fibers in the nerve model. In contrast, with gradient descent, we need only evaluate and backpropagate through the number of fibers in the nerve model. This difference becomes increasingly significant for larger nerve models and larger population sizes, where a single instance of gradient-free optimization can consume all available GPU memory.

We applied our optimization algorithms to spatially selective stimulation. *In vivo* studies demonstrate spatially selective neuromodulation using multi-contact stimulation on the pig cervical vagus nerve for activation of efferent fibers⁴⁷, cat sciatic nerve for control of muscle activation^{48,49}, and human femoral nerve in an intraoperative setting for selective hip and knee flexion and extension⁵⁰. In our first set of optimizations, we constrained stimulation to fixed-duration biphasic rectangular pulses. This is consistent with stimulation delivered by standard implantable pulse generators. However, some commercial stimulators, such as the STG 4008 (Multi Channel Systems, Reutlingen, Germany) support delivery of continuous arbitrary waveforms^{51,52}, and recently, experimental system-on-chip platforms have been developed to deliver fully arbitrary biomimetic signals⁵³. We therefore also conducted a set of experiments where arbitrary waveforms were permitted. The design of arbitrary waveforms demonstrated the performance of S-MF and different optimization algorithms in a high-dimensional parameter space, where prior methods are very computationally expensive and/or ineffective.

Across all 36 optimization trials using S-MF, optimized solutions activated 97.1 ± 10.2 % of the target fascicular area (median 100%, IQR 97.5–100 %) and only 2.8 ± 6.1 % of the off-target fascicular area (median 0%, IQR 0–2.5 %), as evaluated using the NEURON model. S-MF also was highly accurate in predicting whether fibers were activated or not in response to the 36 optimized waveforms (accuracy of 98.0 ± 5.1 %; median 100%, IQR 100–100%) (Supplementary Note 7). Further, we evaluated the neural responses at every 10th step of the gradient-based optimization, and S-MF’s accuracy remained high (Supplementary Note 8). These results indicate that our surrogate model and optimization approach are not biased to poor quality solutions or to regions of low surrogate accuracy.

Our optimization results consistently predicted stimulus configurations of multi-contact electrodes that selectively activated spatially localized fibers. However, the optimized parameters were not unique: the various optimization methods found solutions with comparable

selectivity, but with different numerical values (Supplementary Note 6). Optimizing arbitrary waveform shapes is challenging given the very large degrees of freedom compared to, e.g., simply optimizing stimulation amplitude across contacts for fixed waveform shapes and timing. Nonetheless, gradient descent successfully identified smooth arbitrary waveforms for spatially selective activation (where conventional optimization using NEURON struggled). Multi-contact electrodes for peripheral nerve stimulation will soon be more readily available and common, and thus these methods will be important in improving and accelerating neuroelectronic interface design. We illustrated results for charge-balanced stimulation to achieve spatial selectivity in the vagus nerve with a specific six-contact cuff electrode, but the surrogate model and optimization methods could be applied to any nerve and optimization criteria.

Our surrogate fiber model and optimization approaches efficiently and dramatically reduce the scale of design problems. For example, these methods can quickly generate a selection of high-quality candidate solutions, or they can inform the prioritization and exclusion of different regions of the candidate design space. After this initial reduction, traditional biophysical models can be applied for further analysis at a significantly reduced computational cost. Additionally, using the NEURON model to validate the most promising candidate solutions can maintain higher accuracy when using gradient-free optimization techniques with the surrogate model. Such hybrid methods have shown promise for model-based optimization in other physical domains⁵⁴.

Limitations

Users of S-MF (or any surrogate) must be aware of the scope over which its accuracy has been evaluated, as discussed further below, including the range of training and testing data, the complexity of an optimization task, and validation against *in vivo* data. However, regardless of this specific surrogate model's limitations, our approach and architecture can be leveraged if more accurate target responses become available, either to retrain the model presented herein or to develop a different fiber model.

Successful application of the surrogate model must consider accuracy in scenarios that are outside the range of the training data. S-MF performed well on a range of stimulation waveforms, stimulation protocols (e.g., kilohertz frequency signals and interactions between action potentials), human and pig vagus nerve morphologies, and cuff electrodes outside the training data (Figs. 2, 3). Slightly larger threshold errors occurred with the helical cuff and biphasic waveforms (symmetric rectangular pulse and single period of a sinusoid; Fig. 2d): S-MF slightly underestimated the effect of an anodic second phase in increasing activation threshold in these cases that were not in the training data. S-MF also generalized well to a model of the rat vagus nerve (Supplementary Note 9), which is monofascicular with a diameter $\sim 10\times$ smaller than pig or human vagus nerves. S-MF accuracy decreased for fiber diameters outside of the training range (5.7 to 14 μm), particularly smaller diameter fibers (3 to 4 μm) for which the geometrical parameters are defined by different equations than used for the training data (Supplementary Note 9). Extensive validation of S-MF has been undertaken for cuff-based stimulation paradigms; however, performance may not generalize as well to other conditions such as intraneural stimulation. Additional training data could be used to improve the model's accuracy for other such specialized tasks, such as the effects of electrode design on stimulation responses in a specific species and nerve.

S-MF was slightly less accurate in its predictions when optimizing arbitrary waveforms (median accuracy 100%, IQR 97.4–100%), i.e., an optimization task with a large number of degrees of freedom, compared with optimizing biphasic rectangular waveforms (median accuracy 100%, IQR 100–100%) and thus yielded slightly reduced

selectivity. However, optimization with S-MF generated high-quality solutions (100% median target activation and 0% median off-target activation), and S-MF retained accuracy during optimization even with arbitrary waveforms. We nonetheless recommend validating the results of any surrogate-based optimization against a ground-truth model, especially if performing high-dimensional optimization.

The MRG fiber model that we used as a basis for S-MF has been used to analyze and design neuromodulation therapies that were outside of the scope of the MRG's initial development, including novel energy-efficient stimulation waveforms^{23,24} and kilohertz frequency block²⁵, and there was a strong match between the model and experiments. However, the development and use of any fiber model—whether using our GPU-based approach, NEURON, or another implementation platform—must consider whether the model is sufficiently validated for a given application, or whether additional mechanisms or parameter values must be considered.

Our optimization of stimulation parameters assumed that we have exact knowledge of the nerve's morphology. However, this level of detail is unavailable in clinical settings. The enhanced computational throughput made possible by the surrogate model enables efficient optimization across a range of nerve morphologies, thus accounting for the uncertainties in real-world applications.

Methods

Field models: pig and human vagus nerve stimulation

We implemented and simulated finite element models of pig and human vagus nerves with cuff electrodes using ASCENT v1.2.1 (DOI: 10.5281/zenodo.7627427)⁵⁵, with COMSOL Multiphysics® v5.6 (COMSOL, Inc.; Burlington, MA, USA).

We modeled realistic morphologies of pig ($n=7$) and human ($n=7$) cervical vagus nerves using segmented histology of nerve cross sections^{56–58} that we extruded 50 mm longitudinally (Supplementary Note 1). We made each nerve circular while preserving its cross-sectional area using ASCENT's deformation feature; fascicles were repositioned during deformation to maintain at least 10 μm between fascicles and between fascicles and the nerve boundary.

We modeled two cuff electrodes (Fig. 1). First, we modeled the six-contact ImThera cuff electrode (LivaNova PLC, London, UK). The cuff was 9 mm in length, with an inner diameter of 3 mm and six circular contacts (2 mm diameter each) arranged along two diagonals (Fig. 1a). The ImThera cuff geometry and validated models of pig vagus nerve stimulation (VNS) with the ImThera cuff are published in prior papers^{35,59}. We also modeled the bipolar helical cuff used to treat epilepsy and depression clinically (LivaNova PLC, London, UK) (Fig. 1b), which is available in two sizes. For nerves with a diameter under 3 mm ($n=11$), we used the 2 mm diameter cuff; for larger nerves, we used the 3 mm diameter cuff ($n=3$). In its unexpanded state, the helical metal contacts of the cuffs spanned 338.5 degrees for the 2-mm cuff and 346.5 degrees for the 3-mm cuff. The contacts were positioned 8 mm apart, measured from their centers. Further details on the helical cuff's design and validated models of human VNS using this cuff are previously published⁶⁰. We used the ImThera cuff to generate training data (see "Surrogate model: Training"), and subsequent tests were performed using both the ImThera and helical cuffs.

For both cuff geometries, we positioned the cuff at the midpoint of the nerve's length. Spaces between the cuff and nerve were filled with saline. We modeled a 10 μm saline gap between the ImThera cuff and the nerve and a 100 μm saline gap between the helical cuff and the nerve. We placed the nerve and cuff in a cylinder of muscle (5 mm diameter, 50 mm length). We rotated and shifted each cuff based on the size and position of each nerve's fascicles. Specifically, we determined the centroid of all fascicles ("fascicle centroid") by applying a weighted average (using each fascicle's area) to all individual fascicle centroids, and then found the closest point on the nerve boundary to

that fascicle centroid. We then rotated the center of the cuff to align with the closest point. The cuff was then shifted towards the nerve to leave a 10 μm (ImThera) or 100 μm (helical) gap between the cuff and nerve.

We assigned materials to the geometry as shown in Fig. 1; see Supplementary Note 2 for material and tissue conductivities. Some human nerves feature “peanut fascicles”, where multiple endoneurial bundles (inner perineurium boundaries) are within a single outer perineurium boundary, separated by a perineurium septum (e.g., Supplementary Note 1 H3). We meshed the perineurium for all peanut fascicles, otherwise we modeled the perineurium using a thin layer approximation:

$$\rho_{\text{surface}} = \frac{\text{thk}_{\text{peri}}}{\sigma_{\text{peri}}} \quad (1)$$

where ρ_{surface} is the contact impedance of the perineurium boundary (Ωm^2), σ_{peri} is the conductivity of the perineurium (S m^{-1}), and thk_{peri} is the thickness of the perineurium (m). For human nerve morphologies, we calculated the perineurium thickness for each fascicle based on the original segmentations:

$$\text{thk}_{\text{peri}} = r_{\text{outer}} - r_{\text{inner}} \quad (2)$$

where r_{outer} , r_{inner} are the effective circular radii of the outer and inner perineurium boundaries, respectively. For pig morphologies, we applied a published relationship between fascicle diameter, d_{fasc} (μm), and perineurium thickness for the pig vagus nerve⁵⁶:

$$\text{thk}_{\text{peri}} = 3.44 \mu\text{m} + 0.02547 \times d_{\text{fasc}} \quad (3)$$

We solved Laplace’s equation:

$$\nabla \cdot (\sigma \cdot \nabla V) = 0 \quad (4)$$

using quadratic geometry and solution shape functions once for each electrode contact in each model. When solving for each contact, we applied a 1 mA point source of current in the center of the electrode contact, a floating potential boundary condition for all other electrode contacts, and grounded outer boundaries of the model. Pig finite element models (FEMs) had 35 to 95 million free tetrahedral elements, and human FEMs had 2 to 35 million free tetrahedral elements.

Surrogate model

Design and implementation. The surrogate model (S-MF, Fig. 1g) is a simplified version of the McIntyre-Richardson-Grill (MRG) model of a mammalian myelinated fiber¹² (Fig. 1f) that calculates the transmembrane potential (V_m) and all gating variables (m, h, p, s) at every node of Ranvier and time point. S-MF retains only the nodes of Ranvier and the intracellular space connecting nodes, thus assuming perfectly insulating myelin and omitting additional extracellular field layers. S-MF, like the original MRG model, assumes that the fiber is in a large, highly conductive medium (i.e., equivalent circuit nodes representing extracellular voltage are grounded and have nonzero potential only when an extracellular stimulus is applied). The inputs to S-MF are a 2D array defining the extracellular electric field in space (along the length of the fiber, at each node of Ranvier) and in time (according to the chosen stimulation waveform) (Fig. 1d), the diameter of the fiber, and optionally, any intracellular current sources.

We solved the system of differential equations of the surrogate model using forward Euler discretization and a staggered timestep. To advance the simulation by one timestep, we first update the gating parameter values from time $t - 0.5dt$ to $t + 0.5dt$ using the analytic

solution:

$$x^{n,t+0.5dt} = x_{\infty}(V_m^{n,t}) - \left(x_{\infty}(V_m^{n,t}) - x^{n,t-0.5dt} \right) e^{-dt/\tau_x(V_m^{n,t})} \quad (5)$$

$$x_{\infty}(V_m) = \frac{\alpha_x(V_m)}{\alpha_x(V_m) + \beta_x(V_m)} \quad (6)$$

$$\tau_x(V_m) = \frac{1}{q_{10}^x \times (\alpha_x(V_m) + \beta_x(V_m))} \quad (7)$$

$$q_{10}^x = a q_{10}^x \frac{T_c - \text{bq}_{10}^x}{\text{cq}_{10}^x} \quad (8)$$

where n is the index of the node of Ranvier, $V_m^{n,t}$ is the transmembrane potential for node n at time t , and α_x and β_x are independent rate constant functions, as defined in the MRG paper¹², for gating variable x (where x is m, h, p , or s). q_{10}^x is the q_{10} value for gating variable x , parameterized by scalars $a q_{10}^x, \text{bq}_{10}^x, \text{cq}_{10}^x$, where $T_c = 37^\circ\text{C}$ is the temperature. Using m to indicate $m^{n,t+0.5dt}$ (likewise for all other gating variables h, p , and s), we then calculate the ionic current and advance the transmembrane potential from time t to $t + dt$:

$$I_{\text{ion}}^{n,t} = \bar{g}_{\text{NaF}} m^3 h (V_m^{n,t} - E_{\text{Na}}) + \bar{g}_{\text{Nap}} p^3 (V_m^{n,t} - E_{\text{Na}}) + \bar{g}_{\text{K}} s (V_m^{n,t} - E_{\text{K}}) + \bar{g}_{\text{L}} (V_m^{n,t} - E_{\text{L}}) \quad (9)$$

$$V_m^{n,t+dt} = V_m^{n,t} + dt \times \frac{1}{C_n} \times \left(\frac{1}{R_a} \times \text{CNN}([V_m^{n,t}, V_e^{n,t}]) - I_{\text{ion}}^{n,t} + I_{\text{stim}}^{n,t} \right) \quad (10)$$

$$C_n = c_m \times \pi \times d_{\text{node}} \quad (11)$$

$$d_{\text{node}} = \beta_1^n \times D^2 + \beta_2^n \times D + \beta_3^n \quad (12)$$

$$R_a = \frac{\rho_a \partial_x}{\pi \times \left(\frac{d_{\text{axon}}}{2} \right)^2} \quad (13)$$

$$\partial_x = \beta_1^d \times D^2 + \beta_2^d \times D + \beta_3^d \quad (14)$$

$$d_{\text{axon}} = \beta_1^a \times D^2 + \beta_2^a \times D + \beta_3^a \quad (15)$$

where $\beta^n, \beta^d, \beta^a \in \mathbb{R}^3$ are coefficients of interpolants relating fiber diameter (D) to node diameter (d_{node}), internodal distance (∂_x), and axon internodal diameter (d_{axon}), respectively. We used second-order polynomials for these interpolants to remain consistent with previously developed interpolants for the NEURON MRG implementation⁵⁵ (Supplementary Note 3.5). $\text{CNN}([V_m^{n,t}, V_e^{n,t}])$ is a symmetric linear convolution over concatenated V_m and V_e vectors with a three-node-wide kernel centered on node n at time t . $I_{\text{stim}}^{n,t}$ is the intracellular current injected at node n at time t .

Parameters were initialized to the values used in the MRG NEURON model (Supplementary Note 3), with CNN initialized to the second spatial difference. We used backpropagation and gradient descent to optimize 26 parameters of the surrogate model, as detailed in “Surrogate model: training”, to reproduce the responses of the NEURON model.

We implemented S-MF in PyTorch v2.0⁶¹ and deployed it on an NVIDIA RTX A5000 24GB GPU. All arithmetic operations calculating updates to state variables, determining the presence of action potentials, and performing backpropagation were executed on the GPU.

BOX 1

Algorithm 1 - Training Data Generation

```

Solve 1 pig and 1 human vagus nerve model instrumented with the ImThera cuff (FEM)
ns ← number of fascicles in modeled pig nerve
if pig morphology then
    sample field longitudinally at all fascicle centroids (one fiber location per fascicle)
else
    sample field longitudinally with uniform density at ns fiber locations within fascicles
done
for i in 1 to S do
    for j in 1 to C do
         $\mathbf{v} \in \mathbb{R}^N \leftarrow$  randomly selected voltage distribution from pig or human nerve
         $\mathbf{s}_t \in \mathbb{R}^T \leftarrow$  randomly generated monophasic rectangular pulse,  $A \sim U[-0.2, 0.2]$  mA, ▷ see Fig. 6a
         $\delta \sim U[0, 2]$  ms,  $\nabla \sim U[0, 2]$  ms
         $\mathbf{v}_e^j \in \mathbb{R}^{T \times N} \leftarrow \mathbf{s}_t \otimes \mathbf{v}$ 
    end for
     $D \sim U[5.7, 14]$  μm
     $\mathbf{b}_c \in \mathbb{R}^{T \times N} \leftarrow \sum_{j=1}^6 \mathbf{v}_e^j$ 
    simulate MRG fiber (N nodes, fiber diameter D) with extracellular potential boundary condition  $\mathbf{b}_c$ 
    save field-response pair (all N nodes, all T timesteps, all state variables ( $V_m, m, h, p, s$ ))
end for
    
```

Training. All training data were generated from NEURON simulations of the MRG myelinated fiber¹². All simulations were conducted in NEURON v7.8 and solved using the implicit Euler method with a fixed timestep (dt) of 0.005 ms and tstop of 5 ms.

The ultrastructural geometric parameters for the fiber models were interpolated over the originally published fiber geometries (Supplementary Note 3). The training dataset consisted of 65,536 field-response pairs, where the field was a 2D array representing the extracellular potential at every nodal compartment of the model neuron at each timepoint, and the response was a 3D array representing the states of the neuron ($V_m, m, h, p,$ and s) at every nodal compartment at each timepoint. The training data were generated using two vagus nerve morphologies, one pig (Supplementary Note 1 P1) and one human (Supplementary Note 1 H1), instrumented with the ImThera cuff (Fig. 1a). We followed the procedure outlined in Algorithm 1 (Box 1), where $S=65,536$ is the number of sampled field-response pairs, $C=6$ is the number of contacts in the cuff used to generate the data, $N=53$ is the number of nodes of Ranvier per fiber, and $T=1000$ is the number of simulated timesteps.

All modeled fibers had their central node of Ranvier positioned halfway along the modeled nerve with a random longitudinal offset $\sim U[-\frac{INL}{2}, \frac{INL}{2}]$, where INL is the internodal length. NEURON simulations were distributed over 400 CPU cores on the Duke Compute Cluster, using OpenMPI⁶² for load balancing and parallel HDF5 for efficient data handling. The full training set was -100 GB and took ~2 h to generate. We used 80% of the generated data for training, and 20% for validation.

S-MF was trained by gradient descent on minibatches of data that were randomly sampled from the field-response pairs of the training dataset. Training was performed on an NVIDIA RTX A5000 24GB GPU and took ~5 h. We used double-precision floating-point arithmetic for training to mitigate underflow error during the initial training period; subsequent tests using the trained surrogate were performed using single-precision floating-point arithmetic for improved speed.

Twenty-six parameters of S-MF were optimized during training (i.e., θ , “optimizable surrogate model parameters”): maximum ionic conductance values ($\bar{g}_{NaF}, \bar{g}_{Nap}, \bar{g}_K,$ and g_L), axial intracellular resistivity

(ρ_a), membrane capacitance (c_m), coefficients of the d_{node} interpolant ($\beta_1^a, \beta_2^a, \beta_3^a$), coefficients of the d_{axon} interpolant ($\beta_1^a, \beta_2^a, \beta_3^a$), q_{10} parameters ($aq_{10}^m, aq_{10}^h, aq_{10}^p, aq_{10}^s$), parameters governing K+ dynamics (Supplementary Note 3.3, denoted \mathbf{s}^k below), and the values of the symmetric convolutional kernel of CNN (4 values). The objective function for training was to minimize the mean squared error between S-MF and NEURON predictions:

$$\underset{\bar{g}_{NaF}, \bar{g}_{Nap}, \bar{g}_K, g_L, \rho_a, c_m, aq_{10}^m, aq_{10}^h, aq_{10}^p, aq_{10}^s, \beta^a, \mathbf{s}^k, \text{CNN}}{\text{argmin}} \|\mathbf{y} - \hat{\mathbf{y}}\|_{\text{MSE}} \quad (16)$$

where \mathbf{y} is the NEURON model prediction, and $\hat{\mathbf{y}}$ is the S-MF prediction. See Supplementary Note 4 for discussion of the significance of reparametrizing the membrane potential ODE with a trained convolutional network.

For a given training epoch, let $\langle \mathbf{x}_i \in \mathbb{R}^{B \times N \times T \times 1}, \mathbf{y}_i \in \mathbb{R}^{B \times N \times T \times 5} \rangle$ be a randomly sampled (without replacement) minibatch of field-response pairs from the full training dataset, where $i \in \{1, 2, \dots, \lfloor \frac{\text{training set size} = 65,536}{B} \rfloor\}$ is the minibatch index, B is the size of the sampled minibatch (i.e., number of simulated fibers in the sample; $B=64$), N is the number of modeled nodes of Ranvier (53 nodes), and T is the number of simulated time steps (5.0 ms/0.005 ms = 1000 time steps); the fourth dimension is 1 for the extracellular potentials of \mathbf{x}_i and 5 for the number of state variables of \mathbf{y}_i ($V_m, m, h, p,$ and s).

Let $S: \mathbf{x} \in \mathbb{R}^{B \times N \times T \times 1}, \mathbf{s} \in \mathbb{R}^{B \times N \times 5} \rightarrow \hat{\mathbf{y}} \in \mathbb{R}^{B \times N \times T \times 5}$ represent the mapping performed by our surrogate model S from input \mathbf{x} (field) to predicted output $\hat{\mathbf{y}}$ (fiber response) with initial condition (IC) \mathbf{s} (using the numerical methods described in “Surrogate model: Design and implementation”). Let **SSMRG** (“steady-state MRG”) represent the state of MRG nodal compartments at rest ($V_m = -80$ mV and all four gating variables x at $x_\infty(-80$ mV)).

Let $\text{Adam}(\theta, \nabla_{\theta} \mathcal{L}, \lambda)$ represent a single gradient descent update to optimizable surrogate model parameters θ using the Adam update rule with default hyperparameters⁴⁵ and learning rate λ , where

BOX 2

Algorithm 2 - Disjoint Truncated Backpropagation Through Time

In the following, we use NumPy array indexing syntax, where a colon ‘:’ indicates all indices along that dimension and ‘-1’ indicates the final index

```

for epoch in 1 to 5 do
  for  $i$  in 1 to  $\lceil \frac{\text{training set size}}{B} \rceil$  do
    get  $\mathbf{x}_i, \mathbf{y}_i$                                 ▷ sample minibatch  $i$  without replacement from training set
     $\widehat{\mathbf{y}}\mathbf{c} \leftarrow \text{None}$ 
    for  $c$  in 0 to  $\lfloor T/\Delta \rfloor$  do                ▷ chunk index
       $\mathbf{xc} \leftarrow \mathbf{x}_i[:, :, c \times \Delta : c \times \Delta + \Delta, :]$     ▷ get disjoint chunk of  $\mathbf{x}_i$ , length  $\Delta$ 
       $\mathbf{yc} \leftarrow \mathbf{y}_i[:, :, c \times \Delta : c \times \Delta + \Delta, :]$ 
      if  $c = 0$  then                               ▷ get disjoint chunk of  $\mathbf{y}_i$ , length  $\Delta$ 
         $\mathbf{s} \leftarrow \text{SSMRG}$                              ▷ if first chunk, set  $V_m = -80\text{mV}$  and gating variables at  $\infty$ 
      else
         $\mathbf{s} \leftarrow \widehat{\mathbf{y}}\mathbf{c}[:, :, -1, :]$ 
      done                                           ▷ otherwise, set IC as previous last state predicted by surrogate
       $\widehat{\mathbf{y}}\mathbf{c} \leftarrow S(\mathbf{xc}, \mathbf{s})$                        ▷ calculate surrogate model prediction for input  $\mathbf{xc}$  and IC  $\mathbf{s}$ 
       $L \leftarrow \|\mathbf{yc} - \widehat{\mathbf{y}}\mathbf{c}\|_{\text{MSE}}$                  ▷ calculate MSE loss between surrogate prediction and NEURON
       $\text{Adam}(\theta, \nabla_{\theta} L, \lambda)$                    ▷ perform gradient descent step
    end for
  end for
end for

```

gradients of a scalar valued loss L (mean squared error) with respect to the surrogate model parameters θ are represented by $\nabla_{\theta} L$.

We trained S-MF over 5 epochs using the procedure outlined in Algorithm 2 (Box 2), a variation on truncated backpropagation through time⁶³ with the Adam update rule. Training was performed on temporally sequential disjoint chunks of minibatch data, with chunk length (in terms of number of timesteps) denoted Δ . We used minibatch size $B = 64$ fibers, $\Delta = 50$, and $\lambda = 1e - 5$.

Testing—activation thresholds. We simulated activation thresholds using both S-MF and NEURON models to quantify the accuracy of S-MF across fiber diameters, stimulus waveforms, nerve morphologies, and cuff geometries. We modeled 6 human and 6 pig vagus nerve morphologies (Supplementary Note 1 P2-7, H2-7) each instrumented with the six-contact ImThera cuff and the helical cuff for a total of 24 finite element models for testing; these nerves were constructed from different histological cross sections than those used to train S-MF.

For the ImThera cuff, for each nerve, for each of the six electrode contacts, we randomly (with replacement) selected an associated fascicle. We modeled 8 fiber diameters (5.7 to 14 μm) at the centroid of each selected fascicle and delivered monopolar stimulation from the associated electrode contact. For the helical cuff, for each nerve, we randomly (without replacement) selected up to 10 fascicles (or the number of fascicles in the nerve if the nerve had fewer than 10 fascicles). We modeled 8 fiber diameters (5.7 to 14 μm) at the centroid of each selected fascicle and delivered bipolar stimulation from the cuff. Every modeled fiber consisted of 101 nodes of Ranvier (increased from 53 to reduce the likelihood of end-excitation for optimization and threshold tests), and all modeled fibers had their central node of Ranvier positioned halfway along the modeled nerve.

We simulated activation thresholds for S-MF and NEURON models using a bisection search algorithm to determine a threshold within 1% tolerance, where activation was defined as V_m crossing -20 mV with a rising edge at a node of Ranvier 5 nodes from either end of the fiber. We used a time step of 0.005 ms and simulated a total time (tstop) of 5 ms.

We evaluated six waveforms: monophasic rectangular, symmetric biphasic rectangular, sawtooth, exponential, sinusoid, and Gaussian (Fig. 6). For each waveform, we tested five pulse widths (δ): 0.1, 0.2, 0.5, 0.75, and 1 ms. We compared the activation thresholds predicted by S-MF (T_{SURR}) and by NEURON (T_{NRN}) using the relative percent threshold error ($100\% \times \frac{T_{\text{SURR}} - T_{\text{NRN}}}{T_{\text{NRN}}}$). In total, for nerves instrumented with the ImThera cuff, we tested 17,280 thresholds (12 nerves \times 6 contact-fascicle pairs \times 6 waveforms \times 5 pulse widths \times 8 fiber diameters), and for nerves instrumented with the helical cuff, we tested 25,200 thresholds (12 nerves \times up to 10 fascicles \times 6 waveforms \times 5 pulse widths \times 8 fiber diameters).

We compared the performance of S-MF with the published Peterson threshold estimator¹⁴ designed to predict activation thresholds for monophasic rectangular waveforms. Their approach involves calculating a driving force estimator termed MDF_2 that uses a weighted sum of second spatial differences over 19 adjacent nodes of Ranvier:

$$\text{MDF}_2 = \sum_j W_{|n-j|}(\text{PW}, d) \cdot \Delta^2 V_e^j \quad (17)$$

n is the node index for which MDF_2 is being calculated. Weights $W_{|n-j|}$ are the ratio of depolarization caused by the current injected at node j to the depolarization caused by current injected when $j = n$ for a given pulse width (PW) and fiber diameter d in a linearized MRG model (generated by replacing all nonlinear ion channels with a fixed specific conductance = 0.007 mS cm^{-2}). Rather than using their published MDF_2 values, we reimplemented their procedures to calculate threshold MDF_2 for the specific fiber diameters and pulse widths that we tested for our MRG implementation, to avoid any interpolations of fiber diameter or pulse width. We validated our implementation of the Peterson surrogate model by reproducing the published threshold errors across different electrode placements (Supplementary Note 5).

Testing—other non-linear responses to stimulation. We evaluated S-MF’s accuracy in reproducing non-linear responses to stimulation other than activation thresholds: unidirectional action potential (AP)

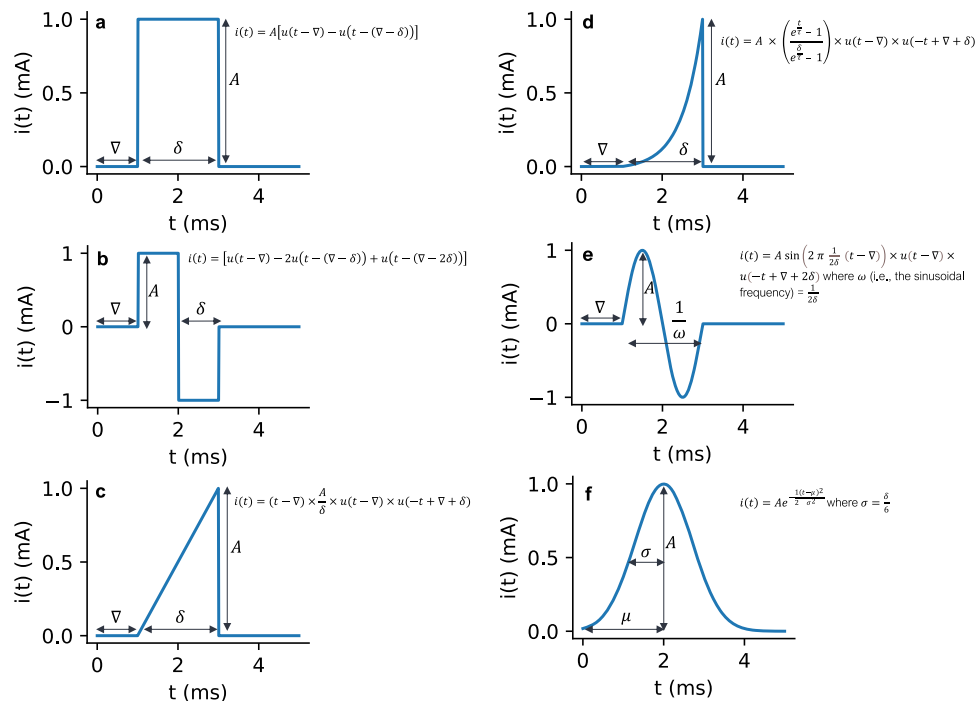


Fig. 6 | Waveform definitions. a Monophasic rectangular. **b** Symmetric biphasic rectangular. **c** Sawtooth. **d** Exponential. **e** Sinusoid. **f** Gaussian.

propagation, bidirectional direct current (DC) block of AP propagation, AP annihilation by collision, excitation and block responses to kilohertz frequency signals, and state-dependent modulation of activity at conventional stimulation frequencies.

For unidirectional propagation and bidirectional DC block, we randomly selected a fascicle and an electrode contact in a pig vagus nerve model instrumented with the ImThera cuff (Supplementary Note 1 P3). We placed a 12 μm diameter fiber in the centroid of the selected fascicle and delivered stimulation using a 0.75 ms cathodic monophasic rectangular pulse from the selected electrode contact. We delivered stimulation at progressively increasing amplitudes to observe transitions in the neuronal response from activation to unidirectional propagation to block to re-excitation (Fig. 3a) in both the NEURON and S-MF models. We used a timestep (dt) of 0.005 ms and a tstop of 7.5 ms.

For kilohertz frequency stimulation, we randomly selected 7 fascicles and 1 electrode contact in a pig vagus nerve model (Supplementary Note 1 P2) instrumented with the ImThera cuff. We placed 5.7, 8.7, and 14 μm diameter fibers at the centroids of the selected fascicles and delivered stimulation from the selected electrode contact. We delivered extracellular sinusoidal stimulation at 1, 2, 5, and 10 kHz starting at $t = 0.5$ ms across a range of amplitudes (0 to 10 mA for the 5.7 μm fibers, 0 to 5 mA for the 8.7 μm fibers, and 0 to 2 mA for the 14 μm fibers). We initiated action potentials at a rate of 100 Hz at one end of each fiber using a 0.1 ms intracellular anodic current pulse (2 nA) starting at $t = 50$ ms and recorded the number of action potentials (defined as V_m crossing -20 mV with a rising edge) arriving at the other end of each fiber after $t = 50$ ms. In all cases, we simulated a total of 100 ms with a dt of 0.001 ms. We compared the population response (mean number of action potentials and 95% confidence interval across the 7 sampled fiber locations) for each stimulation frequency and amplitude for each fiber diameter between NEURON and S-MF (Fig. 3b).

For AP collision, we initiated action potentials at both ends of 5.7, 8.7, 10, and 14 μm fibers using a 2 nA, 0.1 ms intracellular rectangular current pulse. We recorded snapshots of V_m along the length of the fibers at various timepoints to observe how the propagated APs interacted with each other (Fig. 3c).

For state-dependent modulation of activity at conventional frequencies, we modeled the effects of extracellular stimulation across a range of suprathreshold and subthreshold amplitudes on fibers firing intrinsically at different mean frequencies due to random intracellular current pulses. We selected a human nerve model (Supplementary Note 1 H2) instrumented with the helical cuff and randomly selected 7 of its fascicles (without replacement). We modeled 5.7, 8.7, and 14 μm diameter fibers (101 nodes each) at the centroid of each fascicle. We first calculated activation thresholds in NEURON for every fiber for a single 0.1 ms symmetric biphasic pulse using a bisection search algorithm, without any intrinsic activity. For both the NEURON and S-MF models, we then stimulated extracellularly with trains of symmetric biphasic pulses at {30, 50, and 100 Hz}, pulse widths of 0.1 ms, and amplitudes from -50% to $+50\%$ of the NEURON-calculated mean threshold (across the seven sampled locations for the given fiber diameter). We simultaneously delivered intracellular current pulses to one end of the fiber following five Poisson random generated sequences at each of 3 mean rates $\mu_{\text{IFR}} \in \{10, 50, 100 \text{ Hz}\}$, totaling 15 random intrinsic firing patterns for each fiber location and diameter for each extracellular stimulus parameter combination (frequency \times amplitude). We simulated for $t = 100 \times (100/\mu_{\text{IFR}})$ ms using a dt of 0.005 ms with and without extracellular stimulation. We recorded V_m at the fifth node of Ranvier from the end of the fiber distal to the location of intracellular stimulation. We calculated the spike times for APs arriving at the V_m recording node (the timepoints at which V_m crossed -20 mV in the positive direction). We compared spike times with and without extracellular stimulation using the SPIKE-synchronization metric⁶⁴, S_C , which is a mean over aggregated coincidence indicators $C(t_k)$ between two spike trains:

$$S_C = \frac{1}{M} \sum_{k=1}^M C(t_k) \quad (18)$$

where M is the total number of spikes (across both spike trains) and t_k is the k^{th} spike time in the pooled spike train (in the case of an exact match, k counts both spikes). $C(t)$ is calculated per spike in the original

pair of spike trains and subsequently pooled:

$$C(t_i^{(1)}) = \begin{cases} 1 & \text{if } \min_j (|t_i^{(1)} - t_j^{(2)}|) < \tau_{ij}^{(1,2)} \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

$$\tau_{ij}^{(1,2)} = \min \left\{ t_{i+1}^{(1)} - t_i^{(1)}, t_i^{(1)} - t_{i-1}^{(1)}, t_{j+1}^{(2)} - t_j^{(2)}, t_j^{(2)} - t_{j-1}^{(2)} \right\} / 2 \quad (20)$$

where $t_i^{(1)}$ is the i^{th} spike time in train 1 (and likewise for $t_j^{(2)}$). S_C is 0 if and only if the two spike trains do not contain any coincidences, and S_C is one if and only if each spike in each spike train has exactly one matching spike in the other spike train. We used this metric to measure the extent to which the extracellular signal transformed the intrinsic firing signal (Fig. 3d).

Optimization for spatial selectivity

We designed and implemented two optimization methods (gradient-free and gradient-based). We applied both optimization methods to define rectangular pulses and arbitrary waveforms to achieve spatially selective fiber activation in 12 nerve models (Supplementary Note 1 P2-7 and H2-7) instrumented with the six-contact ImThera cuff. We used NEURON to validate all parameters returned by optimizations that used S-MF.

The resting inner diameter of the original ImThera cuff model (which was used to generate training data and activation threshold data for surrogate model validation) was up to 1.9× larger than the diameter of our human nerve specimens (without shrinkage correction). Therefore, for human nerve optimization tasks, we inflated all nerve morphologies by 25% to correct for 20% shrinkage in the tissue preparation process²⁶ and instrumented the resulting nerves with a shrunk version of the ImThera cuff. We reduced the cuff's resting inner diameter from 3 mm to 2.5 mm, but retained the cuff length, insulation thickness, number of contacts, contact size, and lengthwise contact spacing. Edge-to-edge contact spacing of the shrunk ImThera cuff was reduced from 1.05 mm to 0.96 mm. Circumferential inter-contact spacing was reduced from 1.34 mm to 1.11 mm; this spacing was preserved as the cuff expanded to fit around nerves with an inner diameter larger than 2.5 mm.

Optimization for spatial selectivity: tasks

The overall task in all cases was to maximize activation of target fibers and avoid activation of off-target fibers, with different objective functions implemented based on the optimization method (see “Optimization for spatial selectivity: algorithms”). We defined the functional organization of the pig vagus nerves using immunohistochemistry which showed that afferents and efferents are segregated on separate halves of the nerve⁴¹, as confirmed with in vivo imaging⁶⁵ and physiological recordings³⁵. Target and off-target fiber populations were assigned based on this histological designation of sensory/motor (afferent/efferent) fascicles, with sensory fibers targeted for selective activation (Supplementary Note 1 P2-7). Analogous immunohistochemical data are not yet available for human vagus nerves, but this functional organization reflects the gross anatomical structure of the nodose ganglion—where the nodose contains the cell bodies of visceral vagal afferents, next to which efferent fibers pass—and therefore we assumed a similar vagotomy between the pig and human nerves. In human nerves (Supplementary Note 1 H2-7), we randomly partitioned the nerve approximately into two semicircles and randomly assigned all fascicles in each semicircle as either target or off-target. We placed a single 5.7 μm diameter fiber at the centroid of each fascicle, each with 101 nodes of Ranvier and a node placed at the longitudinal midpoint of the nerve.

Stimulation with rectangular pulses. We optimized the stimulation amplitudes delivered through all six contacts of the ImThera cuff for symmetric biphasic rectangular waveforms (PW = 0.4 ms, delay = 0.4 ms). For each nerve, this yielded an optimization problem with 6 degrees of freedom. We compared optimized performance using the ImThera cuff with the best selectivity predicted for bipolar stimulation with the helical cuff. Specifically, we determined the optimum amplitude for stimulation with the helical cuff (using the same biphasic rectangular waveform) by calculating activation thresholds (in NEURON) for every fiber in the nerve and subsequently determining the threshold amplitude that minimized the weighted binary cross-entropy loss (Eq. (21)).

Stimulation with arbitrary waveforms. We optimized the current delivered at each timestep (dt = 0.005 ms) between 0.2 and 1.2 ms (the ‘nonzero period’) from all six contacts of the ImThera cuff. For each nerve, this yielded an optimization problem with 1200 degrees of freedom (200 current amplitudes corresponding to 200 timesteps for each of the six contact). We enforced charge balance for each contact by subtracting the mean of the nonzero period from the nonzero period for each contact before evaluating the neural response.

Optimization for spatial selectivity: algorithms

We developed two optimization methods: gradient-free (using an adapted differential evolution algorithm) and gradient-based (using gradient descent). For the gradient-free method, we compared the performance between the NEURON and S-MF models. Conversely, the gradient-based method leveraged the differentiability of the surrogate; therefore, we only used the surrogate model.

Gradient-free optimization (differential evolution). We implemented an adapted differential evolution (DE) algorithm⁶⁶ to optimize stimulation parameters (Fig. 7). We defined an initial population of candidate solutions (i.e., sets of stimulation parameters) and iteratively updated the candidate solutions based on their performance on the target optimization task (i.e., selective activation of target neurons, as defined in “Optimization for spatial selectivity: tasks”). Thus, the DE approach is stochastic and does not rely on the differentiability of the loss function with respect to the parameters being optimized, i.e., it is gradient-free. We selected DE for its relative ease of implementation, applicability to real-valued loss functions, widespread use, and continued development⁶⁷.

We initialized the population of candidate solutions using random Latin hypercube sampling. For all DE optimizations, we used bounds of (−0.3, 0.3) mA. For a population of size n , a random Latin hypercube was constructed by splitting each of the d dimensions to be sampled into n equivalently sized intervals over the respective bounds. Each dimension corresponded to a parameter to be optimized in our model of extracellular peripheral nerve stimulation: for the rectangular pulses, this corresponded to one amplitude for each contact, yielding a 6-dimensional optimization problem, and for arbitrary waveforms, this corresponded to 200 current amplitudes for each contact, yielding a 1200-dimensional optimization problem. We then randomly selected a value from a uniform distribution in each of the n intervals, and we generated the initial population of candidate solutions by randomly shuffling these values to produce n sets of parameters. We used populations of 100 and 300 candidate solutions for the NEURON and S-MF optimizations, respectively.

For all optimizations using the DE algorithm, we calculated the loss function as the weighted binary cross-entropy (WBCE):

$$f(\mathbf{X}) = WBCE = \sum_{n=1}^N \frac{f_n}{F} \left((1 - A_n) \times \ln(1 - \hat{A}_n) + A_n \times \ln \hat{A}_n \right) \quad (21)$$

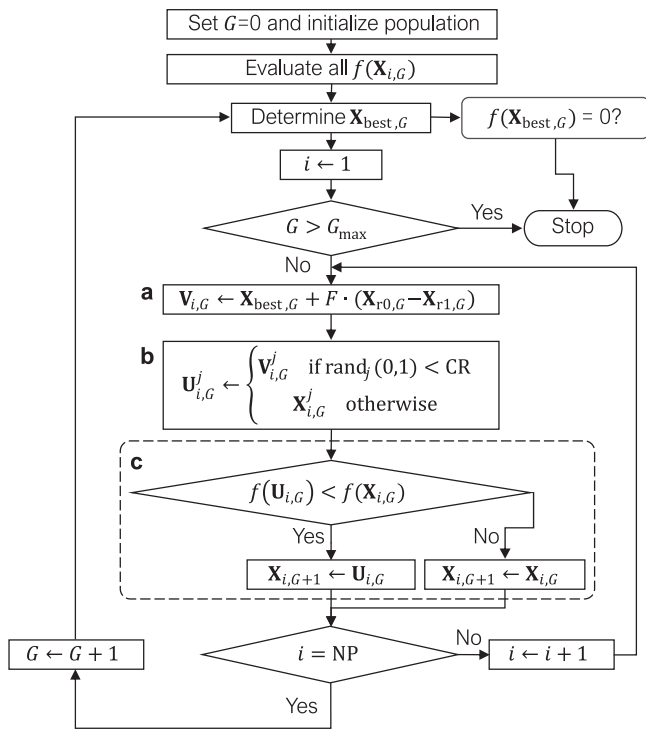


Fig. 7 | Algorithmic flowchart for DE/best/1/bin. G = generation number; G_{\max} = maximum number of generations permitted; $\mathbf{X}_{i,G}$ = population member/candidate solution i in generation G ; $\mathbf{X}_{r0,G}$, $\mathbf{X}_{r1,G}$ = mutually exclusive random members selected from population in generation G ; $\mathbf{X}_{\text{best},G}$ = best population member in generation G ; $\text{rand}(0,1)$ = random number from uniform distribution in interval $[0, 1]$; j = index for each element in $\mathbf{X}_{i,G}/\mathbf{U}_{i,G}/\mathbf{V}_{i,G}$; CR = crossover rate; F = mutation rate; f = loss criterion; NP = population size. **a** Mutation operation. **b** Crossover operation. **c** Selection operation.

where N is the total number of fascicles in the nerve, f_n is the cross-sectional area of fascicle n , F is the summed cross-sectional area of all fascicles in the nerve, $A_n \in \{1,0\}$ is the target activation (activated or not) of the fiber in fascicle n , and $\hat{A}_n \in \{1,0\}$ is the predicted activation (using NEURON or S-MF) of the fiber in fascicle n . Fiber activation was determined by the presence of at least one AP, defined as a rising V_m crossing 0 mV, in either node of Ranvier 5 nodes from each end of the fiber.

Each generation, the population of candidate solutions was updated through a process of mutation (Fig. 7a), crossover (Fig. 7b), and selection (Fig. 7c). A “trial vector” $\mathbf{U}_{i,G}$ was constructed for each i^{th} candidate solution $\mathbf{X}_{i,G}$ that was evaluated in the current generation G using the “best/1/bin”⁶⁸ mutation rule and crossover operation:

$$\mathbf{V}_{i,G} = \mathbf{X}_{\text{best},G} + F \cdot (\mathbf{X}_{r0,G} - \mathbf{X}_{r1,G}) \quad (22)$$

$$\mathbf{U}_{i,G}^j = \begin{cases} \mathbf{V}_{i,G}^j & \text{if } \text{rand}_j(0,1) < \text{CR} \\ \mathbf{X}_{i,G}^j & \text{otherwise} \end{cases} \quad \forall j \in 0, 1 \dots d \quad (23)$$

where $\mathbf{X}_{\text{best},G}$ is the best candidate solution in the current generation, F is the mutation rate, $\mathbf{X}_{r0,G}$ and $\mathbf{X}_{r1,G}$ are mutually exclusive randomly selected candidate solutions from the current generation, $\text{rand}(0, 1)$ is a random sample from a uniform distribution in the interval $(0, 1)$ (resampled for every j), CR is the crossover rate, and d is the dimensionality of the optimization problem. If the resulting trial vector $\mathbf{U}_{i,G}$ exceeded the problem bounds in any dimension, the mutation and crossover operation was repeated up to 10 times, after which those dimensions which remained outside the bounds were randomly

reinitialized to values within the bounds. This is a feasibility-preserving strategy for boundary constraint handling that has shown advantages over Darwinian (i.e., where boundary violations are discouraged by adding a penalty term to the loss function) and repair (i.e., where dimensions in which the boundary has been violated are placed back within the bounds following some rule, such as reflection against the boundary) methods in recent empirical studies using DE⁶⁹. $\mathbf{U}_{i,G}$ then replaced $\mathbf{X}_{i,G}$ in generation $G + 1$ if $f(\mathbf{U}_{i,G}) < f(\mathbf{X}_{i,G})$, i.e., if the loss of $\mathbf{U}_{i,G}$ was less than the loss of $\mathbf{X}_{i,G}$, or equivalently, $\mathbf{U}_{i,G}$ had a better performance.

We used parameter adaptation to improve convergence⁷⁰. At each generation, a crossover rate CR_i was generated independently for each member of the population, sampling from a normal distribution with mean CR and standard deviation of 0.1, and subsequently clipped to $[0, 1]$. Let \mathbf{S}_{CR} be the set of all such crossover rates in the current generation for which the resulting trial vector $\mathbf{U}_{i,G}$ successfully replaced the corresponding target vector $\mathbf{X}_{i,G}$ (i.e., had a lower loss). In the next generation, CR was updated such that $\text{CR} = (1 - c) \times \text{CR} + c \times \text{mean}_A(\mathbf{S}_{\text{CR}})$, where mean_A is the arithmetic mean and c is a positive number between 0 and 1. Likewise, a mutation rate F_i was generated independently for every member of the population by random sampling from a Cauchy distribution with location parameter F and scale parameter 0.1, truncated to 1 if $F_i > 1$ and regenerated if $F_i \leq 0$. Using the same terminology as before, in the subsequent generation, F was updated such that $F = (1 - c) \times F + c \times \text{mean}_L(\mathbf{S}_F)$, where $\text{mean}_L(\mathbf{S}_F)$ is the Lehmer mean = $\frac{\sum_{F \in \mathbf{S}_F} F^2}{\sum_{F \in \mathbf{S}_F} F}$. We initialized F to 0.8, CR to 1.0, and we used a value of $c = 0.1$ as in the original paper.

We terminated the DE operation after 200 or 500 generations (G_{\max}) for the NEURON and S-MF optimizations, respectively, or if a solution with a theoretic minimum WBCE loss (0) was found.

Gradient-based optimization (gradient descent). Gradient descent operates on scalar-valued differentiable loss functions. Surrogate models implemented using AxonML are differentiable and implemented in PyTorch; therefore, neuromodulation parameters can be designed using gradient descent for certain optimization tasks.

We optimized stimulation parameters using the procedure outlined in Algorithm 3 (Box 3). Let $\mathbf{S} : \mathbf{x} \in \mathbb{R}^{B \times N \times T \times 1}, \mathbf{s} \in \mathbb{R}^{B \times N \times 5} \rightarrow \hat{\mathbf{y}} \in \mathbb{R}^{B \times N \times T \times 5}$ represent the mapping performed by S-MF \mathbf{S} from input \mathbf{x} (field), to predicted output $\hat{\mathbf{y}}$ (the fiber response, i.e., membrane potential and gating variables for all nodes at all simulated timepoints) with initial condition \mathbf{s} , where B is the number of fibers in the nerve model, N is the number of modeled nodes of Ranvier per fiber, and T is the number of simulated timepoints. Let \mathbf{SSMRG} (steady-state MRG) represent the state of MRG nodal compartments at rest ($V_m = -80$ mV and all four gating variables x at $x_{\infty}(-80$ mV)).

Let $\text{MF} : \mathbf{p} \in \mathbb{R}^d \rightarrow \mathbf{x} \in \mathbb{R}^{B \times N \times T \times 1}$ represent the mapping from the set of optimizable stimulus parameters \mathbf{p} (where d is the degrees of freedom; $d = 6$ in the waveform constrained case and $d = 1200$ in the arbitrary waveform case) to the S-MF input \mathbf{x} (tensor representing the extracellular field at every node of every fiber at every timestep). Let $\text{Ranger}(\mathbf{p}, \nabla_{\mathbf{p}} L, \lambda)$ represent a single gradient descent update to optimizable stimulus parameters \mathbf{p} using the Ranger update rule⁷¹ (Rectified Adam⁷² with LookAhead⁷³ and Gradient Centralization⁷⁴) and learning rate λ , where gradients with respect to the stimulus parameters \mathbf{p} of a scalar valued loss L (weighted quotient (WQ) loss, Eq. (26)) are represented by $\nabla_{\mathbf{p}} L$. Let $\text{WBCE} : \mathbf{A} \in \{0,1\}^N, \hat{\mathbf{A}} \in \{0,1\}^N \rightarrow \mathbb{R}$ represent the weighted binary cross-entropy (Eq. (21)) between target activations \mathbf{A} and predicted activations $\hat{\mathbf{A}}$ (where activation, in contrast to $\hat{\mathbf{y}}$, is a derived Boolean quantity indicating whether the fiber propagated an action potential, defined, as before, by V_m crossing -20 mV in the positive direction at 5 nodes from either end of

BOX 3

Algorithm 3 - Gradient Descent for Stimulus Parameter Optimization

We used $wd = 0.01 \cdot N$ (where N is the number of fascicles in the nerve), $\gamma_{\max} = 200/N$, $\eta = 0.6$, and $ns = 200$

```

Requires  $wd, \eta, ns, \gamma_{\max}, \mathbf{A}$     ▷ weight decay, learning rate decay, # gradient descent steps, maximum gradient norm, target fiber
                                       activations
 $\mathbf{p} \leftarrow \{0\}^d$                 ▷ initialize stimulus parameters to 0
 $\lambda \leftarrow 2$                   ▷ initialize learning rate to 2
 $\mathbf{p}_{\text{best}}, \omega_{\text{best}}, \mathcal{L}_{\text{best}} \leftarrow \text{None}$ 
for  $i$  in 1 to  $ns$  do              ▷ optimization loop
   $\mathbf{x} \leftarrow \text{MF}(\mathbf{p})$           ▷ construct field from stimulus parameters
   $\hat{\mathbf{y}} \leftarrow \text{S}(\mathbf{x}, \text{SSMRG})$     ▷ use surrogate to predict neural response
   $\omega \leftarrow \text{WBCE}(\mathbf{A}, \hat{\mathbf{A}})$     ▷ calculate weighted binary cross entropy loss
   $\mathcal{L} \leftarrow \text{WQ}(\hat{\mathbf{y}})$         ▷ calculate weighted quotient
  if  $i = 1$  then
     $\mathbf{p}_{\text{best}}, \omega_{\text{best}}, \mathcal{L}_{\text{best}} \leftarrow \mathbf{p}, \omega, \mathcal{L}$   ▷ initialize best params and losses to initial values
  done
  if  $\omega < \omega_{\text{best}}$  then          ▷ if improvement in WBCE
     $\mathbf{p}_{\text{best}}, \omega_{\text{best}} \leftarrow \mathbf{p}, \omega$     ▷ update best params and WBCE
  done
  if  $\omega = \omega_{\text{best}}$  and  $\mathcal{L} < \mathcal{L}_{\text{best}}$  then  ▷ if improvement in WQ and WBCE still best
     $\mathbf{p}_{\text{best}}, \omega_{\text{best}}, \mathcal{L}_{\text{best}} \leftarrow \mathbf{p}, \omega, \mathcal{L}$   ▷ update best params, WBCE, and WQ
    if  $\omega_{\text{best}} < 1$  then
       $\lambda \leftarrow \lambda \cdot \eta$     ▷ reduce learning rate if near minimum
    done
  done
   $\boldsymbol{\gamma}_{\mathbf{p}} \leftarrow \nabla_{\mathbf{p}} \mathcal{L}$     ▷ calculate gradients
   $\boldsymbol{\gamma}_{\mathbf{p}} \leftarrow \min\left(1, \frac{\gamma_{\max}}{\|\boldsymbol{\gamma}_{\mathbf{p}}\|}\right) \boldsymbol{\gamma}_{\mathbf{p}}$   ▷ clip gradient norm
   $\mathbf{p} \leftarrow \mathbf{p} \odot (1 - wd)$     ▷ apply weight decay on stimulus parameters
   $\text{Ranger}(\mathbf{p}, \boldsymbol{\gamma}_{\mathbf{p}}, \lambda)$     ▷ perform gradient descent step on weighted quotient loss
end for
return  $\mathbf{p}_{\text{best}}$                     ▷ return best params (corresponds to min. WQ for min. WBCE)

```

the fiber). Finally, let $\text{WQ} : \hat{\mathbf{y}} \in \mathbb{R}^{B \times N \times T \times 5} \rightarrow \mathbb{R}$ be the weighted quotient loss on S-MF prediction $\hat{\mathbf{y}}$:

$$m_{\text{on}}(\hat{\mathbf{y}}) = \sum_{n \in \mathbf{n}_{\text{on}}} f_n \times m_n(\hat{\mathbf{y}}) \quad (24)$$

$$m_{\text{off}}(\hat{\mathbf{y}}) = \sum_{n \in \mathbf{n}_{\text{off}}} f_n \times m_n(\hat{\mathbf{y}}) \quad (25)$$

$$\text{WQ}(\hat{\mathbf{y}}) = \sqrt{\frac{d}{n_c} \frac{m_{\text{off}}(\hat{\mathbf{y}})}{m_{\text{on}}(\hat{\mathbf{y}})}} \quad (26)$$

where \mathbf{n}_{on} are the indices of all target fascicles, \mathbf{n}_{off} are the indices of all off-target fascicles, f_n is the cross-sectional area of fascicle n , $m_n(\hat{\mathbf{y}})$ is the sum over the m gating parameter in $\hat{\mathbf{y}}$ corresponding to the 20 end nodes of Ranvier (10 at each end, including the terminal nodes) of the fiber in fascicle n , and n_c ($n_c = 6$) is the number of electrode contacts in the model for which there were optimizable parameters.

We leveraged the concurrency of our S-MF GPU implementation to evaluate the neural response and calculate gradients across all 12 nerve models in parallel. This allowed us to solve all 12 optimization problems, with rectangular pulses or arbitrary waveforms, simultaneously with modest overhead.

Statistical analysis

We compared performances of the various optimization methods using a non-parametric paired test (two-sided Wilcoxon signed-rank test) between the WBCE values of generated solutions. We used a non-parametric test because most distributions of WBCE values were non-Gaussian (as verified by the Shapiro–Wilk test). We adjusted significance values (with an uncorrected critical p value = 0.05) using the Holm-Bonferroni method⁷⁵ to account for multiple comparisons; we reported both p and corrected significance values. All Boxplots are displayed using the Tukey method (center line, median; box limits, upper and lower quartiles; whiskers, last point within 1.5× interquartile range of nearest hinge).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Inputs to ASCENT⁵⁵ to build and solve finite element models of vagus nerve stimulation, data required to run example applications in the code repositories, and source data and scripts to generate figures are publicly available in the Duke Research Data Repository (<https://doi.org/10.7924/r48g8tf2476>). Histological data used to construct the nerve models are publicly available in the SPARC Consortium database under accession codes 64 (pig; <https://doi.org/10.26275/maq2-eii457>) and 65 (human; <https://doi.org/10.26275/ofja-ghoz58>). Source data are provided with this paper.

Code availability

Core packages for executing NEURON models and performing gradient-free selective stimulation optimization are available at <https://github.com/minhajh/cajal>⁷⁷. Surrogate model implementation, training, dataset generation, and gradient-based selective stimulation optimization code is available at <https://github.com/minhajh/axonml>²¹.

References

- Raspopovic, S., Petrini, F. M., Zelechowski, M. & Valle, G. Framework for the development of neuroprostheses: from basic understanding by sciatic and median nerves models to bionic legs and hands. *Proc. IEEE* **105**, 34–49 (2017).
- Butson, C. R., Miller, I. O., Normann, R. A. & Clark, G. A. Selective neural activation in a histologically derived model of peripheral nerve. *J. Neural Eng.* **8**, 036009 (2011).
- Kent, A. R. & Grill, W. M. Model-based analysis and design of nerve cuff electrodes for restoring bladder function by selective stimulation of the pudendal nerve. *J. Neural Eng.* **10**, 036010 (2013).
- Schiefer, M. A., Tyler, D. J. & Triolo, R. J. Probabilistic modeling of selective stimulation of the human sciatic nerve with a flat interface nerve electrode. *J. Comput. Neurosci.* **33**, 179–190 (2012).
- Schiefer, M. A., Triolo, R. J. & Tyler, D. J. A model of selective activation of the femoral nerve with a flat interface nerve electrode for a lower extremity neuroprosthesis. *IEEE Trans. Neural Syst. Rehabil. Eng. Publ. IEEE Eng. Med. Biol. Soc.* **16**, 195–204 (2008).
- Couto, J. & Grill, W. M. Kilohertz frequency deep brain stimulation is ineffective at regularizing the firing of model thalamic neurons. *Front. Comput. Neurosci.* **10**, 22 (2016).
- Kilgore, K. L. & Bhadra, N. Nerve conduction block utilizing high-frequency alternating current. *Med. Biol. Eng. Comput.* **42**, 394–406 (2004).
- Pelot, N. A., Behrend, C. E. & Grill, W. M. Modeling the response of small myelinated axons in a compound nerve to kilohertz frequency signals. *J. Neural Eng.* **14**, 046022 (2017).
- Kumbhar, P. et al. CoreNEURON: an optimized compute engine for the NEURON simulator. *Front. Neuroinformatics* **13**, 63 (2019).
- Ben-Shalom, R. et al. NeuroGPU: accelerating multi-compartment, biophysically detailed neuron simulations on GPUs. *J. Neurosci. Methods* **366**, 109400 (2022).
- Akar, N. A. et al. Arbor -- a morphologically-detailed neural network simulation library for contemporary high-performance computing architectures. *2019 27th Euromicro Int. Conf. Parallel Distrib. Netw.-Based Process. PDP 274–282* <https://doi.org/10.1109/EMPDP.2019.8671560> (2019).
- McIntyre, C. C., Richardson, A. G. & Grill, W. M. Modeling the excitability of mammalian nerve fibers: influence of afterpotentials on the recovery cycle. *J. Neurophysiol.* **87**, 995–1006 (2002).
- Warman, E. N., Grill, W. M. & Durand, D. Modeling the effects of electric fields on nerve fibers: determination of excitation thresholds. *IEEE Trans. Biomed. Eng.* **39**, 1244–1254 (1992).
- Peterson, E., Izad, O. & Tyler, D. Predicting myelinated axon activation using spatial characteristics of the extracellular field. *J. Neural Eng.* **8**, 046030 (2011).
- Howell, B., Gunalan, K. & McIntyre, C. C. A driving-force predictor for estimating pathway activation in patient-specific models of deep brain stimulation. *Neuromodulation Technol. Neural Interface* **22**, 403–415 (2019).
- Lagaris, I. E., Likas, A. & Fotiadis, D. I. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans. Neural Netw.* **9**, 987–1000 (1998).
- Baymani, M., Kerayechian, A. & Effati, S. Artificial neural networks approach for solving stokes problem. *Appl. Math.* **1**, 288–292 (2010).
- Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
- Wang, S. et al. Massive computational acceleration by using neural networks to emulate mechanism-based biological models. *Nat. Commun.* **10**, 4354 (2019).
- Egberts, G., Schaaphok, M., Vermolen, F. & Zuijlen, P. van. A Bayesian finite-element trained machine learning approach for predicting post-burn contraction. *Neural Comput. Appl.* **34**, 8635–8642 (2022).
- Hussain, M. A., Grill, W. M. & Pelot, N. A. minhajh/axonml: v1.0.0. *Zenodo* <https://doi.org/10.5281/zenodo.12752387> (2024).
- Romeni, S., Valle, G., Mazzoni, A. & Micera, S. Tutorial: a computational framework for the design and optimization of peripheral neural interfaces. *Nat. Protoc.* **15**, 3129–3153 (2020).
- Wongsarnpigoon, A. & Grill, W. M. Energy-efficient waveform shapes for neural stimulation revealed with genetic algorithm. *J. Neural Eng.* **7**, 046009 (2010).
- Foutz, T. J., Ackermann, D. M. Jr., Kilgore, K. L. & McIntyre, C. C. Energy efficient neural stimulation: coupling circuit design and membrane biophysics. *PLoS ONE* **7**, e51901 (2012).
- Kilgore, K. L. & Bhadra, N. Reversible nerve conduction block using kilohertz frequency alternating current. *Neuromodulation J. Int. Neuromodulation Soc.* **17**, 242–254 (2014).
- Musselman, E. D., Pelot, N. A. & Grill, W. M. Validated computational models predict vagus nerve stimulation thresholds in preclinical animals and humans. *J. Neural Eng.* **20**, 036032 (2023).
- Medina, L. E. & Grill, W. M. Nerve excitation using an amplitude-modulated signal with kilohertz-frequency carrier and non-zero offset. *J. NeuroEngineering Rehabil.* **13**, 63 (2016).
- Patrick, E. E. et al. Modeling the volume of tissue activated in deep brain stimulation and its clinical influence: a review. *Front. Hum. Neurosci.* **18**, 1333183 (2024).
- Sundt, D., Gamper, N. & Jaffe, D. B. Spike propagation through the dorsal root ganglia in an unmyelinated sensory neuron: a modeling study. *J. Neurophysiol.* **114**, 3140–3153 (2015).
- Ben-Menachem, E. et al. Vagus nerve stimulation for treatment of partial seizures: 1. a controlled study of effect on seizures. First International Vagus Nerve Stimulation Study Group. *Epilepsia* **35**, 616–626 (1994).
- Sackeim, H. A. et al. Vagus nerve stimulation (VNS) for treatment-resistant depression: efficacy, side effects, and predictors of outcome. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* **25**, 713–728 (2001).
- Dawson, J. et al. Vagus nerve stimulation paired with rehabilitation for upper limb motor function after ischaemic stroke (VNS-REHAB): a randomised, blinded, pivotal, device trial. *Lancet Lond. Engl.* **397**, 1545–1553 (2021).
- De Ferrari, G. M. et al. Chronic vagus nerve stimulation: a new and promising therapeutic approach for chronic heart failure. *Eur. Heart J* **32**, 847–855 (2011).
- Boon, P., Vonck, K., De Reuck, J. & Caemaert, J. Vagus nerve stimulation for refractory epilepsy. *Seizure* **10**, 448–455 (2001).
- Blanz, S. L. et al. Spatially selective stimulation of the pig vagus nerve to modulate target effect versus side effect. *J. Neural Eng.* **20**, 016051 (2023).
- Ackermann, D. M., Bhadra, N., Foldes, E. L., Wang, X.-F. & Kilgore, K. L. Effect of nerve cuff electrode geometry on onset response firing in high-frequency nerve conduction block. *IEEE Trans. Neural Syst. Rehabil. Eng. Publ. IEEE Eng. Med. Biol. Soc.* **18**, 658–665 (2010).
- Miles, J. D., Kilgore, K. L., Bhadra, N. & Lahowetz, E. A. Effects of ramped amplitude waveforms on the onset response of high-frequency mammalian nerve block. *J. Neural Eng.* **4**, 390–398 (2007).
- Gerges, M. et al. Frequency and amplitude transitioned waveforms mitigate the onset response in high frequency nerve block. *J. Neural Eng.* **7**, 066003 (2010).

39. Peña, E., Pelot, N. A. & Grill, W. M. Non-monotonic kilohertz frequency neural block thresholds arise from amplitude- and frequency-dependent charge imbalance. *Sci. Rep.* **11**, 5077 (2021).
40. Golabek, J., Schiefer, M., Wong, J. K., Saxena, S. & Patrick, E. Artificial neural network-based rapid predictor of biological nerve fiber activation for DBS applications. *J. Neural Eng.* **20**, 016001 (2023).
41. Settell, M. L. et al. Functional vagotomy in the cervical vagus nerve of the domestic pig: implications for the study of vagus nerve stimulation. *J. Neural Eng.* **17**, 026022 (2020).
42. Chu, W., Gao, X. & Sorooshian, S. A new evolutionary search strategy for global optimization of high-dimensional problems. *Inf. Sci.* **181**, 4909–4927 (2011).
43. Owens, J. D. et al. GPU Computing. *Proc. IEEE* **96**, 879–899 (2008).
44. Raina, R., Madhavan, A. & Ng, A. Y. Large-scale deep unsupervised learning using graphics processors. in *Proceedings of the 26th Annual International Conference on Machine Learning* 873–880. <https://doi.org/10.1145/1553374.1553486> (ACM, Montreal Quebec Canada, 2009).
45. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Preprint at <https://doi.org/10.48550/arXiv.1412.6980> (2017).
46. Hinton, G. E. Geoffrey Hinton Neural Networks for machine learning online course. (2014).
47. Thompson, N. et al. Anatomical and functional organization of cardiac fibers in the porcine cervical vagus nerve allows spatially selective efferent neuromodulation. 2024.01.09.574861 Preprint at <https://doi.org/10.1101/2024.01.09.574861> (2024).
48. Veraart, C., Grill, W. M. & Mortimer, J. T. Selective control of muscle activation with a multipolar nerve cuff electrode. *IEEE Trans. Biomed. Eng.* **40**, 640–653 (1993).
49. Grill, W. M. & Mortimer, J. T. Quantification of recruitment properties of multiple contact cuff electrodes. *IEEE Trans. Rehabil. Eng.* **4**, 49–62 (1996).
50. Freeberg, M. J. et al. Intraoperative responses may predict chronic performance of composite flat interface nerve electrodes on human femoral nerves. *IEEE Trans. Neural Syst. Rehabil. Eng.* **27**, 2317–2327 (2019).
51. Taccola, G. et al. Using EMG to deliver lumbar dynamic electrical stimulation to facilitate cortico-spinal excitability. *Brain Stimul. Basic Transl. Clin. Res. Neuromodulation* **13**, 20–34 (2020).
52. Taccola, G. et al. Acute neuromodulation restores spinally-induced motor responses after severe spinal cord injury. *Exp. Neurol.* **327**, 113246 (2020).
53. Culaclii, S. et al. A biomimetic, SoC-based neural stimulator for novel arbitrary-waveform stimulation protocols. *Front. Neurosci.* **15**, 697731 (2021).
54. Li, X. et al. Physics-supervised deep learning-based optimization (PSDLO) with accuracy and efficiency. *Proc. Natl. Acad. Sci.* **120**, e2309062120 (2023).
55. Musselman, E. D., Cariello, J. E., Grill, W. M. & Pelot, N. A. ASCENT (Automated Simulations to Characterize Electrical Nerve Thresholds): a pipeline for sample-specific computational modeling of electrical stimulation of peripheral nerves. *PLoS Comput. Biol.* **17**, e1009285 (2021).
56. Pelot, N. A. et al. Quantified morphology of the cervical and sub-diaphragmatic vagus nerves of human, pig, and rat. *Front. Neurosci.* **14**, 1148 (2020).
57. Pelot, N. A., Goldhagen, G. B., Cariello, J. E. & Grill, W. M. Quantified morphology of the pig vagus nerve. <https://doi.org/10.26275/MAQ2-EI14> (2020).
58. Pelot, N. A. et al. Quantified morphology of the human vagus nerve with anti-claudin-1. <https://doi.org/10.26275/OFJA-GHOZ> (2021).
59. Musselman, E. D. et al. Pig-specific computational models of monopolar vagus nerve stimulation with a six-contact cuff electrode. <https://doi.org/10.26275/HYBX-GGBT> (2023).
60. Musselman, E. D., Grill, W. M. & Pelot, N. A. Validated computational models predict vagus nerve stimulation thresholds in preclinical animals and humans. <https://doi.org/10.26275/VDPW-RJQU> (2023).
61. Paszke, A. et al. PyTorch: An imperative style, high-performance deep learning library. in *Advances in neural information processing systems* **32** (eds. Wallach, H. et al.) 8024–8035 (Curran Associates, Inc., 2019).
62. Gabriel, E. et al. Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation. in *Recent Advances in Parallel Virtual Machine and Message Passing Interface* (eds. Kranzlmüller, D., Kacsuk, P. & Dongarra, J.) vol. **3241** 97–104 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2004).
63. Gruslys, A., Munos, R., Danihelka, I., Lanctot, M. & Graves, A. Memory-Efficient Backpropagation Through Time. (2016).
64. Kreuz, T., Mulansky, M. & Bozanic, N. SPIKY: a graphical user interface for monitoring spike train synchrony. *J. Neurophysiol.* **113**, 3432–3445 (2015).
65. Settell, M. L. et al. In vivo visualization of pig vagus nerve “vago-tomy” using ultrasound. *Front. Neurosci.* **15**, 676680 (2021).
66. Storn, R. & Price, K. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *J. Glob. Optim.* **11**, 341–359 (1997).
67. Ahmad, M. F., Isa, N. A. M., Lim, W. H. & Ang, K. M. Differential evolution: a recent review based on state-of-the-art works. *Alex. Eng. J.* **61**, 3831–3872 (2022).
68. Price, K., Storn, R. M. & Lampinen, J. A. *Differential Evolution: A Practical Approach to Global Optimization*. <https://doi.org/10.1007/3-540-31306-0> (Springer-Verlag, Berlin Heidelberg, 2005).
69. Biedrzycki, R., Arabas, J. & Jagodziński, D. Bound constraints handling in differential evolution: an experimental study. *Swarm Evol. Comput.* **50**, 100453 (2019).
70. Zhang, J. & Sanderson, A. C. JADE: adaptive differential evolution with optional external archive. *IEEE Trans. Evol. Comput.* **13**, 945–958 (2009).
71. Wright, L. Ranger – a synergistic optimizer. *GitHub repository* (2019).
72. Liu, L. et al. On the Variance of the Adaptive Learning Rate and Beyond. Preprint at <https://doi.org/10.48550/arXiv.1908.03265> (2021).
73. Zhang, M. R., Lucas, J., Hinton, G. & Ba, J. Lookahead Optimizer: k steps forward, 1 step back. Preprint at <https://doi.org/10.48550/arXiv.1907.08610> (2019).
74. Yong, H., Huang, J., Hua, X. & Zhang, L. Gradient Centralization: A New Optimization Technique for Deep Neural Networks. Preprint at <https://doi.org/10.48550/arXiv.2004.01461> (2020).
75. Holm, S. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **6**, 65–70 (1979).
76. Hussain, M. A., Grill, W. M. & Pelot, N. A. Data and scripts from: highly efficient modeling and optimization of neural fiber responses to electrical stimulation. *Duke Research Data Repository* <https://doi.org/10.7924/r48g8tf24> (2024).
77. Hussain, M. A., Grill, W. M. & Pelot, N. A. minhajh/cajal: v1.0.0. *Zenodo* <https://doi.org/10.5281/zenodo.12753671> (2024).
78. Davis, C. J., Musselman, E. D., Grill, W. M. & Pelot, N. A. Fibers in smaller fascicles have lower activation thresholds with cuff electrodes due to thinner perineurium and smaller cross-sectional area. *J. Neural Eng.* **20**, 026032 (2023).

Acknowledgements

Funding was provided through NIH SPARC OT2 OD025340 (M.A.H., W.M.G., N.A.P.), NIH SPARC 75N98022C00018 (M.A.H., W.M.G., N.A.P.), and the Duke-Coulter Translational Partnership (W.M.G., N.A.P). We thank the Duke Compute Cluster and Duke Research Computing (Katie Kilroy, Tom Milledge, Mike Newton, and others) for their expertise and assistance in deploying and managing the high-performance hardware. We thank Ishani Raha for implementing the shrunk ImThera cuff. We

thank Nick Merola for validating the installation and deployment processes. We thank Brandon Thio for helpful discussions and suggesting useful validations for the surrogate model. We thank Dr. David Carlson for helpful discussions.

Author contributions

M.A.H, W.M.G., and N.A.P. conceived and designed the research framework. M.A.H. designed and implemented the surrogate modeling approach, NEURON modeling code, and training and optimization algorithms. M.A.H. generated the data. M.A.H wrote the initial draft of the paper. M.A.H., W.M.G., and N.A.P. reviewed and edited the paper.

Competing interests

M.A.H, W.M.G., and N.A.P. are co-inventors on pending patent application 63/318,491 | DU7793PROV: SYSTEMS AND METHODS FOR PREDICTION AND DESIGN OF NEURAL STIMULATION, which relates to surrogate modeling approaches and optimization applications thereof.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-024-51709-8>.

Correspondence and requests for materials should be addressed to Nicole A. Pelot.

Peer review information *Nature Communications* thanks Marco Bonizzato, Mahsa Shoaran and the other, anonymous, reviewer(s) for

their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024