A scalable framework for current steering at single-neuron resolution

Praful K. Vasireddy[1], Ramandeep S. Vilkhu[1], Amrith Lotlikar[1], Jeff B. Brown[1], Alex R. Gogliettino[2], A. J. Phillips[1], Ethan J. Kato[1], Pawel Hottowy[3], Alexander Sher[4], Alan M. Litke[4], Subhasish Mitra[1], Nishal P. Shah[1,5,6] & E. J. Chichilnisky[2]

[1] Department of Electrical Engineering, Stanford University, Stanford CA, USA. [2] Departments of Neurosurgery, Ophthalmology, and Hansen Experimental Physics Laboratory, Stanford University, Stanford CA, USA. [3] AGH University of Krakow, Faculty of Physics and Applied Computer Science, 30-059 Krakow, Poland. [4] Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA, USA [5] Department of Electrical and Computer Engineering, Rice University, Houston TX, USA [6] Neuroengineering Initiative, Rice University, Houston TX, USA

To restore and augment complex functions such as vision, electrical neural interfaces require the ability to selectively activate neurons at cellular resolution, a challenge given the limited electrode density of typical neural interfaces. Current steering – passing current through multiple electrodes simultaneously to target locations between electrodes – could enhance resolution, but is intractable if all stimulation patterns must be tested exhaustively. Here, we develop and experimentally validate a framework for scalable cellular-resolution current steering in the macaque and human retina. Current was passed through three neighboring electrodes simultaneously, revealing a variety of linear and nonlinear neural responses. These observations were validated with biophysical simulations and accurately captured by a model in which spikes are initiated at one or more sites and the neuron produces a spike via a logical OR operation across sites, resulting in greater nonlinearity with more initiation sites. Finally, an adaptive sampling procedure was used to efficiently identify selective stimulation patterns, enabling significant calibration speedups. Together these advances represent the first scalable approach to cellular-resolution current steering for future neural interfaces.

Introduction

Electrical stimulation of neurons can be used to treat loss of nervous system function due to degeneration or trauma. Notable examples include deep brain stimulation for Parkinson's disease [1] and cochlear implants for hearing loss [2]. A major unsolved challenge is the development of high-resolution electronic retinal implants for vision restoration. In retinal degenerative diseases such as age-related macular degeneration (200 million affected worldwide [3]) and retinitis pigmentosa (1.5 million affected worldwide [4]), electrical stimulation can bypass the damaged photoreceptor layer and directly evoke electrical activity in retinal ganglion cells (RGCs) which transmit visual information to the brain. Although the approach has been validated conceptually, epiretinal implants that target RGCs for electrical stimulation have not yet restored high acuity vision. One reason is that vision is mediated by precise spatiotemporal spiking patterns of about 20 RGC types that send distinct visual signals to various brain targets [5–7]. Because the distinct RGC types are intermixed on the surface of the retina, single-cell resolution is required to reproduce naturalistic visual signals. However, this

goal is out of reach of existing epiretinal implants due to their relatively coarse spatial resolution [8].

A potential solution is to use large-scale, high-density arrays of small electrodes, along with closed-loop targeting of distinct RGC types, to more faithfully reproduce the neural code. Indeed, in *ex vivo* experiments using microelectrode array (MEA) stimulation and recording in peripheral macaque and human retinas [9,10], single-cell stimulation resolution is frequently achievable [11–14]. Unfortunately, this selectivity is not always possible, and is more limited in the central retina, which is responsible for high-resolution vision, because the spacing between RGCs is much smaller than typical MEA electrode pitch [12,15].

A possible way to overcome this limitation is *current steering*: simultaneous stimulation using multiple electrodes to more effectively target locations between electrodes and potentially more faithfully mimic the neural code of the retina. Current steering is well-established in neural stimulation at coarse spatial resolution [16–21], but only a few studies have examined the effect of multi-electrode stimulation at single-neuron resolution [16,17,22]. Moreover, the work performed at cellular resolution required either exhaustive searches of all possible multi-electrode stimuli [17] or a specific geometric relationship between cells and electrodes to achieve selective activation [16,22], both of which are difficult to achieve in a clinical implant.

To overcome these challenges, we develop and validate a scalable, quantitative framework for understanding and predicting single-neuron responses to multi-electrode stimulation, and for efficiently searching the large space of all possible stimulation patterns to achieve selective stimulation of a given neuron while excluding others. Multi-electrode stimulation and recording with high-density MEAs was performed in isolated macaque and human retinas, and a multi-site model of neural activation was validated on data from experiments and biophysical simulations. The model permitted rapid calibration of single-neuron responses to patterned electrical stimuli despite nonlinearities due to multi-site activation [17,23,24]. To efficiently identify selective stimuli among the large number of possible multi-electrode stimulation patterns, data-driven regression models coupled with intelligent sampling algorithms were used to speed calibration. These results provide effective methods for employing current steering at cellular resolution and circuit scale to improve the resolution of neural implants.

Results

To probe the potential of current steering for focally activating RGCs, multi-electrode stimulation and recording was deployed in an *ex vivo* preparation of human and macaque retinas using a 512-channel microelectrode array (MEA) system with 30 µm pitch. With this approach, the activation of each targeted RGC in a population of several hundred simultaneously-recorded RGCs was measured. Triphasic, charge-balanced current pulses (50 µs phase duration, 0.1-1.8 µA amplitude) were applied to three adjacent electrodes (an *electrode triplet*) simultaneously, and current amplitudes were varied independently on all three electrodes (Fig. 1A). Three-electrode stimulation was used because two-electrode stimulation can only steer electrical activity on a line between the two electrodes [16]. To characterize the stochastic activation of

neurons (see [25]), each stimulus was repeated 20 times, evoked spikes were identified in the recorded traces using a custom spike sorting algorithm (see Methods), and an activation probability was calculated from these spikes (Fig. 1B). For each cell, the current pulse amplitude combinations in the three electrodes that produced 50% probability of an evoked spike were estimated and defined as the *thresholds* for activation, and the organization of these thresholds in the three-dimensional space of electrical stimuli was then examined (Fig. 1B) to test hypotheses about the nature of activation.

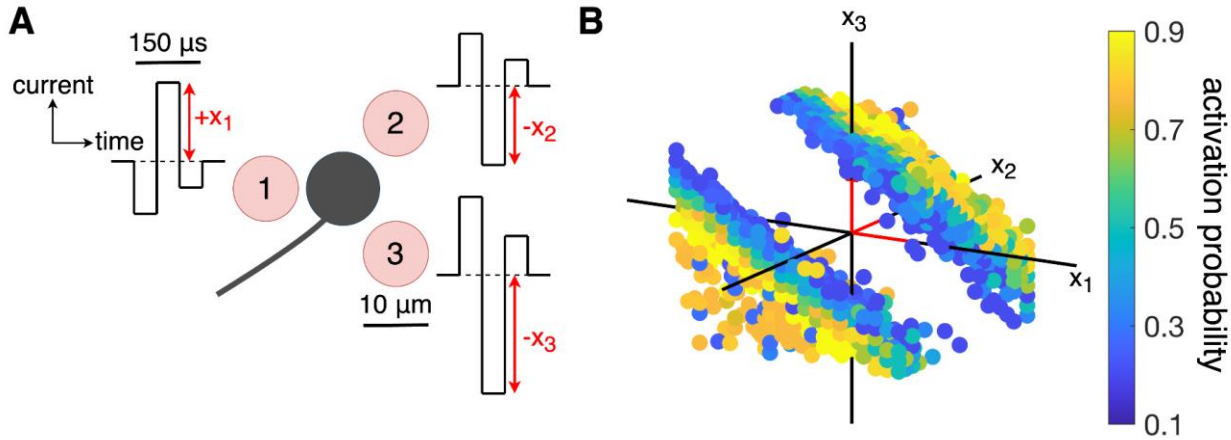

Figure 1. *Illustration of the experimental data collection.* ***A,*** Top-down schematic view of a single neuron (black) against three electrodes in the microelectrode array (light red) with electrode diameter and pitch of 10 µm and 30 µm, respectively. Triphasic, charge-balanced current pulses with 50 µs phase duration, relative phase amplitudes of (2:-3:1), and absolute amplitudes (second phase amplitudes) varying independently across electrodes are applied to all three electrodes simultaneously. ***B,*** Visualization of single-neuron responses to combinations of stimulation currents on three electrodes. Activation probabilities are calculated by averaging spiking responses across repeated trials of the same stimulus current combination. Red sections of axes mark the positive direction and are of length 1 µA. Only intermediate probabilities with 0<p<1 are shown for visual clarity.

Neurons exhibit both linear and nonlinear responses to multi-electrode stimulation

In the simplest case, currents from different electrodes could sum linearly to cause neural activation. Specifically, if currents (x_1, x_2, x_3) from three electrodes sum linearly to evoke a spike, then the collection of stimuli producing an activation probability of 50% satisfy the relation $\sigma(w_0+w_1x_1+w_2x_2+w_3x_3)=0.5$, where σ() is a sigmoidal function of spike probability as a function of current amplitude capturing the approximately Gaussian membrane voltage noise and consequent stochastic spike production [13,25]. Thus, linear summation predicts that the collection of measured thresholds (x_1, x_2, x_3) form a plane.

The results from 14 three-electrode activation tests revealed both linear and nonlinear neural responses to stimulation in various RGC types, retinal eccentricities, and species. In some cases, the thresholds formed approximately planar surfaces in the space of electrical stimuli (Fig. 2A-E). Because a charge-balanced stimulation was used, anodic-peak stimuli (positive current values in Fig. 2) often elicited spiking at similar threshold levels as cathodic-peak stimuli

(negative current values in Fig. 2), resulting in a second planar threshold surface (note that the two planes were not necessarily symmetric about the origin (see Fig. 2A,D)). This two-plane geometry thus indicates that, in some cases, current from the three electrodes summed approximately linearly to evoke spikes. In other cases, the collection of activation thresholds formed a curved surface (Fig. 2F-H), revealing nonlinear summation of currents across electrodes. The nonlinearity was sometimes sufficiently strong that the surface was closed, even within the limited range of current levels available (Fig. 2G,H). This structure confirms and extends previous findings with two-electrode stimulation (see [17]).

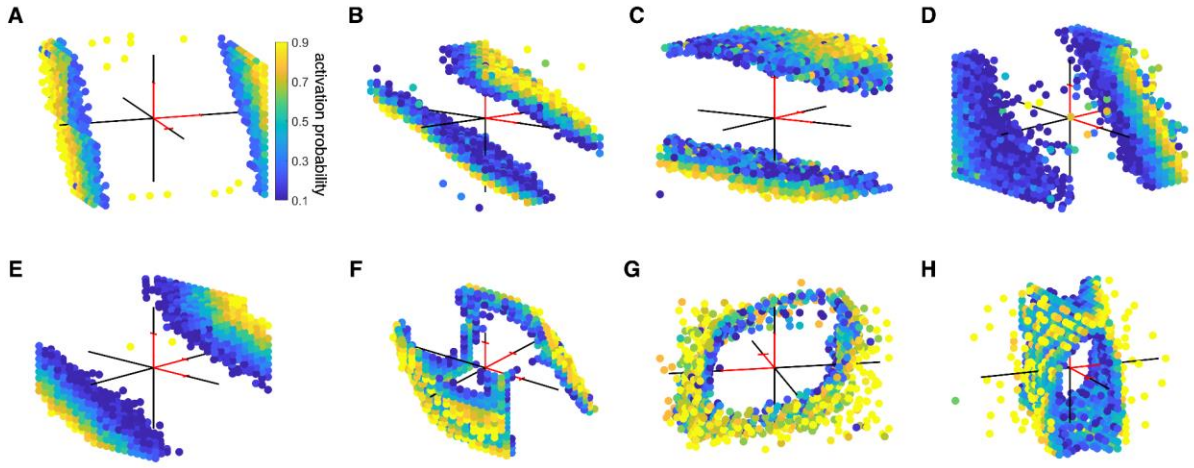

Figure 2. *Experimental data for three-electrode stimulation of individual neurons.* Macaque and human retinal ganglion cells of different cell types from various retinal preparations are stimulated with three electrodes simultaneously at independently varying current amplitudes (x_1, x_2, x_3), and their response profiles are measured. Red sections of axes mark the positive direction and are of length 1 μA. Only intermediate activation probabilities with $0<p<1$ are shown for visual clarity of the threshold surface. ***A,*** Macaque central ON parasol. ***B,*** Macaque peripheral OFF parasol. ***C,*** Macaque central OFF parasol. ***D,*** Human peripheral OFF midget. ***E,*** Human central ON parasol. ***F,*** Macaque peripheral ON parasol. ***G,*** Macaque peripheral ON parasol. ***H,*** Macaque central OFF parasol.

Biophysical simulations reveal linear and nonlinear responses to multi-electrode stimuli depending on electrode position

To test whether the linear and nonlinear activation surfaces observed with different electrode geometries were consistent with known properties of RGCs, a biophysical model was developed with cellular properties and MEA geometry approximately matching the experiments (see Methods). The biophysical model was developed in the NEURON simulation environment [26] and consisted of dendritic, somatic, and axonal compartments. The ion channel parameters in each compartment and cytoplasmic (axial) resistivity were set to match those of primate RGCs, and the temperature was chosen to match experimental conditions. Stimulating electrode geometry and pulse waveforms were arranged to match the experiments (see [23] for details).

Both linear and nonlinear responses were observed in the simulation, and the degree of nonlinearity depended on the electrode-cell geometry in a way that could be understood based on a multi-site activation model. Specifically, electrodes placed near the cell body produced

nearly linear responses (planar surfaces) (Fig. 3B) while electrodes placed near the distal axon produced nonlinear responses (curved surfaces) (Fig. 3A,C). Stimulating electrodes placed near the soma most often initiated spikes at the sodium-channel band, a location on the proximal axon with the highest density of voltage-gated sodium channels, whereas electrodes placed near the distal axon could initiate spikes in several adjacent axonal segments [23]. The rotation angle of the electrodes relative to the neuron affected the orientation of the threshold surfaces but did not significantly alter the degree of nonlinearity (Fig. 3A,B). Increasing the distance between the electrodes and the cell produced increasingly linear response surfaces (Fig. 3C), because electrodes had electrical access to fewer spike initiation sites within a given stimulating current amplitude range [23]. Similar findings were observed using biphasic rather than triphasic stimulating current pulses (data not shown). These features are consistent with multi-site activation of RGCs (see [23]), and confirm that the response of a neuron to multi-electrode stimulation is determined by the geometry of the electrodes relative to the neuron.

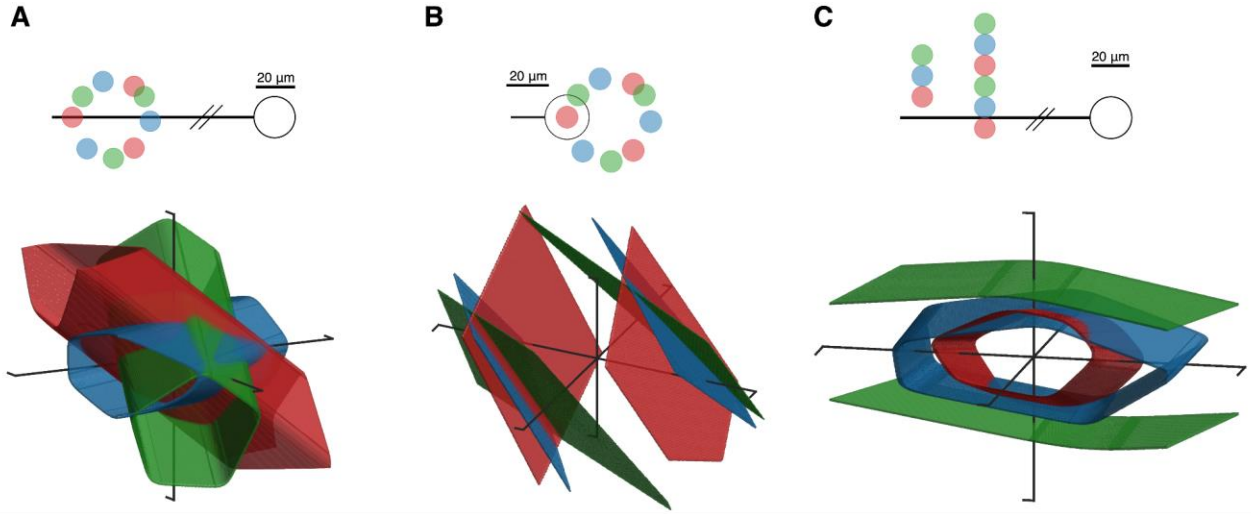

Figure 3. *Biophysical simulation responses for three-electrode stimulation.* Colored surfaces represent the 50% activation probability boundary. Schematic headings show a top-down (x-y) view of the relative geometry between the neuron and stimulating electrodes, with the color of the electrodes corresponding to the color of the 50% activation probability boundary from stimulating with those electrodes. The electrodes lie on a fixed-z plane 12 µm above the distal axon and 40 µm above the soma center. ***A,*** Stimulation at the distal axon with varying rotation angle. ***B,*** Stimulation near the somatic compartment with varying rotation angle. ***C,*** Stimulation at the distal axon with varying transverse offset.

A multi-site model of neural activation quantitatively captures linear and nonlinear behaviors

To test whether linear and nonlinear activation can be understood and harnessed for targeting neurons with electrical stimulation, a parametric model of multi-site activation is presented below. This model is biophysically motivated, can be rapidly fitted to experimental data, and can be used to accurately predict single-neuron responses to untested multi-electrode stimuli.

The multi-site model is informed by the possibility of an electrically elicited spike originating from one of many initiation sites [17,23]. In the model, each site linearly sums the applied stimulation

current on many electrodes and this sum controls activation probability via a sigmoidal relationship, which reflects voltage noise in the cell. Any of the initiation sites can produce a spike. The multi-site model is mathematically described by the predicted probability of spiking:

$$\textbf{(1)}\ p(x) \equiv p(cell\ spikes \mid x) = 1 - \prod_{i=1}^{m} (1 - \sigma(w_{i,0} + w_i^T x)).$$

Here, x is the vector of stimulation currents applied to a collection of electrodes, m is the number of distinct activation sites, $\sigma()$ is the sigmoid function that controls activation probability [13], and $w_{i,0}$ and w_i are biases and weights for each site. Simply put, this model represents a logical OR operation between many linear spike initiation sites.

To test the multi-site model, the three-electrode activation threshold surfaces were fitted with the linear weights and biases as free parameters (see Methods). The fitted model accurately captured both linear (Fig. 4B) and nonlinear (Fig. 4A,C) activation surfaces, with the mean absolute error between model predictions and observed activation probability significantly lower than that of a linear null model (see Methods; Fig. 4D-F). To provide a comparison to a more flexible and general model, Gaussian process regression (GPR) was tested [27,28] (see Methods). In linear cases (Fig. 4B,H), the multi-site model slightly outperformed the GPR model; in nonlinear cases, the two models performed similarly (Fig. 4A,C,G,J), although the multi-site model typically required a larger number of calibrating stimulations to achieve maximum performance (Fig. 4G,J).

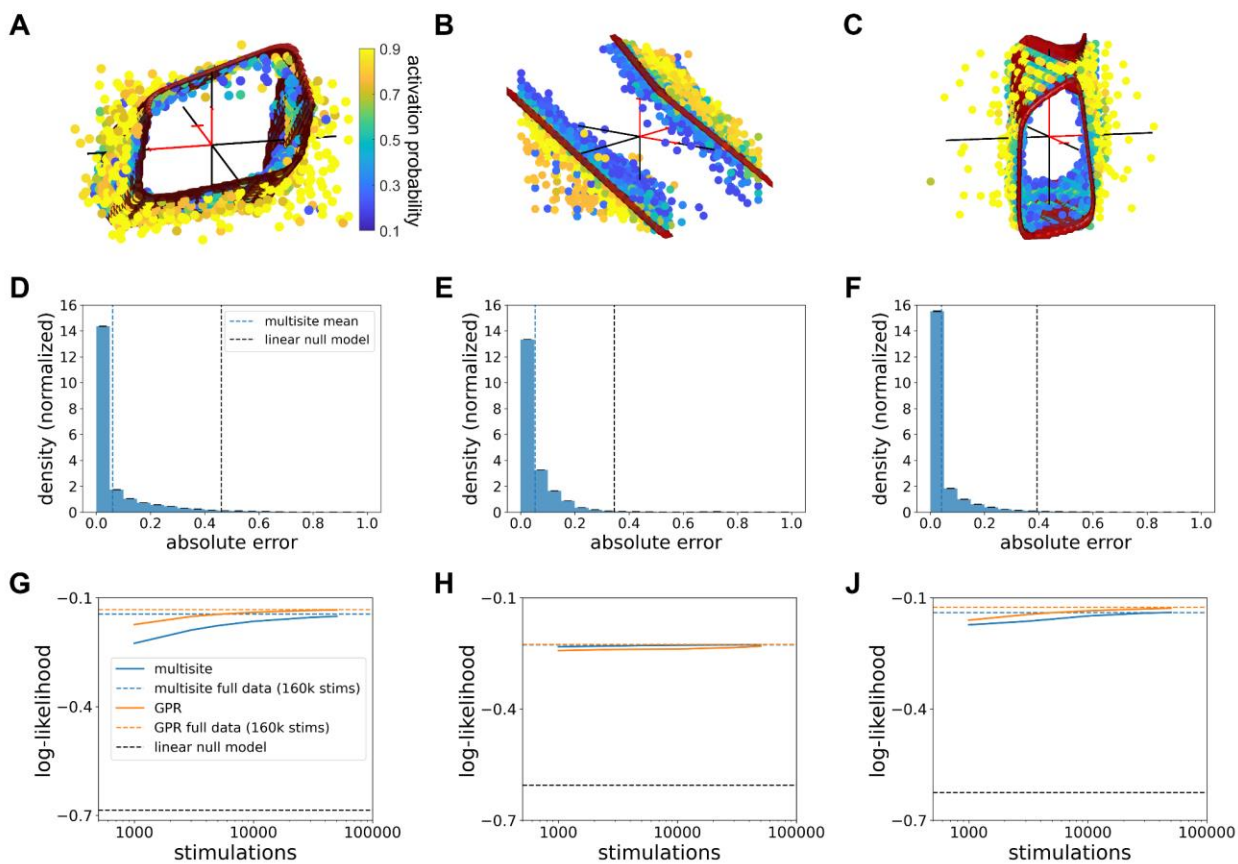

Figure 4. *Multi-site activation model fitting to experimental data.* Red surfaces represent the 50% activation probability surface predicted by the multi-site model fit to the full experimental data. Red

sections of axes mark the positive direction and are of length 1 μA. ***A,*** Nonlinear three-electrode stimulation experimental data (macaque peripheral ON parasol) with multi-site model 50% surface for *m* = 4 activation sites. ***B,*** Linear three-electrode stimulation experimental data (macaque peripheral OFF parasol) with multi-site model 50% surface for *m* = 2 activation sites. ***C,*** Nonlinear three-electrode stimulation experimental data (macaque central OFF parasol) with multi-site model 50% surface for *m* = 4 activation sites. ***D–F,*** Distribution of error between fitted multi-site model predictions and experimental data, for neurons from *A–C*. Blue vertical dashed line represents the mean absolute error achieved by the multi-site model, while black vertical dashed line is the mean absolute error achieved using a linear null model. 100 refittings were used for error bars. ***G–I,*** Log-likelihood between predictions from multi-site model or Gaussian process regression (GPR) model and experimental data as a function of number of calibrating stimulations, for neurons from *A–C*. Colored horizontal dashed lines represent the log-likelihood level achieved by fitting to the full data (160,000 stimulations), and black horizontal dashed line represents the log-likelihood level achieved using a linear null model. 100 samples were used for confidence intervals.

The multi-site model was also tested by fitting to threshold surfaces produced using biophysical simulations with realistic neuron-cell geometries. Simulated surfaces with electrode geometries producing a range of linear (Fig. 5B) to highly nonlinear (Fig. 5A,C) behaviors were accurately fitted by the model. As a test of generalization to larger numbers of electrodes, simulations of four-electrode stimulation were performed. The responses were captured by the multi-site model, and three-dimensional sections of activation surfaces again revealed accurate model performance (see Extended Data Fig. 1).

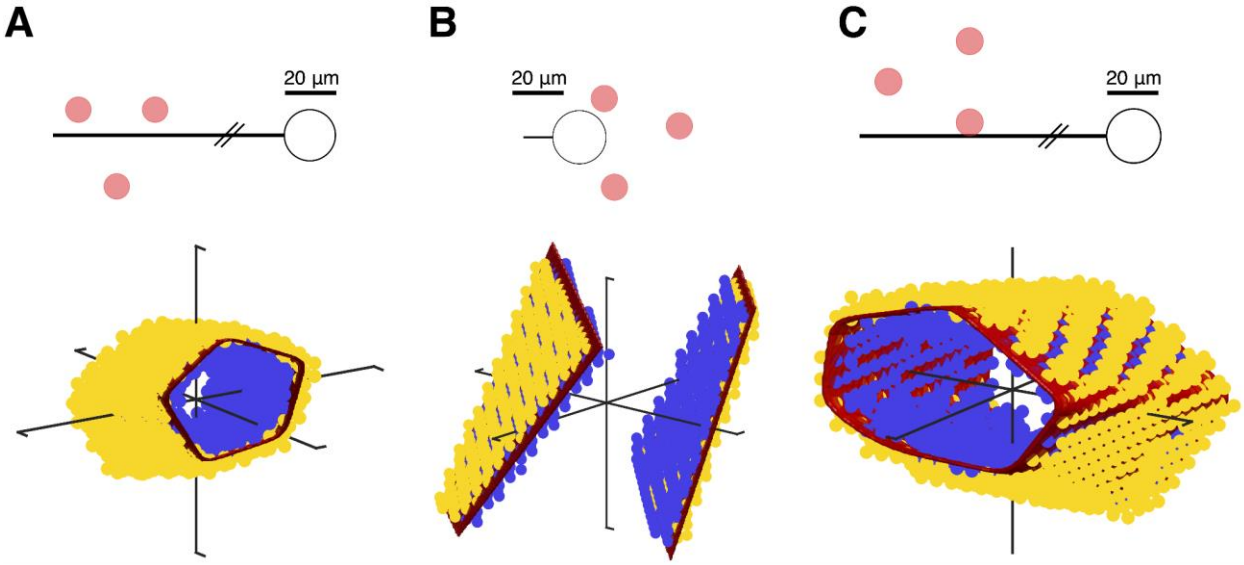

Figure 5. *Multi-site activation model fitting to biophysical simulation data.* Red surfaces represent the 50% activation probability surface predicted by the multi-site model fit to the biophysical simulation data. Schematic headings show a top-down (x-y) view of the electrodes (red) relative to the neuron (black). The electrodes lie on a fixed-z plane 12 μm above the distal axon and 40 μm above the soma center. ***A,*** Nonlinear three-electrode stimulation biophysical simulation data with multi-site model 50% surface for *m* = 8 activation sites. ***B,*** Linear three-electrode stimulation biophysical simulation data with multi-site model 50% surface for *m* = 2 activation sites. ***C,*** Nonlinear three-electrode stimulation biophysical simulation data with multi-site model 50% surface for *m* = 8 activation sites.

Multi-electrode stimulation enables significant improvements in single-neuron selectivity

A central goal of current steering is to selectively activate some neurons without activating others, to most effectively match the neural code. Geometrically, for triplet stimulation, this corresponds to finding regions of the three-dimensional space of current amplitudes that are outside the 50% threshold surface of a target cell but within the threshold surface of other non-target cells (Fig. 6A).

To probe the selectivity that can be achieved with triplet stimulation, a metric of selectivity was developed, and the effectiveness of triplet stimulation was evaluated. Selectivity was quantified by defining the *selectivity index*. For a collection of neurons $c_1,...,c_K$ that can be activated using electrical stimulus A, the selectivity index, $S_{i,A}$, for a given target cell, c_i , with this stimulus is given by:

$$\textbf{(2)}\ S_{i,A} = p_{i,A}(1 - max_{j \neq i} p_{j,A}),$$

where $p_{i,A}$ is the activation probability of the target cell c_i using stimulus A, and $p_{j,A}$ is the activation probability of a given non-target cell c_j using stimulus A, taking the maximum of this value across all non-target cells to represent the worst case. The metric ranges from 0 to 1, with 1 indicating perfect selectivity for the target cell. For example, a stimulus which elicits 70% probability of spiking in the target cell and 30% probability of spiking in the worst-case non-target cell yields a selectivity index of 0.49. The maximum of the selectivity index across all stimuli using a specific electrode triplet was measured. Similarly, the maximum selectivity index across all stimuli using any single electrode in the electrode triplet was measured for comparison. By definition triplet stimulation is always at least as selective as single-electrode stimulation because single-electrode stimulation is a subset of triplet stimulation.

Triplet stimulation frequently provided a significant improvement in selectivity as compared to single-electrode stimulation. For example, in the case of a specific human OFF midget RGC (Fig. 6B), selectivity indices of 0.37, 0.018, and 0.068 were obtained using each of the three electrodes individually. However, using all three electrodes simultaneously led to a stimulation pattern yielding nearly perfect selectivity of 0.95.

The results accumulated across 35 human and macaque RGCs tested with three-electrode vs. single-electrode stimulation suggest that current steering can be very effective (Fig. 6C). Using only single-electrode stimulation, 10/35 cells (28.6%) had a selectivity index greater than 0.5, compared to 19/35 cells (54.3%) using triplet stimulation. The remaining 16/35 cells (45.7%) tested could not be selectively activated with either single- or three-electrode stimulation. This limitation highlights the need for efficient sampling procedures to rapidly identify selective stimulation patterns on many groups of three electrodes (see below).

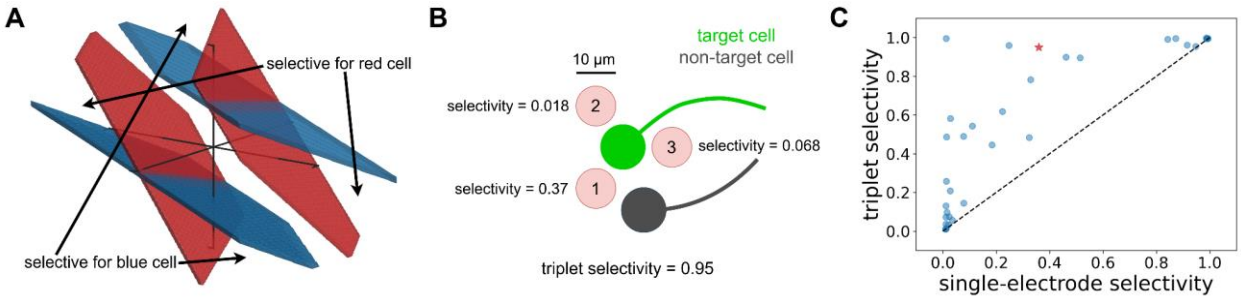

Figure 6. *Selectivity obtained from multi-electrode stimulation.* ***A,*** Schematic of 50% activation probability surfaces (red and blue surfaces) for two neurons stimulated by the same electrode triplet, with selective regions in stimulus space for each cell indicated. Surfaces were extracted from measurements from macaque peripheral RGCs. ***B,*** Schematic of selectivity improvement for an example with two human peripheral RGCs. Red numbered circles denote electrodes, with relative positions of non-target cell (gray; OFF parasol) and target cell (green; OFF midget) shown. Each electrode is labelled with its maximum achievable selectivity index using single-electrode stimulation. ***C,*** Maximum selectivity index obtained with triplet stimulation compared to maximum selectivity index obtained with single-electrode stimulation on the same set of electrodes, accumulated across 35 RGCs from human and macaque retinal preparations. Red star indicates the example shown in panel B.

Rapid identification of selective multi-electrode stimuli using data-driven models

The effectiveness of current steering for selectivity relies on efficient calibration, because it is unknown *a priori* which groups of electrodes will selectively stimulate a neuron, and because exhaustive calibration of current steering across a many groups of electrodes in a large-scale MEA is likely impractical in a clinical setting. Below, a method is developed to rapidly identify selective stimuli, and its performance is evaluated.

First, to reduce dimensionality, a measure of aggregate selectivity was defined to be the selectivity index obtained considering the most selective cell as the target (see Methods Eq. 5). Then, the aggregate selectivity across all tested three-electrode stimuli was computed (see Fig. 7A), defining the space that must be searched to determine the most selective stimuli.

An obvious candidate approach to rapidly identifying selective stimuli would be to use the multi-site activation model fitted to limited experimental data and evaluated on all possible stimuli in this space. To test this possibility, the multi-site model was fitted to either the entire data set or a subset of the data. The model captured the large-scale features of the measured selectivity landscape when it was fitted with all the data (Fig. 7B, Fig. 8B), and its predictions changed qualitatively little even with significant subsampling (Fig. 7C). However, some of the details of the experimentally determined selectivity landscape (Fig. 7A) were not fully captured by the multi-site model, especially in the most relevant highly selective regions (see Fig. 8A).

As a more flexible alternative, a model of selectivity based on GPR was tested by again fitting the selectivity landscape with relatively few measurements and predicting selective stimuli. Compared to the multi-site model, GPR was able to capture more precise details of the selectivity landscape using all the recorded data (Fig. 7D, Fig. 8A). However, due to the

flexibility of the GPR model, random subsampling often led to less accurate fits in highly selective regions (Fig. 7E, Fig. 8A). This finding suggests that GPR fits could be improved by adaptive sampling procedures that preferentially explore informative regions of the selectivity landscape.

Accordingly, the GPR result was used to develop an efficient closed-loop sampling scheme based on a method known as level set estimation (LSE; see Methods). Briefly, little-explored regions of stimulus space with the possibility of yielding high selectivity were preferentially sampled, and previously explored regions with known extreme selectivity (high or low) were avoided to encourage learning about other regions. Using this sampling scheme (denoted as GPR+LSE), highly selective regions were well estimated even with significant subsampling (Fig. 7F, Fig. 8A).

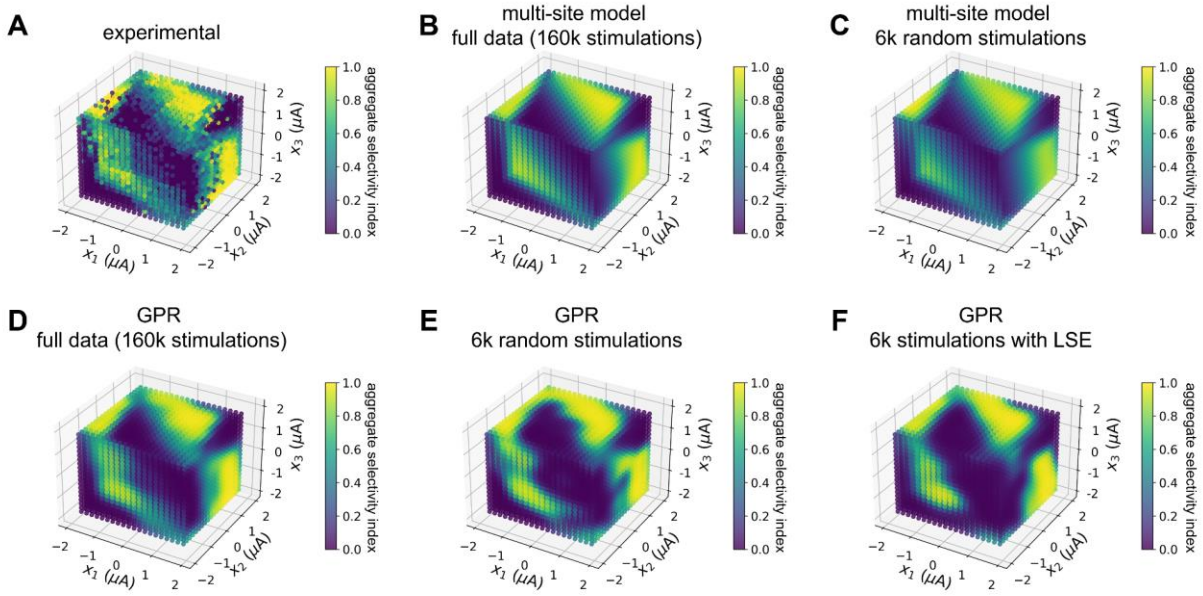

Figure 7. *Aggregate selectivity landscapes fitted with multi-site model and Gaussian process regression (GPR).* ***A,*** Experimental aggregate selectivity profile for three-electrode stimulation, aggregated across three macaque central RGCs stimulated by the same electrode triplet. Remaining panels show fitted selectivity predictions using ***B,*** multi-site model, exhaustive data (160k calibrating stimulations), ***C,*** multi-site model, 6k randomly sampled calibrating stimulations, ***D,*** GPR, exhaustive data (160k calibrating stimulations), ***E,*** GPR, 6k randomly sampled calibrating stimulations, ***F,*** GPR, 6k stimulations sampled with the GPR+LSE procedure.

The GPR+LSE framework substantially reduced total calibration time to accurately determine selective stimulation patterns. Highly selective regions (selectivity index greater than 0.9) were estimated more accurately and with fewer measurements using the GPR+LSE procedure as compared to random sampling (Fig. 8A, Extended Data Fig. 2). In a specific example with three macaque central RGCs stimulated by a specific group of three electrodes, sampling with GPR+LSE required ~10,000 calibrating stimulations to achieve accuracy within 10% of the accuracy from performing GPR on the exhaustive data, compared to ~15,000 stimulations required to reach the same accuracy level with random sampling (Fig. 8A). Both procedures showed substantial improvement over using the multi-site model, requiring ~2,000 stimulations

with GPR+LSE and ~3,000 stimulations with random sampling to achieve the same accuracy as the multi-site model fitted to the exhaustive data (160k stimulations; Fig. 8A). However, because the GPR+LSE procedure prioritized estimating selective regions, unselective regions were more poorly estimated with this method than with random sampling, leading to higher error across the whole selectivity landscape (Fig. 8B). This drawback was alleviated by a hybrid scheme that mixed LSE sampling with uniform random sampling (not shown).

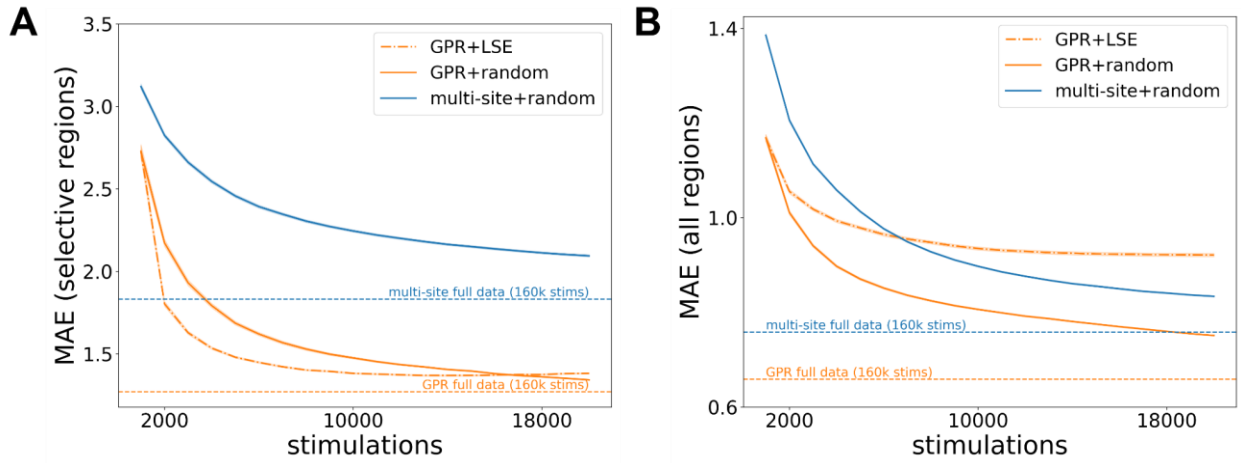

Figure 8. *Gaussian process regression (GPR) and level set estimation (LSE) scheme performance.* ***A,*** Mean absolute error (MAE) in log-odds units between model predictions and experimental data in high selectivity regions (the union of regions with selectivity index greater than 0.9 as estimated by the model and as measured by experimental data) as a function of number of calibrating stimulations, with GPR+LSE sampling compared to GPR with uniform random sampling. MAE of the multi-site model fitted to randomly subsampled data is also shown. 100 restarts were used for confidence intervals. ***B,*** Mean absolute error (MAE) in log-odds units between model predictions and experimental data in all regions as a function of number of calibrating stimulations, with GPR+LSE sampling compared to GPR with uniform random sampling. MAE of the multi-site model fitted to randomly subsampled data is also shown. 100 restarts were used for confidence intervals.

Discussion

Experimental investigation of cellular-resolution current steering with groups of three electrodes yielded two major findings. First, a simple biophysically-inspired model of multi-site activation captured the main trends in linear and nonlinear neural activation, permitting a simple and accurate prediction of neural responses to arbitrary electrical stimulus patterns with three electrodes. Second, the ability to selectively activate one neuron without others was enhanced by current steering, and empirical models permitted efficient sampling to identify the most selective stimuli. These findings suggest that current steering could be used to more precisely reproduce the neural code at cellular resolution in future electronic implants.

The multi-site model of neural activation with multiple electrodes provided a biophysically motivated way to understand the linear and nonlinear properties of electrical activation with relatively few assumptions and measurements. Biological tissue conductivities are approximately linear at the frequencies used here (5-100 Hz) [29,30], meaning that a scalar linear weight can be assigned between the stimulating current on an electrode and the resulting

change in extracellular voltage at a specific region on a passive cell membrane. Other functions of the membrane potential relevant to eliciting spikes extracellularly, such as its second directional derivative along the cell membrane (activating function; [31]), are also linear in the stimulating current. However, in contrast to normal physiological spike initiation, which occurs at the sodium channel band in the proximal axon [32], extracellular stimulation of neurons may initiate spikes at various locations on the cell membrane depending on the geometric arrangement of the stimulating electrodes relative to the cell, as shown by inadvertent stimulation of RGC axons with epiretinal electrodes [11]. The ability to generate a spike at each of several locations leads to nonlinear neural responses to multi-electrode stimuli, even if spike initiation at a single location is linear [23]. Importantly, nonlinear interactions are clearly visible using low currents and short distances between electrodes and neurons, but are much more difficult to observe with electrodes far from the cell [23,33]. The multi-site model can be fitted to such complex nonlinear data accurately, with less computational cost than other nonlinear fitting methods (see Methods).

Current steering with electrode triplets led to substantial increases in single-neuron selectivity, and two methods for estimating selectivity landscapes using subsampled data were contrasted. Using the multi-site model to estimate selectivity landscapes required combining several model fits to individual neurons together to form a selectivity landscape prediction. Though the multi-site model was accurate for fitting stimulation responses of individual neurons, combining the model fits was prone to propagation of fitting errors and generally captured only the large-scale features. In contrast, the flexibility of GPR was leveraged to more accurately estimate precise details of selectivity landscapes, particularly in regions of high selectivity. The uncertainty metric obtained from GPR also lent itself to the design of efficient calibration protocols to find selective stimuli. The tunability and robustness of GPR may allow for extensions to larger numbers of stimulating electrodes and to other neural interfaces.

Due to the lack of precise knowledge of the positions of neurons relative to the MEA combined with the variety of linear and nonlinear responses to multi-electrode stimuli, empirical calibration would likely be required for a future implant to exploit cellular-resolution current steering, and the methods described here could reduce calibration times to clinically feasible timescales. For example, assuming a 10 ms interval between successive stimulations to account for the neural refractory period, the GPR+LSE method required ~1 minute of calibration time per electrode triplet. In selectivity landscapes with no selective regions, the algorithm terminated earlier, preventing unnecessary searches that would yield no selective stimulation patterns. Across the entire MEA used in this work with ~1,000 distinct electrode triplets, assuming a maximum of 1 minute of calibration per triplet using GPR+LSE, the total calibration time is ~16.5 hours, a timescale which could be feasibly divided over several clinical sessions to improve the precision of neural stimulation.

The present work has several limitations and caveats. First, the proposed GPR+LSE algorithm was tested here by offline subsampling of exhaustively collected data rather than online in a true closed-loop experiment. In a true closed-loop approach, performing GPR and deciding which measurements to take next with LSE every iteration requires accurate, real-time spike sorting

with limited data, motivating the need for one-shot spike detection algorithms for electrical stimulation (in the presence of stimulation artifacts) that do not require repeated trials of the same stimulus (A. Lotlikar, personal communication). Online spike sorting and GPR fitting will also add computation time in an experiment, but schemes which batch measurements across distinct electrode triplets could eliminate computational downtime. Second, the present study focused on parasol and midget RCGs. Though these cell types are numerically dominant in the primate retina, there are many other cell types which are likely essential to natural vision, and it is possible that some of the findings would differ for those cell types. Third, the multi-electrode stimuli used were presented simultaneously, ignoring the possible benefits of temporally modulated stimuli for achieving high stimulation selectivity. Finally, while the parametric multi-site model accurately predicted and fitted experimental data, the numerical values of the fitted weights and biases may not have physiological significance, especially in nonlinear cases. Specifically, dendritic tree influences and interactions between regions of the cell membrane, particularly suppressive effects between adjacent regions of the membrane, produce complex stimulation response surfaces which do not correspond to a small number of discrete activation sites.

Multi-electrode current steering may be useful in neuroengineering paradigms outside of the retina, and the calibration scheme developed here could provide a general framework for selective electrical stimulation of individual neurons. Specifically, the multi-site model may be applicable to extracellular stimulation at single-neuron resolution in the brain, because the gray matter is predominantly composed of cell bodies and unmyelinated axons, as is the retina. The efficient and generalizable methods presented here are particularly relevant for the larger-scale neural probes currently being developed. As neural implants are increasingly used to restore and augment complex functionality, efficiently predicting and controlling the responses of individual neurons with patterned multi-electrode stimuli could become a cornerstone of brain-machine interfaces.

Methods

Experimental setup

A custom 512-electrode system [9] was used to stimulate and record from RGCs in isolated rhesus macaque (*Macaca mulatta*) and human retinas. Macaque eyes were obtained from terminally anesthetized animals euthanized during the course of research performed by other laboratories. Human eyes were obtained from brain-dead organ donors through Donor Network West. All procedures were performed in accordance with institutional and national guidelines and regulations. Briefly, the eyes were hemisected in room light following enucleation, the vitreous was removed and the posterior portion of the eye was kept in darkness in warm (33–35°C), oxygenated, bicarbonate buffered Ames solution (Sigma). Then, patches of retina ~2 mm on a side were isolated under infrared light, placed RGC side down on the microelectrode array (MEA), and superfused with oxygenated Ames solution. Central retinal preparations corresponded to regions with 1.5–4.5 mm (4.5–20°) eccentricity, while peripheral preparations ranged from 5–12 mm in eccentricity (22–56° temporal equivalent). The MEA was composed of

512 electrodes (diameter 5-10 μm) in an isosceles triangular lattice with center-to-center distance between adjacent electrodes either 30 μm or 33.5 μm, covering an area of 0.43 mm^2. Within an experiment, platinization of the electrodes produced relatively uniform noise, approximately 6% (standard deviation) across electrodes. A platinum wire encircling the recording chamber (~1 cm diameter) served as the distant return electrode. Voltage recordings were band-pass filtered between 43 and 5,000 Hz and sampled at 20 kHz. Spikes from individual RGCs in the voltage recordings obtained with visual stimulation of the retina were identified and sorted using Kilosort [34].

Visual stimulation and cell type classification

To identify the retinal ganglion cell types recorded, the retina was visually stimulated with a dynamic white noise stimulus, and the spike-triggered average (STA) stimulus was computed for each RGC, as previously described [35,36]. The STA summarizes the spatial, temporal, and chromatic properties of light response. In particular, clustering on the spatial (receptive field size) and temporal (time course) components of the STA was performed to identify distinct cell types, as previously described [37]. This process permitted rapid identification of ON and OFF parasol and midget RGCs, the four numerically dominant RGC types in primates.

Electrical images

The electrical image (EI) represents the average spatiotemporal pattern of voltage deflections produced on each electrode of the array during a spike from a given cell [38]. Electrical images were calculated from data recorded during visual stimulation and served as spatiotemporal templates for the spike waveforms of the cells to be detected during electrical stimulation.

The spatial positions of relevant cell compartments (axon, soma, dendrites) of a recorded cell relative to the electrode array were estimated using the electrical image. The shape of the spike waveform from the EI on an electrode was used for compartment identification: a triphasic waveform was taken to indicate recording from the axon, a biphasic waveform with a positive first phase was taken to indicate recording from the dendrites, and a biphasic waveform with a negative first phase was taken to indicate recording from the soma [38].

Electrical stimulation

Electrical stimulation was provided through one electrode or two or three electrodes simultaneously while recording RGC activity from all 512 electrodes. Stimulating with a given electrode consisted of passing a charge-balanced, triphasic current pulse through that electrode. The cathodic-peak stimulus consisted of a triphasic pulse with anodal/cathodal/anodal phases with relative current amplitudes 2:-3:1 and duration of 50 μs per phase (150 μs total). The anodic-peak stimulus was similar but with the polarity of the phases reversed to be cathodal/anodal/cathodal. All stimuli were delivered in 20 repeated trials at combinations of currents from 20 linearly spaced current amplitudes between -1.8 μA and 1.8 μA. As a result, for three-electrode stimulation, 8000 unique current combinations were tested. The ordering of the stimulation patterns was chosen pseudo-randomly, restricted so that each

stimulating electrode or group of electrodes was far from the previous stimulating electrodes to avoid stimulating the same cell(s) in rapid succession. A time delay of 10 ms was enforced between successive stimulations to ensure that the neural relative refractory period had elapsed and that successive trials could be considered independent. The stimulating electrodes were chosen based on the highest SNR of recorded spikes from the target cell as well as the geometric positioning of the electrodes relative to the target cell.

Responses to electrical stimulation

The spikes recorded during electrical stimulation were analyzed using a custom template matching approach [12]. First, the electrical image (EI) of each cell was calculated from visual stimulation data to serve as a template for the spike waveform of each cell to be detected during electrical stimulation. In the electrical stimulation data, an automated algorithm separated spikes from the electrical artifact by clustering traces obtained from repeated trials of the same electrical stimulus and comparing the waveform difference between these clusters to the EI templates obtained from visual stimulation.

Response probabilities for each cell were calculated for all current combinations passed through the three electrodes. For plotting experimental stimulation responses, only combinations causing response probability $0<p<1$ were used to more clearly display the shape of the response in the three-dimensional space of currents.

Biophysical model

The biophysical model of a spiking parasol RGC was developed in the NEURON [26] simulation environment. Full details are provided in [23]. Briefly, the biophysical properties of the RGC cable model used were adapted from [30] with modifications to the temperature of the simulation environment and the membrane dynamics as detailed in [39]. Additionally, the narrow axon compartment was removed, and the remaining compartment diameters were increased, which allowed for more stable model behavior when simulating electrical stimulation with multiple electrodes. All code used to build the RGC model is online and available at: https://github.com/ramanvilkhu/rgc_simulation_multielectrode

Multi-site model fitting

The multi-site model is described by the predicted probability given in Equation 1. The model weights were fitted using maximum likelihood estimation (MLE) considering the cell spiking responses for each current vector as binary outputs Y_j drawn i.i.d. from a Bernoulli distribution with parameter $p = p(x_j)$. The negative log-likelihood for this Bernoulli distribution with N observations (every independent trial at every current level tested) is given in Equation 3:

$$\textbf{(3)} \quad -l = -\sum_{j=1}^{N} \left[Y_j log(p(x_j)) + (1-Y_j) log(1-p(x_j))\right].$$

Equation 3 can be equivalently written in terms of a weighted sum with the weights being the number of repeated trials at every current level, reducing computation time.

$$\textbf{(3a)} \quad -l = -\sum_{j=1}^{J} T_j\left[P_j log(p(x_j)) + (1-P_j) log(1-p(x_j))\right].$$

Here, P_j is the empirically measured activation probability at current level x_j and T_j is the number of repeated trials performed at x_j to calculate P_j. This negative log-likelihood can be shown to be non-convex for the form of $p(x)$ given in Equation 1. In practice, solutions within parameter constraint bounds were found using second-order methods and regularization. The L-BFGS-B optimizer [40] was used for constrained minimization of Equation 3a with constraints on the bias parameters $w_{i,0}$ to enforce that the model fit produced a sufficiently low probability (0.01) with x = 0 and constraints such that the weight parameters $w_{i,k} \leq 100$. An L2 regularization penalty (λ = 0.5) was also added during optimization.

The process of fitting the multi-site model also involved choosing a hyperparameter: the number of sites (m in Equation 1), which in turn controlled the allowed degree of nonlinearity of fitted surfaces. The value m = 1 corresponded to logistic regression. In practice, a linear model with m = 1 reduces to a null model with only an intercept and no weightings on the currents when fitted to double-planar and nonlinear data. A fixed value of m = 8 was used to account for significant nonlinearities in the three-electrode data. For fitting to data from biophysical simulations of four-electrode stimulation, m = 10 was used. The goodness of fit was quantified using the McFadden pseudo-R^2 metric [41], one of several measures of goodness of fit commonly used for binary regression tasks. The definition of the McFadden pseudo-R^2 is given in Equation 4:

$$\textbf{(4)} \quad R^2_{McF} = 1 - \frac{l_m}{l_0}.$$

Here, l_m is the log-likelihood of the fitted model with number of sites m and l_0 is the log-likelihood of the null model which features only an intercept and no weighting on the input currents x. All code used for fitting the multi-site model can be found at:
https://github.com/pvasired/Vasireddy2025

Gaussian Process Regression and Level Set Estimation

Gaussian process models are nonparametric models describing a probability distribution over possible functions (see [27]). Gaussian process regression (GPR) is an established method for using Gaussian process models to fit a set of observations and predict outputs on novel inputs (see [28]).

In this application, the inputs are the multi-electrode stimulation amplitude vectors x_A using stimulus A. Assuming there are K cells $c_1,...,c_K$ stimulated by stimulus A at different activation probabilities $p_1,...,p_K$, an aggregate selectivity index can be constructed as in Equation 5:

$$\textbf{(5)} \quad S_A = max_i\, p_{i,A} \times (1 - max_{j \neq argmax_k\, p_{k,A}} p_{j,A}).$$

This aggregate metric yields the highest selectivity achievable for any cell using stimulus A, agnostic to which cell is selectively stimulated. This index ranges from 0 to 1, with 1 representing perfect selectivity. Rather than performing GPR directly on the selectivity values from Equation 5 that are bounded between 0 and 1, which would introduce additional complexity to modeling, a logit transformation of the aggregate selectivity index was performed to produce an unbounded real number, $y_A = log\left(\frac{S_A}{1-S_A}\right)$. GPR was then performed on the collection of pairs $\{(x_A, y_A)\}$ across stimulus amplitudes A. To predict activation probabilities of individual neurons

(as in Fig. 4D-F), the logit-transformed activation probabilities for each cell were used, $y_{i,A} = log\left(\frac{p_{i,A}}{1-p_{i,A}}\right)$.

Standard (i.e., non-sparse) GPR was implemented in the GPyTorch framework [42]. A constant mean function prior was assumed, and the RBF kernel was used for the covariance. Hyperparameter optimization of the RBF kernel length scale and the noise variance was performed using the AdamW optimizer [43]. The GPR predictions remained in logit (log-odds) units, but they could be transformed back to the bounded selectivity index using a sigmoid transformation.

With the fitted model from GPR on subsampled data, predictions were made across the entire input space of both the posterior mean function μ and variance σ. Using these predictions, the space of possible inputs was efficiently explored to accurately estimate the high selectivity regions. A variant of the GP-UCB algorithm [44] with similarities to methods from level set estimation (LSE; see [45]) was implemented to encourage exploration of regions that could yield high selectivity while preventing oversampling of regions known with high certainty to yield high selectivity. In the conventional GP-UCB algorithm, the next measurements to take are decided by choosing inputs that maximize the *upper confidence bound* $UCB = \mu + \beta\sigma$ of the current GPR fit, with β a hyperparameter controlling how wide of a confidence bound is considered. Larger values of β in GP-UCB will encourage exploration of unknown regions by strongly weighting uncertainty in the GPR predictions, while smaller values of β will encourage concentrated sampling of regions known to yield high posterior mean function values (see [44] for a fuller discussion of choosing β). In this implementation, the *lower confidence bound* $LCB = \mu - \beta\sigma$ is also used to exclude potential regions to sample. Specifically, during the fitting process, if $LCB > L$ for a given unmeasured input, that input is removed from the set of possible inputs to sample next. Of the remaining set of allowed inputs to sample, the next batch of measurements is chosen randomly amongst inputs which have $UCB > L$. In this implementation, $\beta = 2$, $L = 3$, and a batch size of 1000 measurements per iteration were used. Intuitively, this algorithm prevents sampling points which are nearly certain to yield high selectivity ($LCB > L$) and do not need further exploration while also sampling unexplored regions which could yield high selectivity ($UCB > L$). All code used here for GPR and LSE can be found at: https://github.com/pvasired/Vasireddy2025

Acknowledgements: National Science Foundation Grant No. 1828993 (PKV), ALS Association Milton Safenowitz fellowship (NPS), Excellence Initiative – Research University for the AGH University (PH), Research to Prevent Blindness Stein Innovation Award, Wu Tsai Neurosciences Institute Big Ideas, NIH NEI R01-EY021271, and NIH NEI P30-EY019005 (EJC). We thank Sasi Madugula, Scott Linderman, and the Stanford Artificial Retina team for experimental contributions and useful discussions.

Author Contributions: P.K.V. and E.J.C. conceived of the multi-electrode stimulation experiments and the approach for studying single-neuron responses to current steering. P.K.V., R.S.V., and A.R.G. designed the multi-electrode stimulation experiments and collected the

experimental data. P.K.V., A.L., and E.J.K. performed the downstream analysis and empirical modeling. P.K.V. and J.B.B. conceived of the specific mathematical framing of the multi-site model. A.L. and J.B.B. designed algorithms for spike sorting. R.S.V. implemented the biophysical simulations. A.J.P. assisted with experimental data collection. P.H., A.S., and A.M.L. designed and constructed the multi-electrode array system for stimulation and recording. P.K.V. prepared figures with input from R.S.V., A.J.P., P.H., E.J.K., N.P.S., and E.J.C. P.K.V., N.P.S., and E.J.C. wrote the manuscript with inputs from all authors.

Competing Interests: The authors declare no competing interests.

References

1. Krauss JK, Lipsman N, Aziz T, Boutet A, Brown P, Chang JW, et al. Technology of deep brain stimulation: current status and future directions. Nature Reviews Neurology. 2020;17: 75–87.

2. Zeng F-G, Rebscher S, Harrison W, Sun X, Feng H. Cochlear implants: system design, integration, and evaluation. IEEE Rev Biomed Eng. 2008;1: 115–142.

3. Deng Y, Qiao L, Du M, Qu C, Wan L, Li J, et al. Age-related macular degeneration: Epidemiology, genetics, pathophysiology, diagnosis, and targeted therapy. Genes Dis. 2022;9: 62–79.

4. Cross N, van Steen C, Zegaoui Y, Satherley A, Angelillo L. Retinitis Pigmentosa: Burden of Disease and Current Unmet Needs. Clin Ophthalmol. 2022;16: 1993–2010.

5. Werner JS, Chalupa LM. The New Visual Neurosciences. MIT Press; 2013.

6. Rodieck RW, Watanabe M. Survey of the morphology of macaque retinal ganglion cells that project to the pretectum, superior colliculus, and parvicellular laminae of the lateral geniculate nucleus. J Comp Neurol. 1993;338: 289–303.

7. Meister M, Berry MJ 2nd. The neural code of the retina. Neuron. 1999;22: 435–450.

8. Luo YH-L, da Cruz L. The Argus(®) II Retinal Prosthesis System. Prog Retin Eye Res. 2016;50: 89–107.

9. Hottowy P, Skoczeń A, Gunning DE, Kachiguine S, Mathieson K, Sher A, et al. Properties and application of a multichannel integrated circuit for low-artifact, patterned electrical stimulation of neural tissue. J Neural Eng. 2012;9: 066005.

10. Hottowy P, Dąbrowski W, Skoczeń A, Wiącek P. An integrated multichannel waveform generator for large-scale spatio-temporal stimulation of neural tissue. Analog Integrated Circuits and Signal Processing. 2007;55: 239–248.

11. Grosberg LE, Ganesan K, Goetz GA, Madugula SS, Bhaskhar N, Fan V, et al. Activation of ganglion cells and axon bundles using epiretinal electrical stimulation. J Neurophysiol. 2017;118: 1457–1471.

12. Gogliettino AR, Madugula SS, Grosberg LE, Vilkhu RS, Brown J, Nguyen H, et al. High-Fidelity Reproduction of Visual Signals by Electrical Stimulation in the Central Primate Retina. J Neurosci. 2023;43: 4625–4641.

13. Sekirnjak C, Hottowy P, Sher A, Dabrowski W, Litke AM, Chichilnisky EJ. High-resolution electrical stimulation of primate retina for epiretinal implant design. J Neurosci. 2008;28: 4446–4456.

14. Jepson LH, Hottowy P, Mathieson K, Gunning DE, Dabrowski W, Litke AM, et al. Focal electrical stimulation of major ganglion cell types in the primate retina for the design of visual prostheses. J

Neurosci. 2013;33: 7194–7205.

15. Dragas J, Viswam V, Shadmani A, Chen Y, Bounik R, Stettler A, et al. A Multi-Functional Microelectrode Array Featuring 59760 Electrodes, 2048 Electrophysiology Channels, Stimulation, Impedance Measurement and Neurotransmitter Detection Channels. IEEE journal of solid-state circuits. 2017;52. doi:10.1109/JSSC.2017.2686580

16. Vilkhu RS, Madugula SS, Grosberg LE, Gogliettino AR, Hottowy P, Dabrowski W, et al. Spatially patterned bi-electrode epiretinal stimulation for axon avoidance at cellular resolution. J Neural Eng. 2021;18. doi:10.1088/1741-2552/ac3450

17. Jepson LH, Hottowy P, Mathieson K, Gunning DE, Dąbrowski W, Litke AM, et al. Spatially patterned electrical stimulation to enhance resolution of retinal prostheses. J Neurosci. 2014;34: 4871–4881.

18. Firszt JB, Koch DB, Downing M, Litvak L. Current steering creates additional pitch percepts in adult cochlear implant recipients. Otol Neurotol. 2007;28: 629–636.

19. Butson CR, McIntyre CC. Current steering to control the volume of tissue activated during deep brain stimulation. Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation. 2008;1: 7–15.

20. Janson AP, Baker JL, Sani I, Purpura KP, Schiff ND, Butson CR. Selective activation of central thalamic fiber pathway facilitates behavioral performance in healthy non-human primates. Sci Rep. 2021;11: 23054.

21. Chaturvedi A, Foutz TJ, McIntyre CC. Current steering to activate targeted neural pathways during deep brain stimulation of the subthalamic region. Brain Stimul. 2012;5: 369–377.

22. Fan VH, Grosberg LE, Madugula SS, Hottowy P, Dabrowski W, Sher A, et al. Epiretinal stimulation with local returns enhances selectivity at cellular resolution. J Neural Eng. 2019;16: 025001.

23. Vilkhu R, Vasireddy P, Kish KE, Gogliettino AR, Lotlikar A, Hottowy P, et al. Understanding responses to multi-electrode epiretinal stimulation using a biophysical model. J Neural Eng. 2024. doi:10.1088/1741-2552/ada1fe

24. Maturana MI, Apollo NV, Hadjinicolaou AE, Garrett DJ, Cloherty SL, Kameneva T, et al. A Simple and Accurate Model to Predict Responses to Multi-electrode Stimulation in the Retina. PLoS Comput Biol. 2016;12: e1004849.

25. Rattay F, Tanzer T. A simple model considering spiking probability during extracellular axon stimulation. PLoS One. 2022;17: e0264735.

26. Carnevale NT, Hines ML. The NEURON Book. Cambridge University Press; 2006.

27. Rasmussen CE, Williams CKI. Gaussian Processes for Machine Learning. MIT Press; 2005.

28. Williams C, Rasmussen C. Gaussian Processes for Regression. Advances in Neural Information Processing Systems. 1995;8. Available: https://proceedings.neurips.cc/paper_files/paper/1995/file/7cce53cf90577442771720a370c3c723-Paper.pdf

29. Ayton LN, Barnes N, Dagnelie G, Fujikado T, Goetz G, Hornig R, et al. An update on retinal prostheses. Clin Neurophysiol. 2020;131: 1383–1398.

30. Kish KE, Lempka SF, Weiland JD. Modeling extracellular stimulation of retinal ganglion cells: theoretical and practical aspects. J Neural Eng. 2023;20. doi:10.1088/1741-2552/acbf79

31. Rattay F. The basic mechanism for the electrical stimulation of the nervous system. Neuroscience.

1999;89: 335–346.

32. Werginz P, Fried SI, Rattay F. Influence of the sodium channel band on retinal ganglion cell excitation during electric stimulation – A modeling study. Neuroscience. 2014;266: 162–177.

33. Esler TB, Maturana MI, Kerr RR, Grayden DB, Burkitt AN, Meffin H. Biophysical basis of the linear electrical receptive fields of retinal ganglion cells. J Neural Eng. 2018;15: 055001.

34. Pachitariu M, Steinmetz N, Kadir S, Carandini M, Harris K. Fast and accurate spike sorting of high-channel count probes with KiloSort. NIPS Proceedings. Neural Information Systems Foundation, Inc.; 2016. Available: http://uhra.herts.ac.uk/handle/2299/19969

35. Chichilnisky EJ. A simple white noise analysis of neuronal light responses. Network: Computation in Neural Systems. 2001. pp. 199–213. doi:10.1080/713663221

36. Chichilnisky EJ, Kalmar RS. Functional asymmetries in ON and OFF ganglion cells of primate retina. J Neurosci. 2002;22: 2737–2747.

37. Field GD, Sher A, Gauthier JL, Greschner M, Shlens J, Litke AM, et al. Spatial properties and functional organization of small bistratified ganglion cells in primate retina. J Neurosci. 2007;27: 13261–13272.

38. Litke AM, Bezayiff N, Chichilnisky EJ, Cunningham W, Dabrowski W, Grillo AA, et al. What does the eye tell the brain?: Development of a system for the large-scale recording of retinal output activity. IEEE Trans Nucl Sci. 2004;51: 1434–1440.

39. Fohlmeister JF, Cohen ED, Newman EA. Mechanisms and Distribution of Ion Channels in Retinal Ganglion Cells: Using Temperature as an Independent Variable. J Neurophysiol. 2010 [cited 17 Oct 2023]. doi:10.1152/jn.00123.2009

40. Byrd RH, Peihuang L, Nocedal J. A limited-memory algorithm for bound-constrained optimization. Argonne National Laboratory; 1996 Mar. Report No.: MCS--P404-1293. doi:10.2172/204262

41. McFadden D. Conditional Logit Analysis of Qualitative Choice Behavior. 1973.

42. Gardner J, Pleiss G, Weinberger KQ, Bindel D, Wilson AG. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. Advances in Neural Information Processing Systems. 2018;31. Available: https://proceedings.neurips.cc/paper_files/paper/2018/file/27e8e17134dd7083b050476733207ea1-Paper.pdf

43. Loshchilov I, Hutter F. Decoupled Weight Decay Regularization. 2017. Available: http://arxiv.org/abs/1711.05101

44. Srinivas N, Krause A, Kakade SM, Seeger M. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. 2009 [cited 30 Jan 2025]. doi:10.1109/TIT.2011.2182033

45. Gotovos A, Casati N, Hitz G, Krause A. Active Learning for Level Set Estimation. International Joint Conference on Artificial Intelligence. 2013 [cited 20 Jan 2025]. doi:10.3929/ETHZ-A-009767767

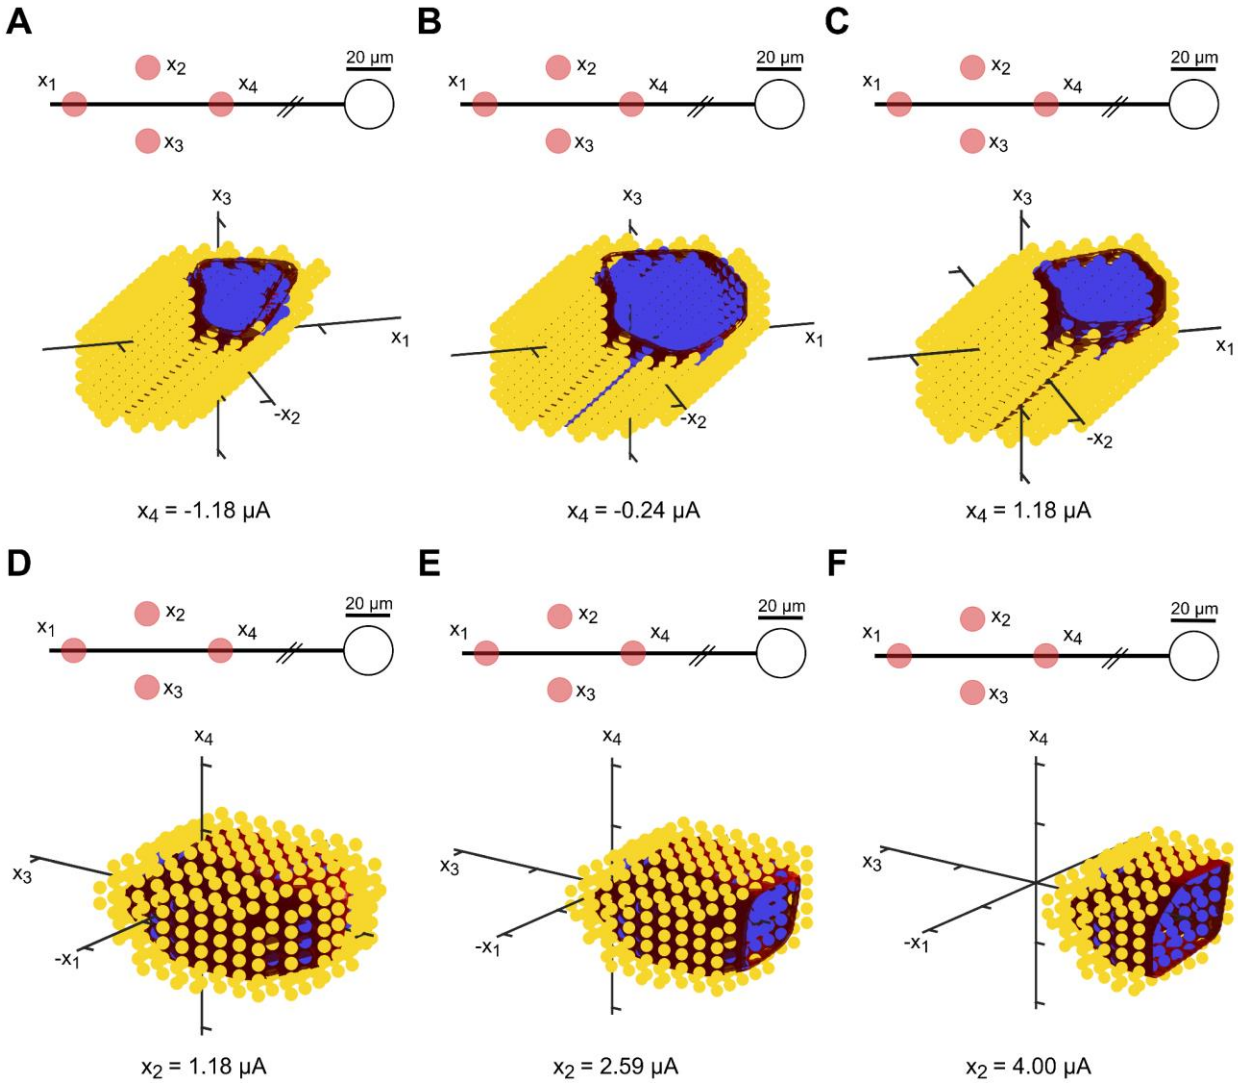

Extended Data Figure 1. *Extension of multi-site model fitting to four-electrode stimulation responses in biophysical simulations.* Nonlinear four-electrode stimulation biophysical simulation data with multi-site model 50% surface for m = 10 activation sites. The multi-site model fit had a McFadden pseudo-R^2 of 0.86, indicating an excellent fit (see Methods). Slices with freely varying amplitudes on three electrodes and fixed fourth electrode amplitude are shown. ***A-C,*** Stimulation responses along three-electrode slices with fixed amplitude fourth electrode along the axon, for various fixed amplitude levels on the fourth electrode. ***D-F,*** Stimulation responses along three-electrode slices with fixed amplitude fourth electrode transversely offset from the axon, for various fixed amplitude levels on the fourth electrode.

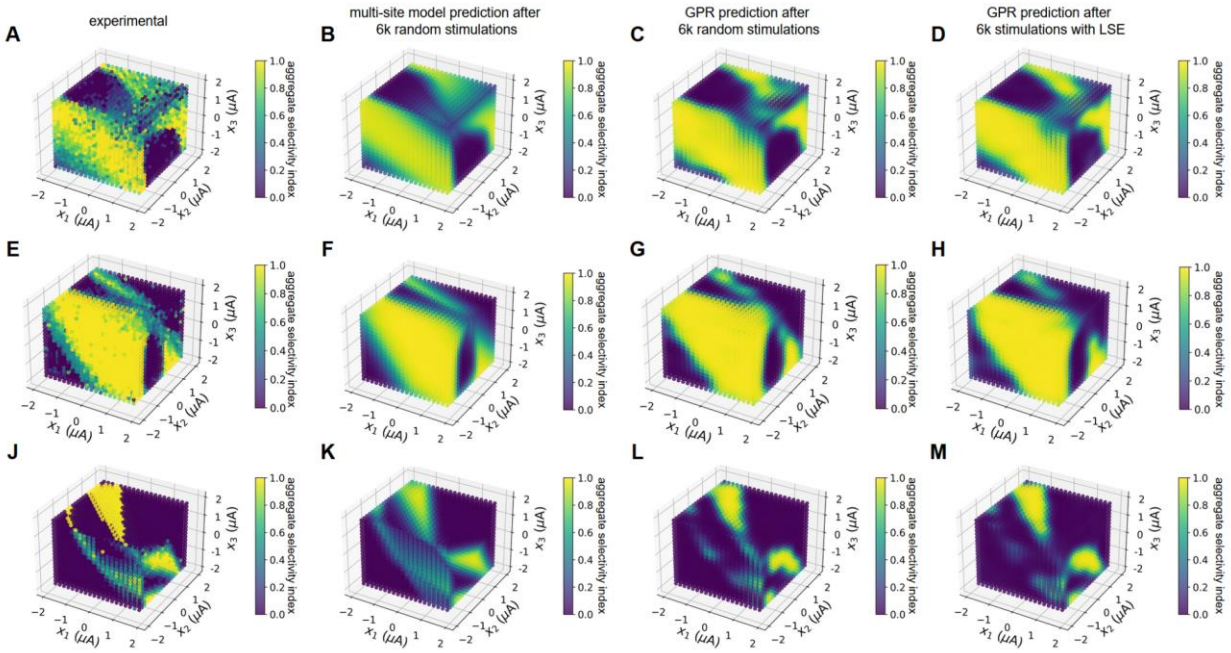

Extended Data Figure 2. *Gaussian process regression (GPR) and multi-site model for estimation of aggregate selectivity profiles.* First column shows experimentally measured aggregate selectivity profiles from full data (160,000 stimulations), second column shows multi-site model prediction after fitting to subsampled data from each cell and computing the aggregate selectivity index, third column shows GPR mean prediction after performing 6,000 random calibrating stimulations, and fourth column shows GPR mean prediction after performing 6,000 calibrating stimulations guided by the GPR+LSE sampling procedure. Each row represents a distinct electrode triplet from a particular retinal preparation. ***A-D,*** Macaque peripheral RGCs (2 cells aggregated). ***E-H,*** Macaque central RGCs (5 cells aggregated). ***J-M,*** Macaque peripheral RGCs (4 cells aggregated).