



Optimizing stimulus waveforms for electroceuticals

Joshua Chang¹ · David Paydarfar^{1,2}

Received: 7 May 2018 / Accepted: 1 August 2018 / Published online: 11 August 2018
© Springer-Verlag GmbH Germany, part of Springer Nature 2018

Abstract

There has been a growing interest in the use of electrical stimulation as a therapy across diverse medical conditions. Most electroceutical devices use simple waveforms, for example sinusoidal or rectangular biphasic pulses. Clinicians empirically tune the waveform parameters (e.g. amplitude, frequency) without altering the fundamental shape of the stimulus. In this article, we review computational strategies that have been used to optimize the shape of stimulus waveforms in order to improve clinical outcomes, and we discuss potential directions for future exploration.

Keywords Electrical stimulation · Bioelectronic medicine · Calculus of variations · Stochastic search · Optimization algorithms

1 Introduction

Electrical stimulation has been a therapeutic modality used across a number of different medical disciplines to revert pathological states back to normal healthy states. In cardiology, electrical stimulation has been the mainstay for defibrillators, cardioverters and pacemakers (Lown 1967; Kouwenhoven 1969). Psychiatry turns to electroconvulsive therapies and deep brain stimulation for refractory depression, mania and other psychiatric conditions (Jimenez et al. 2005; Dettling and Lisanby 2008; Lozano et al. 2008). Neurologists use electrical currents in deep brain stimulators and vagus nerve stimulators to treat Parkinsonian tremors and epilepsy (Boon et al. 2007; Sun et al. 2008), while spinal cord stimulators have shown benefits to patients suffering from chronic back pain (Oakley and Prager 2002; Cameron 2004).

Physiatrists apply electrical stimulation to muscles preventing atrophy in immobilized patients with musculoskeletal or neurological injuries. More recently, electrically stimulating the vagus nerve has shown reduction to pathological levels of inflammation in animal models and proof-of-concept human studies regarding rheumatoid arthritis and other autoimmune or inflammatory disorders (van Maanen et al. 2009; Meregani et al. 2011; Bonaz and Bernstein 2013). In all of these different applications, targeted electrical stimulation is used to cause a physiological state change from pathological to healthy states.

As with all therapies, one of the major challenges in the use of electrical stimulation is determining the appropriate stimulus strength, or dosing protocol. In some of the earliest recorded instances of electrical stimulation as a therapy, live electric rays were the source of electrical shocks given to patients (Stillings 1974; Kane and Taub 1975). Practitioners were only capable of controlling the duration of the stimulus. After the invention of electrostatic generators in the seventeenth century, practitioners were capable of controlling the strength, duration and frequency of the stimulus. In this context, our concept of dosages match that of the pharmaceutical industry, where we can only control the strength and timing of the therapy. One could represent the stimulus as a rectangular pulse given at predetermined instances.

While most studies focus on the optimization of electroceutical dosing protocols with regard to the strength, duration, timing and frequency of simple waveforms, for example, sinusoidal or rectangular (Kuncel and Grill 2004; Wilson and Moehlis 2016; Cassar et al. 2017), less work

Communicated by Peter J. Thomas.

This article belongs to the Special Issue on *Control Theory in Biology and Medicine*. It derived from a workshop at the Mathematical Biosciences Institute, Ohio State University, Columbus, OH, USA.

✉ Joshua Chang
joshua.chang@austin.utexas.edu

David Paydarfar
david.paydarfar@austin.utexas.edu

¹ Department of Neurology, Dell Medical School, The University of Texas at Austin, Austin, TX 78701, USA

² The Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78701, USA

has been done with regard to discovering optimized waveform shapes. Today with digital technology, completely customized stimuli can be generated giving us unprecedented ability to rapidly adjust and tune them in terms of not only their strength, duration and frequency, but also their fundamental shape, optimizing performance both biologically and electrically. As with pharmacological stimulation, weak stimulation fails to elicit the desired health outcome, while excessive stimulation can cause damage to tissue structure and secondary adverse effects (Merrill et al. 2005; Benabid et al. 2009). The protocols with regard to pharmacological dosages are built upon the pharmacokinetics and pharmacodynamics of the particular compound that is used. Unfortunately, such equivalents are not readily available with electroceuticals. While there have been studies examining the ionic mechanisms and dynamics of individual neurons (Hodgkin and Huxley 1952; Izhikevich 2003), we are still in the dark as to what are the optimal stimulus characteristics necessary to achieve specific health outcomes in complex network systems.

To be clear, defining optimality itself is specific to the nature of the problem. For electrical stimulation of tissue, clinicians may need to choose whether to optimize peak energy versus total energy consumption in order to minimize tissue damage, propagation of energy to neighboring regions, or maximize battery life in implantable devices. Others may instead seek to minimize the time to effect of the stimulus, willing to forego energy optimization in order to maximize how quickly the shock will effect a change in the patient (Moehlis et al. 2006; Nabi and Moehlis 2012). Still others may be interested in the power consumption required from the stimulus on the electrical circuitry (Wongsarnpi- goon and Grill 2010). Sengupta and Stemmler (2014) provide an interesting and in-depth review of the relevance of energy efficiency in biological systems both from the perspective of stimulation as well as system design that examines the importance of optimizing around system responses and not just stimuli characteristics. For the purposes of this article, we highlight various approaches that have been and are being used to find optimal stimulus waveforms. Most of the algorithms and discussions are centered on energy efficiency optimization for the purposes of simplicity, but the techniques discussed can be easily adapted to other optimization objectives.

2 Mathematical framework

Mathematically, one can set up this problem as follows. Given a nonlinear set of equations representing a biological system:

$$\dot{x} = f[x(t), u(t)] \quad (1)$$




Fig. 1 Visualization of solution space in 3D (left) and topologically (right). Each point in the x - y plane represents a unique solution. The z axis represents how well the stimulus has performed as measured by the electrical properties of the stimulus and/or the system's response. We convert the 3D space to a topological plot for easier analysis in subsequent figures. The orange dot represents the global optimal solution

where $x(t)$ describes an m -dimensional system and $u(t)$ describes an n -dimensional external stimulus to the system, we seek an optimal stimulus, $u(t)$, that minimizes the scalar performance index, J , such that:

$$J = \int L[x(t), u(t), t] dt + \text{penalty} \quad (2)$$

where $L[x(t), u(t), t]$ is a performance metric, and the penalty term penalizes the performance index if the system fails to achieve a desired state transition. For much of the work done in electrical stimulation, the performance index is often structured around energetics, seeking to minimize the stimulus' energy, or mathematically, the L^2 -norm of the stimulus:

$$L[x(t), u(t), t] = u^2. \quad (3)$$

As an example, one may want to find the optimal stimulus waveform that causes a quiescent neuron to fire an action potential using the least amount of energy. In this example, a desired state transition may be defined such that:

$$x_1(0 \leq t \leq 50 \text{ ms}) \geq 15 \text{ mV}, \quad (4)$$

where we are searching for the optimal stimulus that causes the voltage membrane potential of the neuron to cross a threshold (15 mV) by 50 ms after the start of the stimulus, penalizing the performance metric if this condition fails. Here, the *penalty* term would be binary: zero if the condition is met and an extremely large value if the condition is not met. The penalty value is often chosen to be much larger than what the performance metric would normally be such that it essentially rules out any stimulus that fails to achieve the desired state transition.

Because we are working with digital technology, one can represent the stimulus, $u(t)$ as a vector of discrete intensities over time, $u[N]$, where N represents a discretized time point. Each of these discrete intensities could be viewed as a

unique dimension in a multidimensional space. Thus, a 30-ms stimulus sampled at a 0.1-ms resolution would result in a 300-dimensional space. Figure 1 graphically represents this space using a 2-dimensional reduction of the stimulus space.

3 Systematic search

The bulk of the work done in optimizing electrical stimulation is through systematic experimentation. Because there are an infinite number of waveforms, researchers limit the search space to a predetermined shape, optimizing specific parameters (e.g. duration, amplitude) associated with that shape. For instance, a sinusoidal waveform may be assumed, following the mathematical notation:

$$u(t) = A \sin(Bt) \quad (5)$$

where A represents the amplitude of the stimulus waveform and B represents the frequency. These parameters are then systematically explored for optimality.

For more than a century and a half, practitioners of electrical stimulation often use this systematic approach to finding optimal protocols in a variety of fields. Garratt (1858) was one of the first to use electrotherapy to provide dental analgesia in his Boston practice. At that time, he was able to vary only the intensity of the current, and he remarked after performing a number of studies, that in cases of low intensity, slight pain was felt due to the stimulus itself, while at higher levels of intensity, a painless sensation was reported. We can imagine that a standard shape of the stimulus was used here, and Garratt systematically tried a number of intensities for his patients.

This is still one of the more prevalent methods that most clinicians use today. For example, deep brain stimulators used to control symptoms of Parkinsonism generate a fundamental waveform, usually a rectangular biphasic fixed frequency pulse train. Clinicians adjust only a few specific parameters to optimize outcomes, empirically guided by the patient's response to stimulation.

One of the first pushes away from the mold of using standard pulses was in the field of defibrillation. In the mid-twentieth century, there was a lot of interest in using electrical stimulation to treat cardiac disorders (Gurvich and Yuniev 1947; Lown et al. 1962). While sinusoidal waveforms were the mainstay initially through alternating current, direct current became more popular over time as the waveforms themselves could be modified and adjusted to improve therapeutic results. Lown and his colleagues constructed many different analog circuits, generating a number of different waveforms in order to optimize the success rate of cardioversion (Lown 1967; Kouwenhoven 1969). Similar studies have been carried out more recently in controlling arrhythmias in




Fig. 2 While systematically search through the entire solution space (left) results in a more optimal solution, it requires evaluation of an exponentially larger number of solutions as compared to systematically searching through a restricted search space (right). The rectangular box in the right panel represents the restricted search space constructed by assuming a fundamental waveform shape

the heart using both computational models as well as animal models (Luther et al. 2011; Li et al. 2009). In neurostimulation, Foutz and McIntyre (2010) have studied the effects of a few different fundamental waveforms (rectangular, exponential, triangular, Gaussian and sinusoidal pulse shapes), tuning the parameters of these fundamental waveforms and comparing their energy consumption and efficacy.

Because there is an infinite number of possible waveform shapes, the restriction to a predetermined shape simplifies the problem making the number of possibilities more manageable, but it also precludes finding the global optimal solution. Figure 2 demonstrates this idea by illustrating first the systematic search through the entire solution space, and then the systematic search only through a restricted portion of the space. When the system is restricted to a small portion of the complete solution space, it biases the solution, at times toward a completely different result.

4 Analytical approaches to solving computational models

In order to address this problem, researchers have sought to define some of these problems as a set of mathematical equations that can be solved. Computational models have been useful surrogates for biological systems in that they provide a basis for mathematical analysis. Grill (2015) reviews the benefits of using computational models in the design of efficient neural stimulation waveforms, while Fishler (2000) demonstrates their usefulness in predicting theoretical optimal waveforms for cardiac defibrillation later verified through experimentation (Malkin et al. 2006).

Analytical methods for finding optimal control fall into roughly two categories: indirect methods and direct methods. The indirect methods construct a set of necessary requirements for the optimal solution using calculus of variations (Gelfand et al. 2000) or Pontryagin's Maximum Principle (Pontryagin et al. 1962) while the direct methods use only




Fig. 3 Calculus of variations finds optimality by defining additional constraints that define where optimality occurs. The lines represent constraints defined by the Lagrange multipliers as determined using the results from calculating P_u , P_x , and P_λ

the system of equations themselves and attempt to find optimal solutions using an iterative descent approach (Bryson and Ho 1975; Golfetto and Silva 2012). Both the indirect method as well as the direct method have been used in finding optimal waveforms, and we discuss them in greater detail below.

In calculus of variations, the performance metric is expanded by including the system equations as constraints. Thus, instead of our optimization to be solely around (3), we are now optimizing around

$$P = L[x(t), u(t), t] + \lambda (\dot{x} - f[x(t), u(t)]) \tag{6}$$

where λ represents a Lagrange multiplier, L represents the original performance metric (e.g. L^2 -norm), and P represents our new expanded performance metric. Now with this expression, Euler equations can be applied:

$$P_u = \frac{dP_{u'}}{dt}, P_x = \frac{dP_{x'}}{dt}, P_\lambda = \frac{dP_{\lambda'}}{dt} \tag{7}$$

to generate the constraints for the optimal solution. These equations are the foundation of the indirect method as they now form a new system of equations that defines the optimal stimulus as seen in Fig. 3.

Offner (1946) was one of the first to optimize neural stimulation using calculus of variations. He began with the idea that the optimal waveform was one which stimulated the cell with the least amount of injury, which he proposed was proportional to heat generated. To simplify his initial set of equations, he assumed that only brief currents were used, and thus accommodation or inhibition was not essential to describing the system. Using these assumptions, he applied calculus of variations to the system equations, determining that the optimal waveform to excite the tissue with the least amount of injury was a rising exponential. From an engineering perspective at the time of his work, Offner acknowledged that exponentially rising current was technically challenging, and thus he analyzed square waves, exponentially falling

currents, and sine waves, technically simpler designs, determining the energy-optimal requirements for each waveform.

Since then, our understanding of the various biological systems has continued to advance, which has furthered our ability to develop mathematical models of their mechanisms and processes. Advances in recording and stimulating electrode technology has also provided the basis for implementing electroceutical therapies. Hodgkin and Huxley (1952) studied the ionic channels in a squid giant axon, formulating the standard model of axonal nerve conduction. As technology advanced, computational modeling has been able to construct more accurate and more realistic systems of equations that better model ionic mechanisms underlying neuronal excitation.

Unfortunately, as these computational models became more accurate, they also became more complex, and as such they no longer retained the simplicity that Offner benefited from when he derived the closed form for optimal waveforms. Recently, Forger et al. (2011) applied the Euler–Lagrange equations to the Hodgkin–Huxley model. They chose a set of terminal conditions for the state variables after a rectangular pulse was given. With these equations and terminal conditions, they obtained a numerical approximation of a special case of optimal excitation of the neuron. Clay et al. (2012) later showed that the specific shape of the optimal waveform was directly related to the mechanisms of the ionic channels in the Hodgkin–Huxley model.

The indirect method derives a set of first-order constraints by which the optimal stimulus needs to follow. Forger et al. recognized that a closed-form solution could not be derived for the four-dimensional Hodgkin–Huxley model. Thus, numerical analysis techniques like the shooting method or another boundary value problem (BVP) solver were required to approximate a solution from the set of equations. If found, the solution can be verified quickly against constraints derived through calculus of variations. The difficulty with the indirect methods is that the region of convergence is extremely small, especially with the addition of the Lagrange multipliers. Practically, it becomes more likely than not for a BVP solver to diverge away from the optimal solution. Developing the initial guess for the BVP solver becomes an extremely important, but difficult problem to solve.

The direct method does not create a surrogate system of equations, but instead marches iteratively toward optimality as seen in Fig. 4 (Kelley 1962; Bryson and Ho 1975). The algorithm begins with an estimate of the stimulus variable, $u(t)$. Using this estimate and the system’s response to it, it calculates two influence functions:

$$\dot{p}^T = -\frac{\partial L}{\partial x} - p^T \frac{\partial f}{\partial x}, \tag{8}$$

$$\dot{R} = -\left(\frac{\partial f}{\partial x}\right)^T R. \tag{9}$$

In these calculations, $p(t)$ represents the effect that changes to the stimulus have on the performance index, while $R(t)$ represents the effect that changes to the stimulus have on the error in terminal conditions. These two influence functions describe how changes in the stimulus affect the performance index and the distance from the expected terminal conditions. From here, a multiplier is determined that weighs the effect of both influence functions:

$$v = -Q^{-1}[\delta x(t_f) + g] \tag{10}$$

where

$$Q = k \int R^T \frac{\partial f}{\partial u} \left(\frac{\partial f}{\partial x} \right)^T R dt, \tag{11}$$

$$g = k \int R^T \frac{\partial f}{\partial u} \left\{ \left(\frac{\partial f}{\partial u} \right)^T p + \left(\frac{\partial L}{\partial u} \right)^T \right\} dt, \tag{12}$$

$$\delta x(t_f) = \epsilon[x(t_f) - x_f]. \tag{13}$$

This multiplier is then used to calculate an improved estimate of the stimulus:

$$\delta u = -k \left\{ \left(\frac{\partial f}{\partial x} \right)^T [p + vR] + \left(\frac{\partial L}{\partial u} \right)^T \right\}, \tag{14}$$

Because it is an iterative approach, it allows for a greater degree of flexibility in the initial seed compared to the indirect method. Chang and Paydarfar (2014) demonstrated the robustness of this algorithm to randomly generated initial seeds. The limitations of the direct method is that it is ultimately an approximation of the optimal solution. For their work, a first-order gradient system was used, which allowed for the increased robustness in the early stages. Unfortunately, this method can have difficulty converging when it gets close to the optimal solution due to the first-order gradients. Furthermore, there is no guarantee that the solution found using the direct method is indeed the optimal solution. As can be seen from Fig. 4, these methods are susceptible to getting stuck in locally optimal solutions. This can be mitigated by starting from multiple different starting points. Lastly, unlike the indirect method, there is no set of equations defining optimality that can validate the results. For a more in-depth review on direct and indirect methods, we refer the reader to Stryk and Bulirsch (1992) and Rao (2009).

With either direct or indirect methods, these approaches become difficult to manage when dealing with complex mathematical models. As more information becomes available regarding the biological systems in question, they are often incorporated as equations into the mathematical model, increasing its analytical complexity. One approach to studying these complex systems has been through dimensional reduction of the models to more manageable systems of equations. While the study of this process falls out of this review’s




Fig. 4 Gradient-based algorithms derive first-order gradients given the system of equations, and iteratively converge toward locally (right) or globally (left) optimal solutions. As can be seen, depending on the starting solution, the algorithm may converge toward a locally optimal solution using gradient-based algorithms

domain, we refer the reader to Indic et al. (2006) and Forger (2017)). An example of this process as applied to optimal stimulus waveforms can be found in Wilson and Moehlis (2014), where they used calculus of variations on a phase reduced neuron model in order to simplify the mathematical model into something more manageable. As seen in this study, dimensional reduction can help simplify some complex mathematical models into simpler systems that can be solved more easily using analytical means.

5 Model-independent approaches to finding optimal solutions

An alternative approach to dealing with these increasingly complex systems has been through the use of model-independent approaches of finding optimality. One of the interesting techniques recently developed has been the use of phase resetting curves to calculate the energy-optimal DBS stimulus necessary to desynchronize a population of neurons. Wilson and Moehlis (2014) derive a method using calculus of variations that requires only knowledge of a neuron’s phase resetting curve. Once that phase resetting curve is captured, which can be done experimentally, optimal control can be calculated for desynchronizing a population of neurons. This particular method does well with balancing the use of analytical techniques with some experimental exploration to determine the phase response curves of the neurons.

In most situations however, the phase resetting curve is not easily obtained. Stochastic search approaches have been used in other fields as a means to find solutions in black box systems, using just the inputs and outputs of the system to gain insight into where optimality should be. For good reviews of stochastic search approaches and other methods to solve high-dimensional, black box systems see Spall (2003) and Shan and Wang (2010).

Forger and Paydarfar (2004) examined the possibility of using spike-triggered averaging as a noise-driven search process, comparing it to the results of calculus of variations (Paydarfar et al. 2006). From their studies, the results are extremely close, showing the potential for spike-triggered averaging to approximate the optimal solution. Further studies showed that spike-triggered averaging did indeed come close, but depending on the characteristics of the noise and the solution space (e.g. multiple local extrema), the spike-triggered averaging may be biased away from the optimal solution (Chang and Paydarfar 2015).

Wongsarnpigoon and Grill (2010) tackled this problem using genetic algorithms, one of the more common approaches in stochastic search algorithms. Genetic algorithms mimic the biological process of evolution by treating solutions as chromosomes, allowing solutions to learn from each other by “breeding together.” Each stimulus is considered a chromosome and the stimulus intensity at each time point represents a single gene. In their study, the genetic algorithm was applied to a population of 100 parallel McIntyre, Richardson, and Grill (MRG) axons distributed uniformly within a cylinder. The stimulation was modeled to be delivered through a point current source within the center of the cylinder. They determined the optimized waveform given different pulse widths, and showed that the optimal solution had a Gaussian-like structure. To confirm their work, they applied the optimized waveform to the sciatic nerve of an adult male cat, showing that the stimulus waveform did indeed successfully trigger an action potential, while also outperforming rectangular and decaying exponential waveforms. This study showed the power of using stochastic search algorithms to finding optimized waveforms. While the search was done in a computational model, they were able to validate the results in an animal model. One of the interesting discussion points made in their study was that the optimized waveform was slightly different depending on the model used. The MRG model of the neuron used stands in contrast to the simpler Hodgkin–Huxley model in that it accounts for paranodal sections as well as more ion channel dynamics. The differences in the model used led to the unique optimal waveforms found.

When considering the patient population, each patient’s pathophysiology or underlying cause for their disease may indeed be slightly different from the next patient who may have the same disease, but whose underlying states (e.g. their neuronal wiring) may be slightly different. What is optimized for one patient, may not necessarily be optimized for another patient. Thus, the challenge becomes not just whether we can find an optimal solution, but can we find an optimal solution for the specific patient in front of us. Wongsarnpigoon and Grill’s work showed that using a genetic algorithm, one did not have to understand the underlying system in order to develop an optimal waveform.

One of the main advantages of genetic algorithms is their ability to search across large solution spaces, not getting caught in locally optimal solutions, through the use of crossover. This process treats each individual gene as an independent entity. In biology, this assumption is often not correct as biological stimuli are more often smooth and not sharp and jagged. Stochastic hill-climbing is a more rudimentary approach that is often more efficient than genetic algorithms, but does carry the risk of getting caught in locally optimal solutions. Systematic studies and comparisons have been made between these two approaches in computer science (Mitchell et al. 1994; Ross and Corne 1995; Prugel-Bennett 2004).

This challenge was recently studied using a stochastic hill-climbing approach (Chang and Paydarfar 2018), mitigating the problem of getting caught in locally optimal solutions by running multiple instances of the algorithm from many different randomly generated starting seeds. Stochastic hill-climbing starts with a randomly generated stimulus, searches a few of the neighboring stimuli, picks the best one and uses that as the starting point for the next iteration of search. It continues to do this over and over again, always choosing the best seen stimulus as the starting point for the next iteration. The results and a sample of the intermediate iterations can be seen in Fig. 5.

One of the biggest challenges using this method, and perhaps the cause of the large number of stimuli required for testing in stochastic search approaches in general, is the “curse of dimensionality”, which states that as the number of dimensions increases linearly, the search space increases exponentially (Bellman 1961). When searching for optimal stimuli, the stimulus intensity at each time point becomes a unique dimension. A 30-ms stimulus sampled at a 0.1-ms resolution, becomes a 300-dimension problem. The use of extrema as key features points was proposed to reduce the dimensionality of the stimulus. In searching for neighboring stimuli, only the extrema of the starting stimulus was perturbed and the intermediate points were linearly transformed, reducing the dimensionality of the search space from the number of discrete time points in the stimulus to the number of extrema in the stimulus. This method was able to match closely with results generated from gradient approaches.

Furthermore, one of the interesting aspect of this algorithm is that it iteratively learns and seeks more optimal solutions, adapting to the plasticity inherent in biological systems. By using the best seen stimulus as the starting seed for each iteration, it retains some, but not complete, knowledge about previously explored stimuli. This adaptability is important as biological systems are known to adapt to a repeated stimulus over time. For instance, if a spinal cord stimulator is giving the exact same stimulus over and over again, the body may adapt to that particular stimulus, decreasing its efficacy over time. Algorithms that are constantly learning and adapting




Fig. 5 Evolution of the stimulus waveform using an extrema featured stochastic hill-climbing approach. Adapted from Chang and Paydarfar (2018) with permission

may well be able to change and find new waveforms so that the efficacy remains high in these patients. Some stochastic search approaches utilize the results from all previously explored stimuli to construct a probabilistic map of the solution space (Shahriari et al. 2015) as opposed to retaining just the best stimulus between iterations in order to gain insights into how to converge more quickly. While these approaches do allow for more efficient convergences, they are unable to adapt if the biological system in question undergoes a major change as it relies on data collected under a different paradigm.

6 Discussion

We have reviewed a number of different ways researchers are tackling the problem of finding optimal stimulus waveforms for electrical stimulation as a medical therapy. Much work has been done using systematic approaches, but creative approaches are being leveraged to try and break the mold, finding waveforms tailored to the patient's needs and requirements. The ultimate goal is to one day have these protocols built into the electroceutical system itself, such that the device would find optimal stimulus waveforms, adapting constantly to the patient's responses.

There are a few different means to approach this problem. First, better computational models can enlighten our ability to monitor and predict patient's responses to stimuli. Many researchers are looking to develop these new models, in cardiology [examining cardioversion or defibrillation (Fain et al. 1989; Fishler 2000; Kodoth et al. 2011)], pulmonology [stimulating respiratory rhythms in patients with apnea (Bloch-Salisbury et al. 2009)], neurology (understanding neuronal connections and information transfer in the brain), and pharmacology [studying optimal dosing patterns for drug treatments (Tzafiriri et al. 2005; Wong et al. 2008)] among others. With advances in electrode technology and devices, we are able to capture larger amounts of biological data than ever before, at higher resolutions, leading to more accurate computational models. With the exponential growth of processing power, these computational models may lead the way in developing more optimal waveforms to test on in vitro and in vivo systems in the future. Furthermore, modeling these systems may no longer be restricted to just the biological tissue in question, but potentially the entire system including the electrical interfaces. This inclusion may address some of the issues that Offner and others noted in that even if the stimulus itself is determined to be energetically optimal, the energy requirements to generate that stimulus may overwhelm the energy savings of the stimulus itself (Offner 1946; Foutz and McIntyre 2010).

Certain problems however, may be too complex or too large to model accurately, and thus stochastic search approaches may be more useful. Although they have been shown to find optimal stimulus waveforms, these approaches can be improved further to increase their efficiency. First, most of these algorithms inject randomness into the stimulus waveform to search for more optimal solutions. Very little information is passed on from one iteration to the next. The inclusion of some sort of memory may help these algorithms avoid testing and re-testing solutions from the same area. For instance, in the extrema featured stochastic search approach, only the best solution is passed on from one iteration to the next. Every failed stimulus is ignored and removed. As stated previously, this process is beneficial for the algorithm regarding adaptability, but one could assume that there is an ideal medium between adaptability and efficiency that could be studied in greater depth. Using Bayesian analysis, one may be able to develop a decaying uncertainty boundary around the stimulus space such that regions of the space that have failed in the recent past may be sampled from less than regions that have had greater successes. Some of this work has already been explored in drug combinations studies (Park and Pillow 2012; Park et al. 2013), but they often explore spaces with a smaller number of dimensions. Future work in Bayesian active learning or a related field may open up approaches that leverage statistical learning so that every

stimulus, whether failed or successful, will be useful in finding optimal waveforms.

Another aspect of stochastic searches that needs to be addressed in future studies is its efficacy under noisy conditions. Because information gained from each stimulus is assumed to be true in all subsequent iterations, there is an assumption that the biological system is the same from one test stimulus to the next. This is rarely the case in biology. Future research may develop new algorithms that can find optimal stimulus waveforms in spite of stochastic conditions. In fact, some of the latest work in reinforcement learning may provide some insight into how to reframe the question. Current algorithms seek to find an optimal stimulus waveform to take a patient from a state A to state B. Reinforcement learning algorithms would develop rules, or “policies”, that pair a patient’s given state with a stimulus, or an “action”, which leads to a more optimal state (Barto and Sutton 2017). These methods have had great success recently in problems like the game Go, in which a complete solution is not tractable, but optimized actions can be determined for various states of the board (Silver et al. 2016, 2017). One could learn policies that gave smaller bursts of customized stimuli based on the “state” the patient was in. These methods described so far only touch upon a few of the potentially relevant algorithms and techniques that are being used in the broader context of optimization. As the field of electroceuticals continues to advance in electrode technology and neurological mapping, these methods both past and future may open new doors and treatments focused on intelligent and effective waveforms.

Acknowledgements This work was supported by the Clayton Foundation for Research and NIH R01 GM104987.

References

- Barto AG, Sutton RS (2017) Reinforcement learning: an introduction. MIT Press, Cambridge
- Bellman R (1961) Adaptive control processes: a guided tour. Princeton University Press, Princeton
- Benabid AL, Chabardes S, Mitrofanis J, Pollak P (2009) Deep brain stimulation of the subthalamic nucleus for the treatment of Parkinson’s disease. *Lancet Neurol* 8(1):67–81
- Bloch-Salisbury E, Indic P, Bednarek F, Paydarfar D (2009) Stabilizing immature breathing patterns of preterm infants using stochastic mechanosensory stimulation. *J Appl Physiol* 107(4):1017–27
- Bonaz BL, Bernstein CN (2013) Brain–gut interactions in inflammatory bowel disease. *Gastroenterology* 144(1):36–49
- Boon P, Vonck K, De Herdt V, Van Dycke A, Goethals M, Goossens L, Van Zandijcke MV, De Smedt T, Dewaele I, Achten R, Wadman W, Dewaele F, Caemaert J, Van Roost D (2007) Deep brain stimulation in patients with refractory temporal lobe epilepsy. *Epilepsia* 48(8):1551–1560
- Bryson AE, Ho YC (1975) Applied optimal control: optimization, estimation and control. CRC Press, Boca Raton
- Cameron T (2004) Safety and efficacy of spinal cord stimulation for the treatment of chronic pain: a 20-year literature review. *J Neurosurg Spine* 100(3):254–267
- Cassar IR, Titus ND, Grill WM (2017) An improved genetic algorithm for designing optimal temporal patterns of neural stimulation. *J Neural Eng* 14(6):066013
- Chang J, Paydarfar D (2014) Switching neuronal state: optimal stimuli revealed using a stochastically-seeded gradient algorithm. *J Comput Neurosci* 37(3):569–582
- Chang J, Paydarfar D (2015) Optimal stimulus waveforms for eliciting a spike: How close is the spike-triggered average? In: 2015 7th international IEEE/EMBS conference on neural engineering (NER), pp 414–417
- Chang J, Paydarfar D (2018) Evolution of extrema features reveals optimal stimuli for biological state transitions. *Sci Rep* 8:3403
- Clay J, Forger D, Paydarfar D (2012) Ionic mechanism underlying optimal stimuli for neuronal excitation: role of Na⁺ channel inactivation. *PLoS ONE* 7(9):e45983
- Dettling M, Lisanby SH (2008) Electroconvulsive therapy for depression. *New Engl J Med* 358(6):645–646
- Fain ES, Sweeney MB, Franz MR (1989) Improved internal defibrillation efficacy with a biphasic waveform. *Am Heart J* 117(2):358–364
- Fishler MG (2000) Theoretical predictions of the optimal monophasic and biphasic defibrillation waveshapes. *IEEE Trans Biomed Eng* 47(1):59–67
- Forger D (2017) Biological clocks, rhythms, and oscillations: the theory of biological timekeeping. MIT Press, Cambridge
- Forger DB, Paydarfar D (2004) Starting, stopping, and resetting biological oscillators: in search of optimum perturbations. *J Theor Biol* 230(4):521–532
- Forger DB, Paydarfar D, Clay JR (2011) Optimal stimulus shapes for neuronal excitation. *PLoS Comput Biol* 7(7):e1002089
- Foutz TJ, McIntyre CC (2010) Evaluation of novel stimulus waveforms for deep brain stimulation. *J Neural Eng* 7(6):066008
- Garratt A (1858) Galvanism for counteracting pain in the extraction of teeth. *Boston Med Surg J* 59:32–34
- Gelfand I, Fomin S, Silverman R (2000) Calculus of variations. Courier Dover Publications, Englewood Cliffs
- Golfetto WA, da Silva Fernandes S (2012) A review of gradient algorithms for numerical computation of optimal trajectories. *J Aerosp Technol Manag* 4(2):131–143
- Grill W (2015) Model-based analysis and design of waveforms for efficient neural stimulation. *Prog Brain Res* 222:147–162
- Gurvich N, Yuniev G (1947) Restoration of heart rhythm during fibrillation by a condenser discharge. *Am Rev Sov Med* 4(3):252–256
- Hodgkin A, Huxley A (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 52(4):500–544
- Indic P, Gurdziel K, Kronauer RE, Klerman EB (2006) Development of a two-dimension manifold to represent high dimension mathematical models of the intracellular mammalian circadian clock. *J Biol Rhythm* 21(3):222–232
- Izhikevich E (2003) Simple model of spiking neurons. *IEEE Trans Neural Netw* 14(6):1569–1572
- Jimenez F, Velasco F, Salin-Pascual R, Hernandez JA, Velasco M, Criales JL, Nicolini H (2005) A patient with a resistant major depression disorder treated with deep brain stimulation in the inferior thalamic peduncle. *Neurosurgery* 57(3):585–592
- Kane K, Taub A (1975) A history of local electrical analgesia. *Pain* 1(2):125–138
- Kelley HJ (1962) Methods of gradients. In: Leitmann G (ed) Optimization techniques, chap 6, 5th edn. Academic Press, Inc, New York, pp 206–254

- Kodoth V, Castro NC, Glover BM, Anderson JM, Escalona OJ, Lau E, Manoharan G (2011) Waveform optimization for internal cardioversion of atrial fibrillation. *J Electrocardiol* 44(6):689–693
- Kouwenhoven WB (1969) The development of the defibrillator. *Ann Intern Med* 71(3):449–458
- Kuncel AM, Grill WM (2004) Selection of stimulus parameters for deep brain stimulation. *Clin Neurophysiol* 115(11):2431–2441
- Li W, Ripplinger CM, Lou Q, Efimov IR (2009) Multiple monophasic shocks improve electrotherapy of ventricular tachycardia in a rabbit model of chronic infarction. *Heart Rhythm* 6(7):1020–1027
- Lown B (1967) Electrical reversion of cardiac arrhythmias. *Br Heart J* 29(4):469–89
- Lown B, Neuman J, Amarasingham R, Berkovits B et al (1962) Comparison of alternating current with direct electroshock across the closed chest. *Am J Cardiol* 10(2):223–233
- Lozano AM, Mayberg HS, Giacobbe P, Hamani C, Craddock RC, Kennedy SH (2008) Subcallosal cingulate gyrus deep brain stimulation for treatment-resistant depression. *Biol Psychiatry* 64(6):461–467
- Luther S, Fenton FH, Kornreich BG, Squires A, Bittihn P, Hornung D, Zabel M, Flanders J, Gladuli A, Campoy L, Cherry EM, Luther G, Hasenfuss G, Krinsky VI, Pumir A, Gilmour RF Jr, Bodenschatz E (2011) Low-energy control of electrical turbulence in the heart. *Nature* 475(7355):235–239
- Malkin RA, Jackson SR, Nguyen J, Yang Z, Guan D (2006) Experimental verification of theoretical predictions concerning the optimum defibrillation waveform. *IEEE Trans Biomed Eng* 53(8):1492–1498
- Meregnani J, Clarençon D, Vivier M, Peinnequin A, Mouret C, Sinniger V, Picq C, Job A, Canini F, Jacquier-Sarlin M, Bonaz B (2011) Anti-inflammatory effect of vagus nerve stimulation in a rat model of inflammatory bowel disease. *Auton Neurosci Basic Clin* 160(1–2):82–89
- Merrill DR, Bikson M, Jefferys JG (2005) Electrical stimulation of excitable tissue: design of efficacious and safe protocols. *J Neurosci Methods* 141(2):171–198
- Mitchell M, Holland JH, Forrest S (1994) When will a genetic algorithm outperform hill climbing. *Adv Neural Inf Process Syst* 6:51–58
- Moehlis J, Shea-Brown E, Rabitz H (2006) Optimal inputs for phase models of spiking neurons. *J Comput Nonlinear Dyn* 1(4):358
- Nabi A, Moehlis J (2012) Time optimal control of spiking neurons. *J Math Biol* 64(6):981–1004
- Oakley JC, Prager JP (2002) Spinal cord stimulation: mechanisms of action. *Spine* 27(22):2574–2583
- Offner F (1946) Stimulation with minimum power. *J Neurophysiol* 9(5):387–390
- Park M, Pillow J (2012) Bayesian active learning with localized priors for fast receptive field characterization. *Adv Neural Inf Process Syst* 25(25):2357–2365
- Park M, Nassar M, Vikalo H (2013) Bayesian active learning for drug combinations. *IEEE Trans Biomed Eng* 60(11):3248–3255
- Paydarfar D, Forger DB, Clay JR (2006) Noisy inputs and the induction of on–off switching behavior in a neuronal pacemaker. *J Neurophysiol* 96(6):3338–3348
- Pontryagin LS, Boltyanskii V, Gamkrelidze R, Mishchenko E (1962) *Mathematical theory of optimal processes*. Wiley, New York
- Prugel-Bennett A (2004) When a genetic algorithm outperforms hill-climbing. *Theor Comput Sci* 320:135–153
- Rao AV (2009) A survey of numerical methods for optimal control. *Adv Astronaut Sci* 135(1):497–528
- Ross P, Corne D (1995) Comparing genetic algorithms, simulated annealing, and stochastic hillclimbing on timetabling problems. In: *AISB workshop on evolutionary computing*, pp 94–102
- Sengupta B, Stemmler MB (2014) Power consumption during neuronal computation. *Proc IEEE* 102(5):738–750
- Shahriari B, Swersky K, Wang Z, Adams RP, Freitas ND (2015) Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE* 104(1):1–24
- Shan S, Wang GG (2010) Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions. *Struct Multidiscip Optim* 41(2):219–241
- Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M, Dieleman S, Grewe D, Nham J, Kalchbrenner N, Sutskever I, Lillicrap T, Leach M, Kavukcuoglu K, Graepel T, Hassabis D (2016) Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587):484–489
- Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A, Chen Y, Lillicrap T, Hui F, Sifre L, Van Den Driessche G, Graepel T, Hassabis D (2017) Mastering the game of Go without human knowledge. *Nature* 550(7676):354–359
- Spall JC (2003) *Introduction to stochastic search and optimization: estimation, simulation, and control*. Wiley, Hoboken
- Stillings D (1974) A survey of the history of electrical stimulation for pain to 1900. *Med Instrum* 9:255–259
- Sun FT, Morrell MJ, Wharen RE (2008) Responsive cortical stimulation for the treatment of epilepsy. *Neurotherapeutics* 5(1):68–74
- Tzafirri AR, Lerner EI, Flashner-Barak M, Hinchcliffe M, Ratner E, Parnas H (2005) Mathematical modeling and optimization of drug delivery from intratumorally injected microspheres. *Clin Cancer Res* 11(21):826–834
- van Maanen MA, Vervoordeldonk MJ, Tak PP (2009) The cholinergic anti-inflammatory pathway: towards innovative treatment of rheumatoid arthritis. *Nat Rev Rheumatol* 5:229–232
- von Stryk O, Bulirsch R (1992) Direct and indirect methods for trajectory optimization. *Ann Oper Res* 37(1):357–373
- Wilson D, Moehlis J (2014) Optimal chaotic desynchronization for neural populations. *SIAM J Appl Dyn Syst* 13(1):276–305
- Wilson D, Moehlis J (2016) Toward a more efficient implementation of antifibrillation pacing. *PLoS ONE* 11(7):1–28
- Wong PK, Yu F, Shahangian A, Cheng G, Sun R, Ho CM (2008) Closed-loop control of cellular functions using combinatory drugs guided by a stochastic search algorithm. *Proc Natl Acad Sci* 105(13):5105–5110
- Wongsarnpigoon A, Grill WM (2010) Energy-efficient waveform shapes for neural stimulation revealed with a genetic algorithm. *J Neural Eng* 7(4):1–20