


SCIENTIFIC REPORTS



OPEN

Evolution of extrema features reveals optimal stimuli for biological state transitions

Joshua Chang ^{1,2} & David Paydarfar^{2,3}

The ability to define the unique features of an input stimulus needed to control switch-like behavior in biological systems is an important problem in computational biology and medicine. We show in this study how highly complex and intractable optimization problems can be simplified by restricting the search to the signal's extrema as key feature points, and evolving the extrema features towards optimal solutions that closely match solutions derived from gradient-based methods. Our results suggest a model-independent approach for solving a class of optimization problems related to controlling switch-like state transitions.

Switch-like transitions play an important regulatory role in biological systems. Cell division, differentiation and apoptosis are controlled by complex molecular signaling networks that can abruptly switch from one pathway to another^{1–4}. Electrical signaling of multicellular networks switch between distinct electrical states – quiescence, repetitive firing, and bursting patterns – that are essential for controlling normal movement^{5,6}, autonomic function⁷, and behavioral states⁸ or perhaps triggering pathological states like migraines^{9,10}. Therapeutic interventions work to exploit switch-like shifts from one state to another, for example, inducing immunity through vaccination with small doses of sub-pathologic viruses¹¹, or applying electrical stimuli to defibrillate the heart^{12–14}, suppress parkinsonian tremors^{15–17}, and abort epileptic seizures^{18,19}. Normal and pathological switching are typically governed by highly nonlinear mechanisms, with small input signals inducing large and rapid system responses.

This idea that tiny perturbations can induce large state changes has been the focus of a growing interest in critical transitions^{10,20–23}. Some have examined these critical transitions in the context of moving along parameter continua (e.g. transitioning from a monostable system to a bistable system)²², while others have examined transitions in a single system from one state to another state²⁰. When examining transitions from one state to another, there are a number of different forms, as seen in Fig. 1. Some perturbations induce a specific response trajectory (e.g. an action potential), while others cause a transition from one stable state to another stable state (e.g. limit cycle or fixed point). In all of these cases, an abrupt transition is occurring in the state space caused by a small input stimulus.

A major challenge to understanding these transitions is recognizing what are the key stimulus waveform features that trigger them. While many different stimulus waveforms might push a system across a tipping point to cause a transition in state, they are not all optimal. Defining optimality is highly specific to the system of interest. For electrical stimulation of tissue, clinicians may wish to minimize the peak stimulus intensity to avoid causing tissue damage or energy leakage into neighboring regions, producing detrimental side effects. Other applications may instead seek to minimize total energy consumption, to maximize battery life for implantable devices. In vaccines or chemotherapy, one may seek to minimize the total injected dose to prevent toxicity. In general, therapeutic optimality seeks to achieve a beneficial health outcome while minimizing adverse effects.

Analytical methods for deriving optimal control signals traditionally use calculus of variations^{24–26}; broad application of these techniques to biology is limited because validated mathematical models are available in only a few systems, and variational solutions to complex models can be elusive^{27,28}. To deal with complex models, researchers have used phase reduction models^{28–32}, parameterized stimuli³³, or simpler conceptual models^{24,25}. The majority of biological systems however, do not have models readily available, and as such empirical approaches have explored the system's responses to simple stimulus waveforms (e.g. rectangular, sinusoidal)^{34–37}.

¹Department of Neurology, University of Massachusetts Medical School, Worcester, Massachusetts, 01604, USA.

²Department of Neurology, Dell Medical School, The University of Texas at Austin, Austin, Texas, 78701, USA. ³The Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, Texas, 78701, USA. Correspondence and requests for materials should be addressed to J.C. (email: joshua.chang@austin.utexas.edu) or D.P. (email: david.paydarfar@austin.utexas.edu)

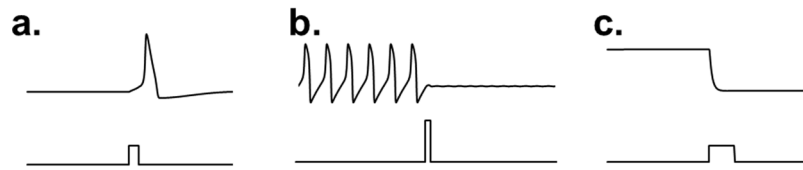


Figure 1. Models illustrating different types of transitions. Panel a shows the Hodgkin-Huxley model transitioning between a stable quiescent state and an unstable active state (action potential), induced by a brief depolarizing pulse. Panel b shows the FitzHugh-Nagumo model transitioning between a stable oscillatory state (repetitive firing) and a quiescent state (with damped subthreshold oscillations), induced by a brief perturbation. Panel c shows the genetic toggle switch model transitioning between two stable steady states of protein production, induced by a pulse of isopropyl β -D-1-thiogalactopyranoside (IPTG). For an explanation of the models, see Supplementary Text S1.

The stimulus parameters (e.g., pulse width, frequency) are tuned to achieve the optimal parameters for achieving an outcome. Unfortunately, this approach is also limited because it requires certain assumptions about the optimal shape, leading to a biased search, potentially missing more optimal solutions.

There has been interest in developing unbiased methods to search for optimal stimuli that do not require *a priori* knowledge of the underlying mechanisms or mathematical equations defining the system^{38–40}. Since biological control systems are often highly nonlinear and respond to fluctuations over broad ranges of intensity, duration, and timing, finding the optimal control signal can be challenging. One approach is to encode the stimulus as a vector, representing stimulus amplitude over time³⁹. The number of dimensions a stimulus requires is calculated by multiplying the duration of the stimulus by the resolution of the stimulus generator. For example, the solution space for a 1-second stimulus with 1-dimensional amplitude generated at a resolution of 1-millisecond has a dimension of 1000. Such high dimensionality of the solution space can rapidly exceed the limits of computational tractability, an example of the “curse of dimensionality”^{41,42}.

Here we present a novel evolutionary method for estimating optimal stimuli based on the hypothesis that the optimization problem is vastly simplified by restricting the search algorithm to the signal extrema as key feature points. The algorithm is initially seeded by random signals whose extrema are iteratively perturbed and pruned, thereby sculpting the stimulus waveform towards an optimal shape. We find that this process allows for many degrees of freedom in the early stages, but as the algorithm progresses and extraneous extrema are removed, the degrees of freedom decrease, increasing the probability of finding more optimal solutions. Multiple, independent stochastically-generated initial seeds enable the algorithm to search for both local and global optima. By comparing to numerical solutions from analytically-derived optimal approaches, we demonstrate and validate the algorithm in three models representing different types of transitions (Fig. 1): from quiescence to a single action potential in the Hodgkin-Huxley model⁴³; between sustained oscillation and quiescence in the FitzHugh-Nagumo model^{44,45}; and between two quiescent states in a genetic toggle switch model⁴⁶. We also apply the algorithm to a more complex problem of finding an optimal stimulus that desynchronizes a group of coupled Hodgkin-Huxley neuronal oscillators, which is not readily solved by traditional approaches.

Methods

Stochastic searching with extrema features. This algorithm is based on the hypothesis that optimization of control signals can be aided by restricting the search space to the signal extrema as key feature points, rather than allowing every point of the stimulus to be a unique dimension. The use of extrema in simplifying data has been introduced in other problems in signal processing. In the field of pattern recognition, certain algorithms use the coordinates of robust extrema to decrease the number of points to match⁴⁷. In computer vision, extrema are used as key feature markers for object recognition and multi-scale representation of the original images^{48,49}. In many of these algorithms, the extrema are used to highlight key feature points, and thus construct a simplified representation of the original signal retaining only the most important components. This makes the manipulation and usage of the signals more efficient and simpler. We hypothesize that this concept of using extrema to simplify signals is applicable to stochastic search algorithms.

The steps of our algorithm are as follows:

1. Let μ_0 be our initial stimulus, where $\mu = (\mu[1], \mu[2], \dots, \mu[M])$, M being the product of the number of samples per unit time and the duration of the stimulus. This vector is initially randomly generated using a uniform random number generator at each time point.
2. Define the function, f , to represent the system in question. Thus, $f(\mu)$ is the response of the stimulus to a given stimulus.
3. Define T to be the set of times $\{T_1, T_2, \dots, T_N\}$ at which extrema occur in the stimulus. We define an extremum to be an element such that at time T_n :

$$\mu[T_n - 1] < \mu[T_n] \text{ and } \mu[T_n + 1] < \mu[T_n] \text{ or } \mu[T_n - 1] > \mu[T_n] \text{ and } \mu[T_n + 1] > \mu[T_n]. \quad (1)$$

4. Create a set of intervals, I , to be defined as $T_n - T_{n-1}$ for $n = 2, \dots, N$. Multiply each of these intervals by a random number generated from a Gaussian distribution:

$$I_{temp, n} = I_n * \mathcal{N}(1, \sigma_I). \quad (2)$$

- Rescale the new time intervals such that the total duration matches the original duration:

$$I_n' = I_{temp, n} * \Sigma I / \Sigma I_{temp}. \quad (3)$$

- Use these new intervals to redefine a new set of times, T' .
- Perturb the amplitude of each of the extrema by adding to each peak a randomly generated number:

$$E_n' = E_n * \mathcal{N}(0, \sigma_E) \quad (4)$$

- Linearly transform all the data points between the original extrema using the new perturbed extrema end-points to generate a “neighboring stimulus” of the starting seed. Mathematically, a point $(t, \mu(t))$ between the two old extrema $(T_n, \mu(T_n))$ and $(T_{n+1}, \mu(T_{n+1}))$, transforms to

$$\left(\left(\frac{T_{n+1}' - T_n'}{T_{n+1} - T_n} \right) * ((t - T_n) + T_n'), \left(\frac{\mu'(T_{n+1}') - \mu'(T_n')}{\mu(T_{n+1}) - \mu(T_n)} \right) * ((\mu(t) - \mu(T_n)) + \mu'(T_n')) \right) \quad (5)$$

where $(T_n', \mu'(T_n'))$ and $(T_{n+1}', \mu'(T_{n+1}'))$ are the new perturbed extrema.

- Apply this new stimulus to the system and measure its system response.
- Calculate the performance metric (e.g. L^2 -norm, $J = \int \mu^2 dt$) of the neighboring stimulus, penalizing the performance metric if the stimulus does not successfully cause a state transition. For example, we used L^2 -norm:

$$J = \int \mu^2 dt + \text{penalty} \quad (6)$$

- Repeat steps 3–7 10 times to produce 10 neighboring solutions based on the starting stimulus.
- Compare the original starting stimulus with the neighboring stimulus to find the solution with the best performance.
- Repeat steps 2–11 for a predetermined number of iterations, L , using the best performing stimulus found from the previous iteration as the new starting seed in step 2.
- After L iterations, output the stimulus with the best performance metric as the optimal stimulus.

Figure 2 illustrates one iteration of the signal perturbation carried out between Steps 3–6 in a search for the optimal stimulus waveform to trigger an action potential in the Hodgkin-Huxley model. For the uniform random number generator at Step 1, the uniform random number generator should be set such that the noise given to the system just barely causes a switch in states. If an excessive amount of noise is given, state switches may be more easily obtained, but there is a lot more energy to reduce in the iterative process. We used a Gaussian distribution to randomly shift the extrema amplitudes and inter-extrema intervals. The variance of this distribution, the number of neighboring stimuli generated for each iteration, and the number of iterations used before ending the algorithm were empirically tuned to optimize computational performance of the algorithm. If the variance of the distributions are very large, the neighboring stimuli become more spread out, which may lead to better exploration at the beginning of the process, but fails to converge as smoothly in the later iterations. Conversely, if the variance of the distributions are very small, the iterative perturbations from one solution to the next is much smaller allowing for smoother convergences, but in the early stages the algorithm takes a lot longer to converge.

In most cases, L^2 -norm was used as the optimization metric, for example to achieve an energy-optimal electrical stimulus waveform for causing a single action potential. For other applications described in this study we defined different optimization metrics (e.g., peak or total protein concentration for the genetic toggle switch), which can include multiple performance measures (e.g., combining stimulus L^2 -norm and a desynchronization index for optimizing a stimulus that maximally desynchronizes a group of coupled neurons).

Validation using a stochastically-seeded gradient algorithm. Calculus of variations⁵⁰ is a well-established analytical framework used to address optimization problems. We solved the variational problem directly using a stochastically-seeded gradient algorithm²⁷, applied to three models (defined in Supplementary Text S1) representing different types of transitions: from quiescence to a spike in the Hodgkin-Huxley model; between sustained oscillation and quiescence in the FitzHugh-Nagumo model; and between two steady states in a genetic toggle switch model. The stochastically-seeded gradient algorithm generates a series of random stimuli, and using a gradient-based approach, calculates how the system will respond to minute changes in the stimulus. The algorithm iteratively uses these small gradients, or slopes, to converge towards the optimal solution. Stochastic seeding again enables the algorithm to explore a large solution space, revealing both the global optimal waveform as well as local optima²⁷.

Results

Stochastic search using extrema features convergences quickly. As described in the Methods section, this algorithm began with a stochastically generated stimulus. We first applied the algorithm to the Hodgkin-Huxley model to find the optimal stimulus waveform necessary to trigger a single action potential. Figure 3 shows the progression of this signal as the algorithm distorted the intervals between and amplitudes of the extrema. It is interesting to note that the algorithm converges very quickly in first few iterations and, within

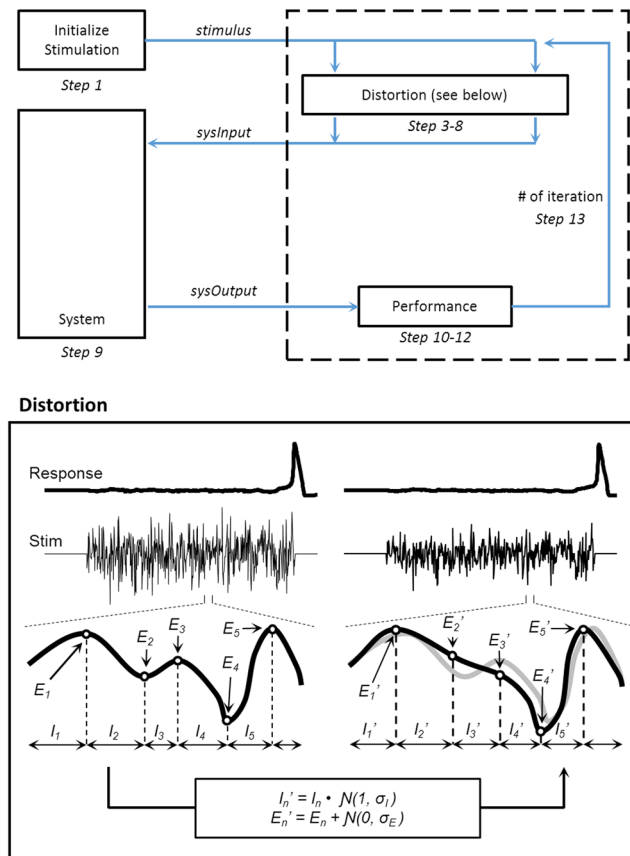


Figure 2. Schematic of the extrema featured stochastic search algorithm (top) along with a detailed representation of the distortion component (bottom). The distortion inset showcases both extrema interval and amplitude distortions before (bottom left) and after (bottom right) a single iteration. E_n and E_n' represent the n^{th} extrema of the stimulus before and after the iteration, respectively, and I_n and I_n' represent the n^{th} intervals between the respective extrema.

10–20 iterations, the fundamental shape of the optimal solution appears. The next few thousand iterations further prune extraneous extrema and smooths the signal to converge towards the optimal solution.

From a different perspective, one can observe the extrema, both in terms of their temporal locations and amplitudes as seen in Fig. 4. The algorithm removed extraneous extrema over the course of the algorithm.

Using extrema improves search efficiency. We reasoned that our focus on extrema would improve computational efficiency because it decreased the search space as we approached optimality. To evaluate this idea, we reconfigured our algorithm to use every single data point, i.e., stochastic perturbation of all-points instead of only the extrema points. As seen in Fig. 5A, the improvement of L^2 -norm across the 35 snippets using just the extrema converged towards an optimal solution faster than the improvement of L^2 -norm when perturbing all the points. The average number of iterations to drop below an L^2 -norm of 25 was 31.9 (S.D.: 5.5) using all-points versus 19.3 (S.D.: 5.2) using just extrema. To drop below an L^2 -norm of 20 took 51.6 iterations (S.D.: 15.6) using all-points versus 29.4 (S.D.: 7.8) using just extrema. To get below an L^2 -norm of 18 took 91.7 (S.D.: 34.5) iterations using all-points while it took only 36.8 iterations (S.D.: 9.4) using extrema only. Moreover, the average L^2 -norm when perturbing all points was 15.902 with a standard deviation of 0.171. The best result had an L^2 -norm of 15.683. Figure 5B shows that the best result in all-points perturbation retained a considerable level of noise as compared to the best result found when using just extrema.

Both results supported our hypothesis that using extrema features increases efficiency, while maintaining a high level of accuracy.

Validation of stochastic search using extrema features. We applied this technique to three unique models to display the adaptability of this algorithm to a range of different systems. We studied two computational models used in neuroscience: the Hodgkin-Huxley model for triggering a single action potential, and the FitzHugh-Nagumo model for suppressing repetitive firing. We have also applied this algorithm to a model used in synthetic biology, the genetic toggle switch, for causing a switch from one steady state to another. The methods and models used are described in greater detail in Supplementary Text S1.

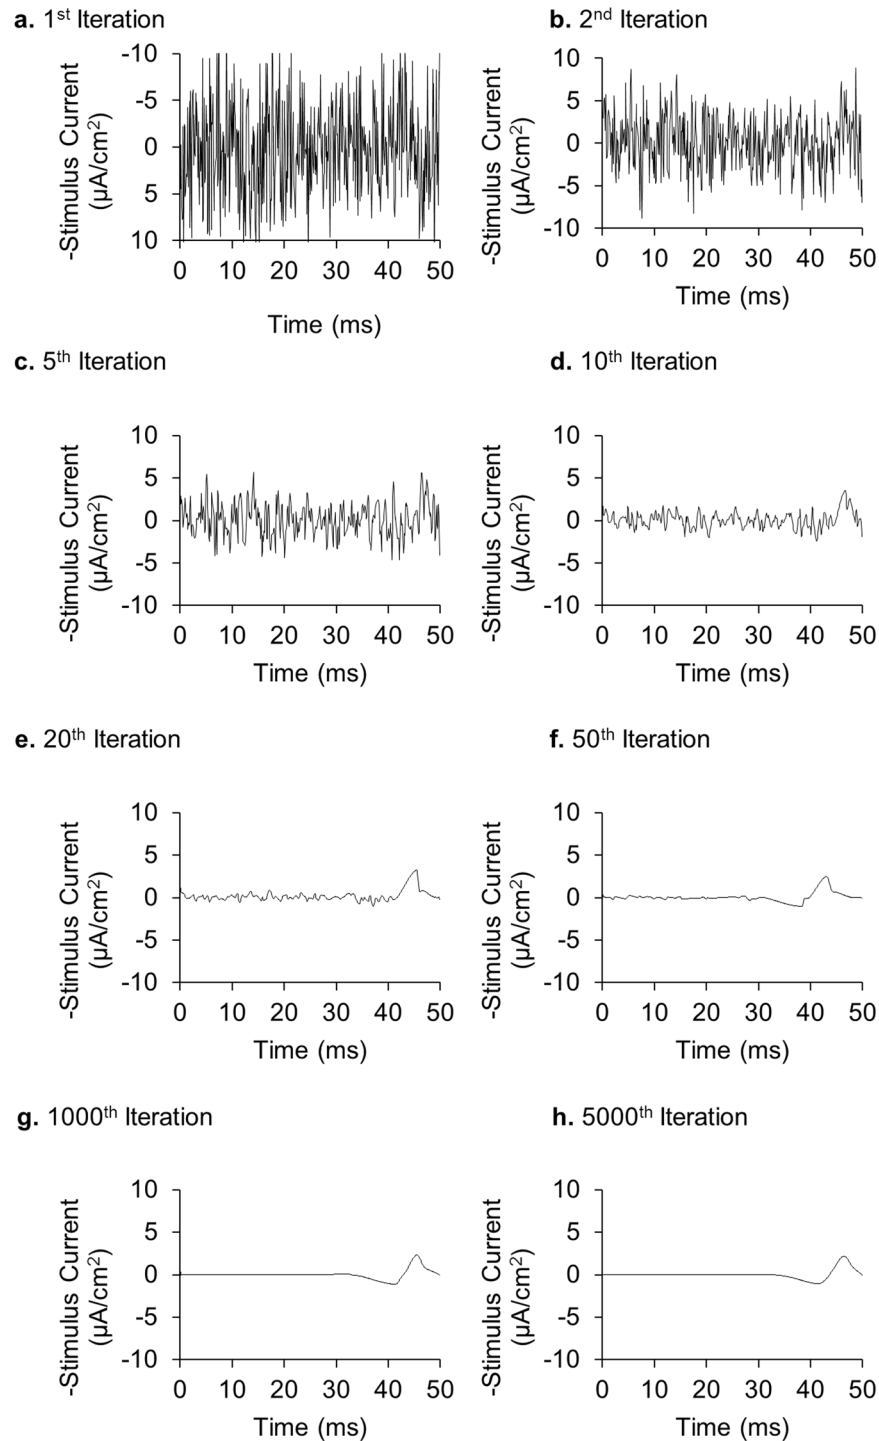


Figure 3. Convergence of the stimulus shape using an extrema feature stochastic search algorithm. We searched for the optimal 50-ms stimulus waveform necessary to trigger a single action potential. Waveforms are evaluated based on their L^2 -norm, or energy consumption. See Supplementary Text S1 for implementation details.

We applied the algorithm to each of these models, and compared the resulting waveform obtained using the stochastically seeded gradient algorithm. Figure 6 shows the best solution from multiple runs with different starting seeds.

For the Hodgkin-Huxley model, this algorithm found solutions that had an average L^2 -norm of 15.6 (standard deviation of 0.149) compared to gradient algorithm at 15.4. For the FitzHugh-Nagumo model, results from focusing on extrema features had a mean of 7.2×10^{-3} (standard deviation: 1.8×10^{-3}) compared to gradient algorithm's 1.12×10^{-2} . For the genetic toggle switch, the results had a mean of 2.97×10^{-8} (standard deviation: 2.25×10^{-9}) compared to gradient algorithm's 2.49×10^{-8} .

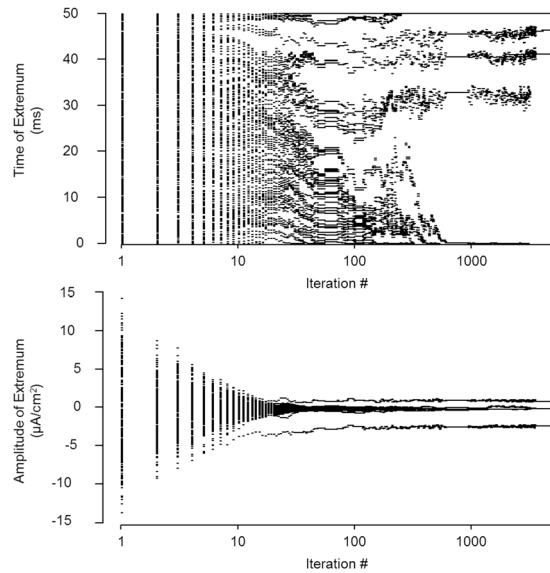


Figure 4. Tracking of the stimulus extrema. Every single point on both the top and the bottom panel represent a single extremum, plotted based either on when the extrema occurred in time (top) or in amplitude (bottom).

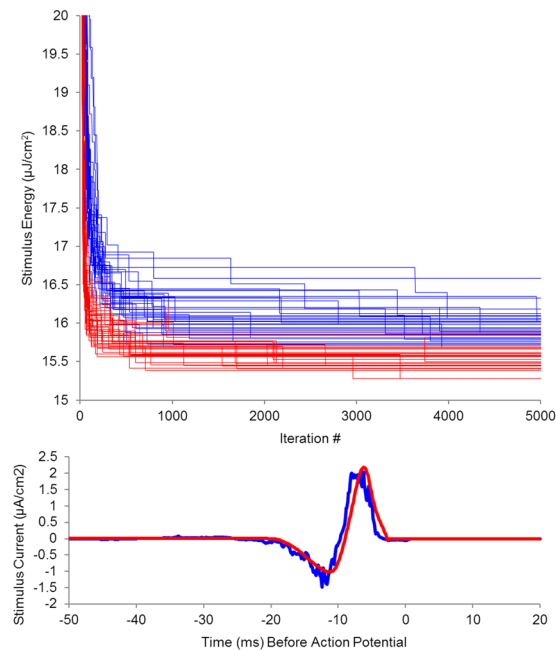


Figure 5. Comparison of an all-points search algorithm (blue) versus an extrema feature search algorithm (red). Top panel shows improvement in energy level (L^2 -norm) over the course of 500 iterations for 35 randomly seeded stimuli for each of the two algorithms. Bottom panel shows the stimulus for each algorithm having the least L^2 norm after 500 iterations.

In the FitzHugh-Nagumo and the Hodgkin-Huxley models, we noted that in certain instances, stochastic search did better than the gradient algorithm. Exploring the Hodgkin-Huxley model in greater depth, we noted that this technique only beat the gradient algorithm once in a set of 35 instances, having an L^2 -norm of 15.3. This is possible because the gradient algorithm is an approximation technique, utilizing first-order gradients to approach the optimal solution, and as such, depending on the topology of the system, it can have difficulty converging tightly to the optimal.

While this explained the discrepancy that occurred rarely with the Hodgkin-Huxley model, we were surprised that in the FitzHugh-Nagumo model, the stochastic search outperformed the gradient algorithm on a regular basis. As we investigated this further, we noticed that the extrema featured stochastic search found the optimal solution achieving our constraint, the suppression of repetitive firing, without necessarily having reached the mathematical stable point.

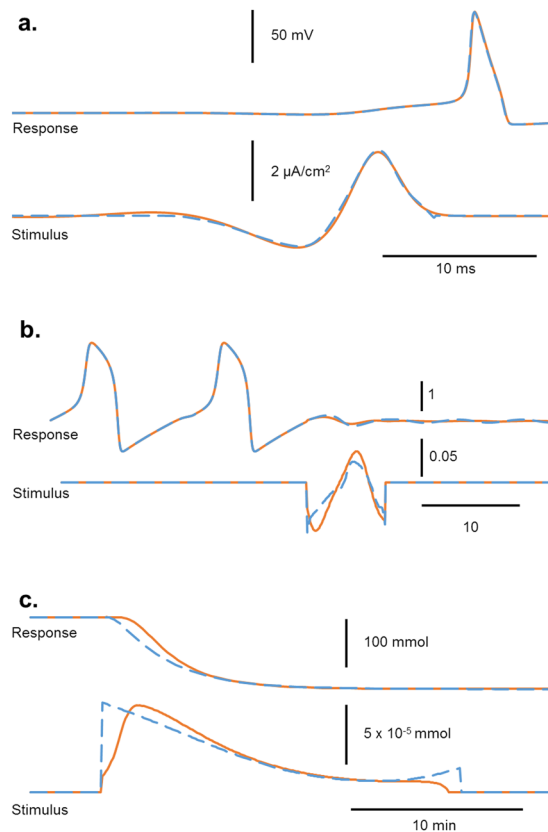


Figure 6. Comparison of optimal waveforms for eliciting a transition in three model. A single action potential in the Hodgkin-Huxley model (top), suppression of repetitive firing in the FitzHugh-Nagumo model (middle), and induction of a switch in protein synthesis in the genetic toggle switch model (bottom), using the gradient algorithm (dashed blue) and our extrema search algorithm (solid orange).

The FitzHugh-Nagumo model is interesting in that there is a stable limit cycle and a stable fixed point, with an unstable cycle in between the two states. The basin around the stable fixed point is very shallow, meaning that the system takes a long time within this basin to converge towards the fixed point. If the stimulus caused the system to cross the threshold from repetitive firing into quiescence, there will still be a period of oscillation depending on the state's proximity to the stable fixed point. The gradient algorithm calculated the optimal solution for the system to reach very close to the fixed point by the end of the simulation, while our algorithm sought to simply cross the threshold from repetitive firing to quiescence in the same time frame. We demonstrated the energy needed to cross the threshold is much smaller than the energy necessary to take the system completely to the fixed point.

Flexibility of stochastic search. Based on what we noticed in the FitzHugh-Nagumo model, we further investigated how changing the constraints affected the results. Could we tighten or loosen the constraints that define a successful suppression, and how would that affect the optimal solution? Fig. 7 shows the results of that experiment. We tightly suppressed the terminal oscillations in one experiment, with a resulting stimulus waveform that matched the gradient algorithm result much more closely. Similar, we also loosened the constraint such that only the bare minimum was required to suppress repetitive firing of action potentials. The result showed that a large energy difference could be achieved by this simple change in constraints. Using the gradient algorithm, we could not make these adjustments without defining the entire unstable limit cycle quantitatively and searched for optimal waveforms to every point along the unstable limit cycle. Here, our algorithm adjusted itself quickly to our changes in constraints.

Not only were we able to adjust our constraints, but we found that we could adjust our performance metrics to more real-world quantities in ways that are difficult to achieve mathematically. For instance, in a genetic toggle switch, the total concentration of molecular substrate used may not be as relevant as the peak dosage at any given point in time. It is difficult to write a mathematical formula defining peak dosage at any given point in time for the purposes of the gradient algorithm, but relatively straight forward with using our algorithm. Instead of using the L^2 -norm in our performance metric, we would use the maximal dosage amount in the stimulus as our performance metric. Figure 8 shows the differences between the two results using our algorithm: optimization for total concentration versus maximal dosage.

Application of stochastic search using extrema features to complex systems – Desynchronization.

We were able to also construct compounded requirements for the optimal stimulus waveform necessary to suppress synchronization of multiple repetitively firing neurons. This concept has been hypothesized to be important

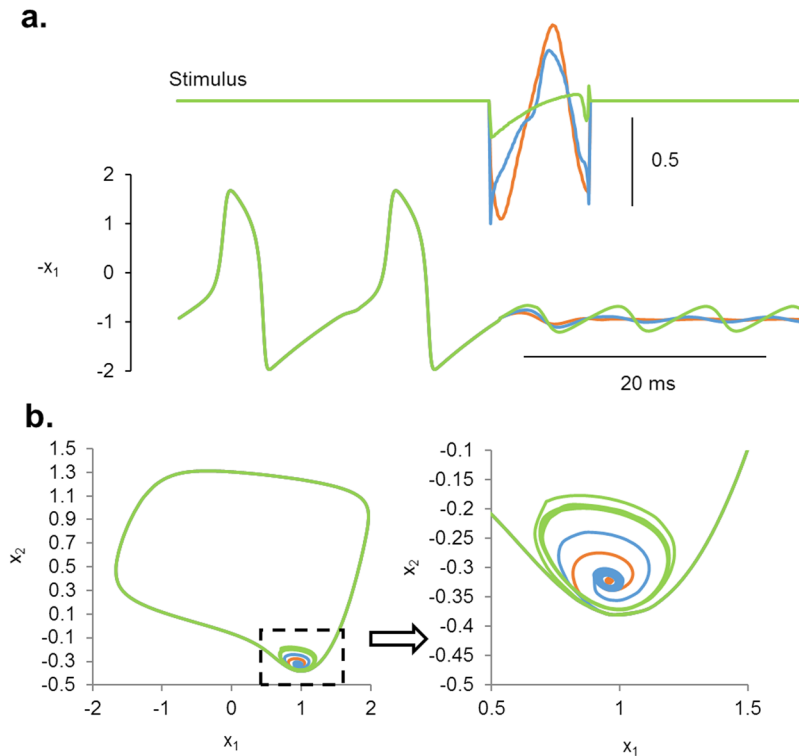


Figure 7. Searching for optimality using different levels of terminal condition constraints. Stricter constraints on the terminal condition in the FitzHugh-Nagumo model (blue) led to results that match more closely to the results of the gradient algorithm (gray) demonstrating a steadier quiescent state as compared to when looser constraints on the terminal condition were applied (orange).

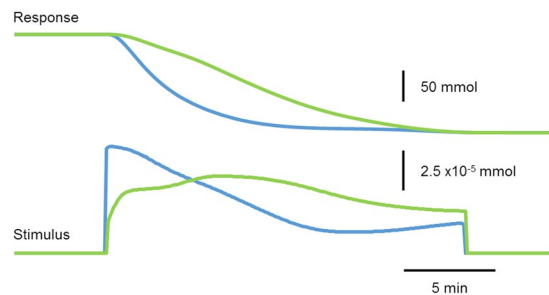


Figure 8. Optimal stimulus waveform is different depending on the performance metric defined. In the genetic toggle switch, minimizing peak dosage (red) revealed a different waveform than minimizing total concentration (blue).

from a clinical perspective in the suppression of seizure activity^{17,19,28,51}. Finding the optimal waveform for a complete ionic model of multiple coupled neurons is difficult; analytical approaches have attempted to solve simplified phase models^{19,28,29}, but solutions to ionic models have not been reported. Our algorithm approached this easily by constructing a compound performance metric that simultaneously (1) minimized stimulus energy and (2) maximized neuronal desynchronization. In the example shown, the algorithm was programmed to maximize neuronal desynchronization with a higher weight over the minimization of stimulus energy. Figure 9 shows the resulting optimal waveform pattern to desynchronize a group of five coupled Hodgkin-Huxley neurons. We used the mean field voltage as our marker for desynchronization, i.e., the lower the mean field voltage, the greater the neurons are desynchronized¹⁹. The model and methods can be found in Supplementary Text S1.

We noted that the optimal stimulus exhibits a rapid reversal of polarity just before the expected burst of action potentials from the synchronized neurons. We believe that this is because the stimulus is being given close to a point of discontinuity on the phase resetting curve. Thus, as seen in Fig. 10, when the stimulus hits the system at this point, the tiny phase differences in the initial conditions are suddenly magnified, and the phases are dispersed.

Previous work has shown that the phase response of nonlinear oscillators to discrete perturbing stimuli can show similar phase dispersion in response to precisely timed stimuli. These studies use topological classification

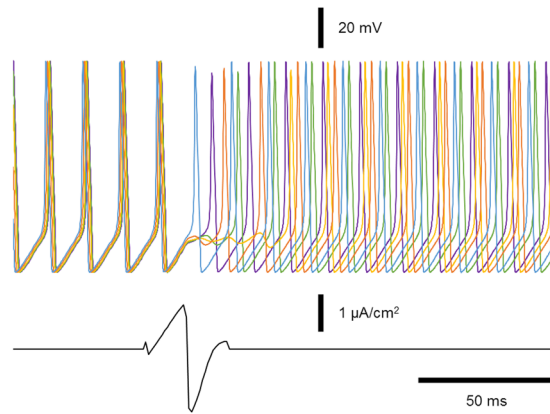


Figure 9. Desynchronization of coupled oscillator network. The stimulus shown here (bottom) is the optimal found after 1,000 iterations that maximized desynchronization seen in the coupled oscillators (top). The algorithm searched to minimize the mean voltage while also minimizing energy consumption. Positive axis for the stimulus reflect depolarization. While a numerical solution based on analytically-derived approaches for optimality for this model could not be calculated, we systematically tested a wide range of rectangular pulses with varying widths and amplitudes to find the optimal rectangular pulse. We found the optimal rectangular waveform had an L^2 -norm of $154.88 \mu\text{J}/\text{cm}^2$ as compared to $11.80 \mu\text{J}/\text{cm}^2$ for the optimal determined by our extrema-driven stochastic search algorithm.

of phase-resetting curves to guide how stimulus strength and timing should be iterated towards a stimulus that induces maximum phase dispersion^{52,53}. One limitation of the topological method is the presence of apparent discontinuities in phase resetting curves, which defies classification and complicates the search for successful stimuli^{54–56}. Our results identify critically timed and shaped waveforms that are also optimized energetically, without the need for topological classification of resetting behavior.

Optimal solutions can be phase dependent. It is important to note that in all of our optimization algorithms, the stimulus found is optimized specifically for the parameters by which we set up the problem. In the FitzHugh-Nagumo model, the optimization protocol that we examined is for a particular stimulus of a specified duration starting at a specific phase of the limit cycle. In the coupled oscillator model, the optimization protocol that we examined is for a particular stimulus of a specified duration starting at a particular set of initial conditions for synchronization. In both of these applications, the state transition occurs only within a narrow range of phases. When the stimulus is given at other phases, there is no state transition. Figure 10 illustrates this narrow window by plotting out the phase resetting curves of the five coupled oscillators in response to the stimulus. As we can see from this figure, the stimulus causes the 5 neurons to exhibit highly variable phases following stimulation (cophase, θ), only when the optimized stimulus is given within a narrow window of phases in the cycle (old phase, φ). When the stimulus is given at all other phases, the reset times (θ) of the 5 neurons do not exhibit dispersion of phases. Standard definitions of phase and cophase were used⁵².

Optimal waveforms are non-rectangular. Using the method we described in our previous work²⁷, we determined an optimal rectangular waveform pulse for each of the models we used. The optimal rectangular pulse energy for inducing an action potential in the Hodgkin-Huxley model is $49 \mu\text{J}/\text{cm}^2$, compared to the optimal energy using our algorithm, $15.6 \mu\text{J}/\text{cm}^2$. Suppressing repetitive firing in the FitzHugh-Nagumo model using an optimal rectangular pulse intensity (L^2 -norm) of 6.4×10^{-2} compared to waveform computed from our algorithm, 7.3×10^{-3} . In the model of a network of coupled oscillators, the optimal rectangular pulse was $154.88 \mu\text{J}/\text{cm}^2$, while the optimal we found was $11.80 \mu\text{J}/\text{cm}^2$.

When examining the optimal stimulus waveform for the Hodgkin-Huxley model, we do note that there is a hyperpolarizing component followed by a depolarizing component. The ionic mechanisms of this particular shape are explored more fully by Clay *et al.*⁵⁷. As previously suggested, these biphasic waveforms in ionic systems have resonant-like properties⁵⁸.

Video illustrating the evolution of optimal waveforms and the system responses in the four models are available in Supplementary Videos S1–S4.

Discussion

One of the challenges in understanding critical transitions is determining the ideal dynamical features of inputs necessary to induce them. This study sought to determine these ideal features by finding optimal stimulus waveforms necessary to induce state transitions by focusing on extrema features using a stochastic optimization approach. In three separate biological models, our approach developed solutions that closely match those found using traditional variational techniques.

There have been alternative model-independent techniques used in the past to find these optimal stimulus waveforms. Spike-triggered averaging was one such stochastic technique that computed the average input stimulus preceding each spike generated from broad-band stimulation^{26,59}, though it was shown that the topology of

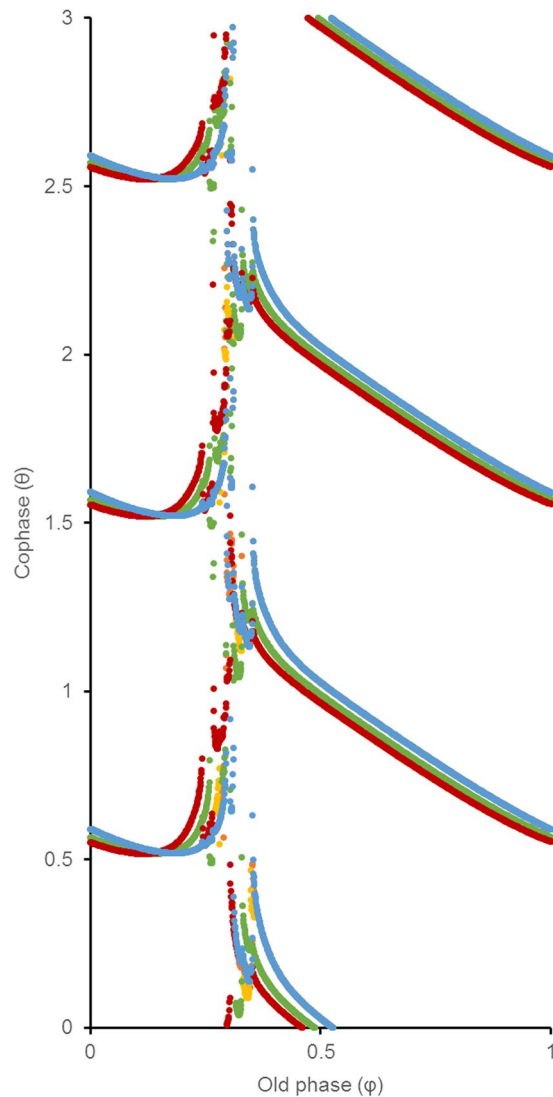


Figure 10. Phase resetting curves of the coupled oscillator network. Old phase (φ) is defined as the time, normalized to one cycle length, between the onset of the burst to the onset of the stimulus. Cophase (θ) is defined as the time, normalized to one cycle length, between the end of the stimulus and the next action potential for each of the neurons. We define the cycle length as the time between the onsets of two consecutive bursts immediately prior to stimulation. The data for each neuron is depicted with a different color.

the solution space had a large biasing effect away from optimality⁴⁰. Stochastic search algorithms that leveraged evolutionary computation, like genetic algorithms, have also been applied to finding the optimal stimulus waveform³⁹. This previous study used every point of the stimulus waveform as a distinct dimension, and thus required many iterations to converge towards a solution.

In order to reduce the dimensionality of the solution space, not of the model, we explored the idea of feature extraction. In the field of time-series analysis, feature extraction is a necessary step in solving problems like time series subsequence data classification^{60–62}, signal pattern recognition^{63–65}, and anomaly detection^{66,67} as often the data set is so large that huge amounts of computational power would be necessary to compare and analyze every data point in the set. While many studies use features like Fourier coefficients or eigenvectors^{64,68}, there has been increasing interest in the use of extrema because of its ease of capture, relationship to variation in the data, and resilience in the face of disruptions^{49,65,69}. To our knowledge, ours is the first study on the use of extrema applied to search for optimal stimulus waveforms.

As we have shown, extracting the extrema and using them as the key feature points of the stochastic search algorithm has greatly improved the results of the algorithm. From this foundation, we believe that the algorithm can be further developed for greater improvement: initial noise reduction, statistical learning, and further feature analysis and extraction. We discuss in Supplementary Text S2 some of the preliminary work that we did in the reducing the noise of the starting seed, our findings and areas of further exploration. With regards to statistical learning, the algorithm as it currently stands utilizes a memoryless Gaussian distribution in order to choose the next stimulus to try. If we can incorporate memory into the algorithm, we may be more efficient with our search

protocols such that the algorithm explores unknown solutions more frequently than repeating waveforms similar to ones that it has tested previously. Finally, in our existing protocol, we utilize only the first order extrema points (e.g. local minima and maxima) causing us to lose some information. We have not characterized or quantified the information loss by using only the first-order extrema, but future studies may help us gain a better appreciation of how much information is lost, whether greater orders of extrema (such as the inflection points between extrema) are necessary to retain as much characterization of the waveform while still maintaining a certain level of computational efficiency. Furthermore, by gaining a better understanding of the robustness of each data point, we may gain more insight with regards to other potential landmarks to use.

In our exploration into computational biology, we have found many more biological models that can be studied, from cardiology (examining cardioversion or defibrillation from arrhythmias^{70–72}) to respiratory (stimulating respiratory rhythms in patients with apnea)⁷, pharmaceutical sciences (studying optimal dosing patterns for drug treatments^{73,74}). Furthermore, it may be interesting to model not just the behavior of the biological system in response to electrical stimuli, but also the electrical and physical components that interface with it (e.g. electrodes, circuits, chips). As some researchers have noted, while we may be able to cut down on the energy consumption of the signal itself, the energetic requirements of the hardware may overwhelm any energy savings on the stimulus waveform itself³⁴. Searching for optimality in the entire system from electrode to biology may yield further insights into how we can improve the biological response to exogenous stimuli.

The use of extrema features in stochastic search could also potentially be used to not just search for transitions between multiple distinct states, but also perhaps to shift systems along a continuum, like in the phase resetting of clocks. What is the optimal light therapy necessary to minimize jet lag or adverse effects of shift work^{53,75,76}? While some are tackling this from an analytical and quantitative perspective using models, perhaps different insights may be gained from approaching the problem with a model-independent perspective.

References

1. Ferrell, J. E. & Machleder, E. M. The biochemical basis of an all-or-none cell fate switch in *Xenopus* oocytes. *Science* **280**, 895–898 (1998).
2. Weinberg, R. A. The retinoblastoma protein and cell cycle control. *Cell* **81**, 323–330 (1995).
3. Zhang, D. W. *et al.* RIP3, an energy metabolism regulator that switches TNF-induced cell death from apoptosis to necrosis. *Science* (80-.). **325**, 332–336 (2009).
4. Tyson, J. J., Chen, K. C. & Novak, B. Sniffers, buzzers, toggles and blinkers: Dynamics of regulatory and signaling pathways in the cell. *Curr. Opin. Cell Biol.* **15**, 221–231 (2003).
5. Lilly, J. C., Austin, G. M. & Chambers, W. W. Threshold movements produced by excitation of cerebral cortex and efferent fibers with some parametric regions of rectangular current pulses (cats and monkeys). *J. Neurophysiol.* **15**, 319–341 (1952).
6. Earhart, G. & Bastian, A. Form switching during human locomotion: traversing wedges in a single step. *J. Neurophysiol.* **84**, 605–615 (2000).
7. Bloch-Salisbury, E., Indic, P., Bednarek, F. & Paydarfar, D. Stabilizing immature breathing patterns of preterm infants using stochastic mechanosensory stimulation. *J. Appl. Physiol.* **107**, 1017–27 (2009).
8. Saper, C. B., Fuller, P. M., Pedersen, N. P., Lu, J. & Scammell, T. E. Sleep state switching. *Neuron* **68**, 1023–1042 (2010).
9. Dahlem, M. A. *et al.* Towards dynamical network biomarkers in neuromodulation of episodic migraine. *Transl. Neurosci.* **4**, 282–294 (2013).
10. Scheffer, M., van den Berg, A. & Ferrari, M. D. Migraine strikes as neuronal excitability reaches a tipping point. *PLoS One* **8**, 1–4 (2013).
11. Segel, L. A., Jäger, E., Elias, D. & Cohen, I. R. A quantitative model of autoimmune disease and T-cell vaccination: Does more mean less? *Immunol. Today* **16**, 80–84 (1995).
12. Lown, B. Electrical reversion of cardiac arrhythmias. *Br. Heart J.* **29**, 469–89 (1967).
13. Walcott, G. P. *et al.* Choosing the optimal monophasic and biphasic waveforms for ventricular defibrillation. *J. Cardiovasc. Electrophysiol.* **6**, 737–750 (1995).
14. Keener, J. & Panfilov, A. A biophysical model for defibrillation of cardiac tissue. *Biophys. J.* **71**, 1335–45 (1996).
15. Hammond, C., Bergman, H. & Brown, P. Pathological synchronization in Parkinson's disease: networks, models and treatments. *Trends Neurosci.* **30**, 357–64 (2007).
16. Rubin, J. & Terman, D. High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model. *J. Comput. Neurosci.* **16**, 211–235 (2004).
17. Hauptmann, C., Popovych, O. & Tass, P. A. Effectively desynchronizing deep brain stimulation based on a coordinated delayed feedback stimulation via several sites: A computational study. *Biol. Cybern.* **93**, 463–470 (2005).
18. Loddenkemper, T. & Pan, A. Deep brain stimulation in epilepsy. *J. Clin. Neurophysiol.* **116**, 217–34 (2001).
19. Tass, P. A. A model of desynchronizing deep brain stimulation with a demand-controlled coordinated reset of neural subpopulations. *Biol. Cybern.* **89**, 81–88 (2003).
20. Scheffer, M. *Critical Transitions in Nature and Society* (Princeton University Press, 2009).
21. Scheffer, M. *et al.* Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
22. Kuehn, C. A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics. *Physica D: Nonlinear Phenomena* **240**, 1020–1035 (2011).
23. Olde Rikkert, M. G. M. *et al.* Slowing down of recovery as generic risk marker for acute severity transitions in chronic diseases. *Crit. Care Med.* **44**, 1–6 (2016).
24. Offner, F. Stimulation with minimum power. *J. Neurophysiol.* **9**, 387–390 (1946).
25. Jezernik, S. & Morari, M. Energy-optimal electrical excitation of nerve fibers. *IEEE Trans. Biomed. Eng.* **52**, 740–743 (2005).
26. Forger, D. B. & Paydarfar, D. Starting, stopping, and resetting biological oscillators: In search of optimum perturbations. *J. Theor. Biol.* **230**, 521–532 (2004).
27. Chang, J. & Paydarfar, D. Switching neuronal state: optimal stimuli revealed using a stochastically-seeded gradient algorithm. *J. Comput. Neurosci.* **37**, 569–582 (2014).
28. Wilson, D. & Moehlis, J. Locally optimal extracellular stimulation for chaotic desynchronization of neural populations. *J. Comput. Neurosci.* 243–257 (2014).
29. Danzl, P., Nabi, A. & Moehlis, J. Charge-balanced spike timing control for phase models of spiking neurons. *Discret. Contin. Dyn. Syst.* **28**, 1413–1435 (2010).
30. Moehlis, J., Shea-Brown, E. & Rabitz, H. Optimal inputs for phase models of spiking neurons. *J. Comput. Nonlinear Dyn.* **1**, 358 (2006).
31. Nabi, A. & Moehlis, J. Time optimal control of spiking neurons. *J. Math. Biol.* **64**, 981–1004 (2012).

32. Ellinger, M., Koelling, M. E., Miller, D. A., Severance, F. L. & Stahl, J. Exploring optimal current stimuli that provide membrane voltage tracking in a neuron model. *Biol. Cybern.* **104**, 185–195 (2011).
33. Tahayori, B. & Dokos, S. Optimal stimulus current waveshape for a hodgkin-huxley model neuron. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS* 4627–4630 (2012).
34. Foutz, T. J. & McIntyre, C. C. Evaluation of novel stimulus waveforms for deep brain stimulation. *J. Neural Eng.* **7**, 66008 (2010).
35. Barnett, W., O'Brien, G. & Cymbalyuk, G. Bistability of silence and seizure-like bursting. *J. Neurosci. Methods* **220**, 179–89 (2013).
36. Feng, X. J., Greenwald, B., Rabitz, H., Shea-Brown, E. & Kosut, R. Toward closed-loop optimization of deep brain stimulation for Parkinson's disease: Concepts and lessons from a computational model. *J. Neural Eng.* **4**, L14–L21 (2007).
37. Kuncel, A. M. *et al.* Clinical response to varying the stimulus parameters in deep brain stimulation for essential tremor. *Mov. Disord.* **21**, 1920–1928 (2006).
38. Spall, J. C. *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control* (John Wiley and Sons, 2003).
39. Wongsarnpigoon, A. & Grill, W. Energy-efficient waveform shapes for neural stimulation revealed with a genetic algorithm. *J. Neural Eng.* **7**, 1–20 (2010).
40. Chang, J. & Paydarfar, D. Optimal stimulus waveforms for eliciting a spike: How close is the spike-triggered average? *Int IEEE EMBS Conf Neural Eng. 2015* 2015–July, 414–417 (2015).
41. Donoho, D. L. Aide-Memoire. High-dimensional data analysis: The curses and blessings of dimensionality. *Am. Math. Soc. Lect. Challenges 21st Century* 1–33 (2000).
42. Bertsekas, D. *Dynamic Programming*. (Athena Scientific, 2005).
43. Hodgkin, A. L. & Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **52**, 500–544 (1952).
44. FitzHugh, R. Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.* **1**, 445–466 (1961).
45. Nagumo, J., Arimoto, S. & Yoshizawa, S. An active pulse transmission line simulating nerve axon. *Proc. IRE* **50**, 2061–2070 (1962).
46. Gardner, T. S., Cantor, C. R. & Collins, J. J. Construction of a genetic toggle switch in *Escherichia coli*. *Nature* **403**, 339–342 (2000).
47. Vemulapalli, P. K., Monga, V. & Brennan, S. N. Optimally robust extrema filters for time series data. *Proc. Am. Control Conf.* 2189–2195 (2012).
48. Lindeberg, T. Image matching using generalized scale-space interest points. *Proc. Int. Conf. Scale Sp. Var. Methods Comput. Vis.* **7893**, 335–367 (2013).
49. Witkin, A. Scale-space filtering: A new approach to multi-scale description. *ICASSP '84. IEEE Int. Conf. Acoust. Speech, Signal Process.* **9**, 150–153 (1984).
50. Gelfand, I., Fomin, S. & Silverman, R. *Calculus of Variations* (Courier Dover Publications, 2000).
51. Jaseja, H. EEG-desynchronization as the major mechanism of anti-epileptic action of vagal nerve stimulation in patients with intractable seizures: Clinical neurophysiological evidence. *Med. Hypotheses* **74**, 855–856 (2010).
52. Winfree, A. T. *The Geometry of Biological Time: Interdisciplinary Applied Mathematics* (Springer-Verlag, 2001).
53. Forger, D. *Biological Clocks, Rhythms, and Oscillations: The Theory of Biological Timekeeping* (MIT Press, 2017).
54. Glass, L. & Winfree, A. Discontinuities in phase-resetting experiments. *Am. J. Physiol.* **246**, R251–R258 (1984).
55. Krogh-Madsen, T., Butera, R., Ermentrout, G. B. & Glass, L. Phase resetting neural oscillators: Topological theory versus the real world. *Phase Response Curves in Neuroscience* **6** (2012).
56. Krogh-Madsen, T., Glass, L., Doedel, E. J. & Guevara, M. R. Apparent discontinuities in the phase-resetting response of cardiac pacemakers. *J. Theor. Biol.* **230**, 499–519 (2004).
57. Clay, J., Forger, D. & Paydarfar, D. Ionic mechanism underlying optimal stimuli for neuronal excitation: Role of Na⁺ channel inactivation. *PLoS One* **7**, e45983 (2012).
58. Paydarfar, D., Forger, D. B. & Clay, J. R. Noisy inputs and the induction of on-off switching behavior in a neuronal pacemaker. *J. Neurophysiol.* **96**, 3338–3348 (2006).
59. Forger, D. B., Paydarfar, D. & Clay, J. R. Optimal stimulus shapes for neuronal excitation. *PLoS Comput. Biol.* **7**, e1002089 (2011).
60. Subasi, A. EEG signal classification using wavelet feature extraction and a mixture of expert model. *Expert Syst. Appl.* **32**, 1084–1093 (2007).
61. Li, L. *et al.* Data mining techniques for cancer detection using serum proteomic profiling. *Artif. Intell. Med.* **32**, 71–83 (2004).
62. Übeyli, E. D. Adaptive neuro-fuzzy inference system for classification of ECG signals using Lyapunov exponents. *Comput. Methods Programs Biomed.* **93**, 313–321 (2009).
63. Berg, A. C., Berg, T. L. & Malik, J. Shape matching and object recognition using low distortion correspondences. *IEEE Conf. Comput. Vis. Pattern Recognit.* **1**, 26–33 (2004).
64. Olszewski, R. T. *Generalized feature extraction for structural pattern recognition in time-series data.* (2001).
65. Vemulapalli, P. K., Monga, V. & Brennan, S. N. Robust extrema features for time-series data analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1464–1479 (2013).
66. Yairi, T., Kato, Y. & Hori, K. Fault detection by mining association rules from house-keeping data. *Proc. 6th Int Symp on Artificial Intelligence, Robotics, and Automation in Space* **18**, 21 (2001).
67. Keogh, E., Lonardi, S. & Chiu, B. 'Yuan-chi'. Finding surprising patterns in a time series database in linear time and space. *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.* 550–556 (2002).
68. Morchen, F. Time series feature extraction for data mining using DWT and DFT. *Tech. Report, No. 33, Dep. Math. Comput. Sci. Univ. Marburg, Ger.* 1–31 (2003).
69. Perng, C.-S., Wang, H., Zhang, S. R. & Parker, D. S. Landmarks: A new model for similarity-based pattern querying in time series databases. *Data Eng. 2000. Proceedings. 16th Int. Conf.* 33–42 (2000).
70. Fain, E. S., Sweeney, M. B. & Franz, M. R. Improved internal defibrillation efficacy with a biphasic waveform. *Am. Heart J.* **117**, 358–364 (1989).
71. Fishler, M. G. Theoretical predictions of the optimal monophasic and biphasic defibrillation waveshapes. *IEEE Trans. Biomed. Eng.* **47**, 59–67 (2000).
72. Kodoth, V. *et al.* Waveform optimization for internal cardioversion of atrial fibrillation. *J. Electrocardiol.* **44**, 689–693 (2011).
73. Tzafirri, A. R. *et al.* Mathematical modeling and optimization of drug delivery from intratumorally injected microspheres. *Clin. Cancer Res.* **11**, 826–834 (2005).
74. Wong, P. K. *et al.* Closed-loop control of cellular functions using combinatory drugs guided by a stochastic search algorithm. *Proc. Natl. Acad. Sci. USA* **105**, 5105–10 (2008).
75. Serkh, K. & Forger, D. Optimal schedules of light exposure for rapid correcting circadian misalignment. *PLoS Comput. Biol.* **10**, e1003523 (2014).
76. Dean, D. A., Forger, D. B. & Klerman, E. B. Taking the lag out of jet lag through model-based schedule design. *PLoS Comput. Biol.* **5**, e1000418 (2009).

Acknowledgements

We thank Peter Grigg, Premananda Indic and Zhiping Weng for their advice, discussions and guidance through this project. This work was submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at the University of Massachusetts Medical School. This work was supported by NIH R01 GM104987

and the Clayton Foundation for Research. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author Contributions

J.C. participated in the algorithm conception and development, wrote the MATLAB code, collected the data, and drafted the manuscript. D.P. participated in the algorithm conception and development and helped draft the manuscript. All authors gave final approval for publication.

Additional Information

Supplementary information accompanies this paper at <https://doi.org/10.1038/s41598-018-21761-8>.

Competing Interests: The authors declare no competing interests.

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 4.0 International License, which permits use, sharing, adaptation, 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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018