Published in final edited form as:

Int IEEE EMBS Conf Neural Eng. 2015 April ; 2015: 414–417. doi:10.1109/NER.2015.7146647.

Optimal stimulus waveforms for eliciting a spike: How close is the spike-triggered average?

Joshua Chang and

University of Massachusetts Medical School, Worcester, MA 01604 USA

David Paydarfar

Department of Neurology, University of Massachusetts Medical School, Worcester, MA 01604 USA and the Wyss Institute for Biologically Inspired Engineering, Harvard University, Boston, MA 02115 USA

Abstract

Computing the average input stimulus preceding a spike, the spike-triggered average (STA), has been a powerful tool for discovering a neuron's 'preferred' stimulus feature that enables efficient encoding of sensory information. Recent work in the squid giant axon has shown that STA waveforms can be remarkably similar to the energetically optimal stimulus waveforms for eliciting a spike. In the present study, we show using the Hodgkin-Huxley model that the STA can deviate from the global optimal solution if there is averaging of multiple solutions across different time scales and of multiple modes of spike induction. These findings inform attempts to develop model-free stochastic algorithms for finding energy-optimal stimuli, which is relevant to the efficient delivery of exogenous therapeutic stimuli in neurological diseases.

I. Introduction

For nearly half a century, spike-triggered averaging (STA) has been used in experimental neuroscience to gain insight into neuronal responses to stimulus features [1]–[6]. STA using broadband stochastic stimulation has revealed mechanisms that govern how inherently noisy sensory information might be encoded by neurons.

STA waveforms rendered from noisy current stimulation are remarkably similar to optimal stimuli that elicit an action potential with the least amount of applied current energy [7]–[9], which is relevant to the efficient delivery of exogenous therapeutic stimuli in neurological conditions [10], [11].

In the present study, we analyze how noise-based STA waveforms relate to energy-optimal waveforms for eliciting a spike. In order to explore a range of neuronal threshold levels, we studied the Hodgkin-Huxley (HH) model [12] under two conditions: the standard condition in which there is no persistent current, and the condition of a large persistent current, which causes the membrane to be at a heightened state of excitability near a sub-critical Hopf bifurcation [13].

II. Methods

A. Hodgkin-Huxley Model

The HH model is a four-dimensional system that captures the ionic mechanisms underlying the generation of an action potential. By adjusting the persistent current, we can raise the excitability of the membrane such that much less stimulus energy is required to fire an action potential. Using different levels of persistent current, we are able to study differences in stimulus optimality related to a range of membrane thresholds, for example between high threshold interneurons as compared to low thresholds seen in some cortical neurons. Also, raising the persistent current affects the resonance properties of the membrane such that the neuron is more responsive to oscillatory stimuli [13].

B. Determining optimality using a stochastically seeded gradient algorithm

In order to provide a gold standard for optimality, we rely on analytic methods. In this paper, we calculate energy-optimal waveforms by minimizing the L^2 -norm of the stimulus. Traditionally, optimization of signals has been conducted using calculus of variations [14], in which an optimal functional is derived and solved as a boundary-value problem. This approach has been used to calculate energyoptimal stimuli for the HH model [8]. An alternative approach has solved this problem directly using a stochastically-seeded gradient algorithm [15]. We use this approach for the HH model, to determine the optimal stimulus necessary to trigger a single action potential in both high and low persistent current conditions.

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. These small gradients, or slopes, are then used to determine a better solution. This algorithm iterates forward until it converges to a single solution. Stochastic seeding enables the algorithm to explore a large solution space. In the case of the standard HH model (with no persistent current), we previously showed convergence to a global optimal waveform [15].

C. Finding the spike triggered average using different noise levels

In order to mimic experimental spike-triggered averaging, we applied white noise current stimulation, sampled every 0.1 ms, to the Hodgkin-Huxley model neuron. The current was chosen randomly using a Gaussian distribution at each of the sampled points. The amplitude of the noise was tuned such that the action potentials would fire at a rate of <1 Hz. The white noise was given continuously until we reached 3,000 action potentials. Each time an action potential occurred, the stimulus signal preceding the spike (up to 100 ms) was stored. The average of all 3,000 stimulus snippets (aligned to the peak of the action potential) was scaled to the minimum amplitude for triggering an action potential as previously described [7], [8].

We ran the above protocol for the HH model under standard conditions (no persistent current) and the near-bistable condition (a persistent current of $6\mu\text{A}/\text{cm}^2$). The duration of the stimulus prior to each spike was 30 ms for the standard condition and 100 ms for the

near-bistable condition; the longer averaging period was used to capture the full oscillatory stimulus pattern in the neuron near its sub-critical Hopf bifurcation.

All of the simulations and data analysis was done in MATLAB (The Mathworks, Inc, Natick, MA) on a Dell Precision T7500 (Intel Xeon Core E5520 @ 2.27 GHz, 2.26 GHz processors and 24 GB RAM). The code for the gradient algorithm can be obtained directly from the authors or from Physionet (www.physionet.org).

III. Results

A. Gradient-based calculation of stimulus optimality for eliciting an action potential

As seen in Fig. 1, the standard HH model fires an action potential most efficiently when the stimulus includes a single hyperpolarization followed by depolarization phase. On the other hand, when the neuron is more highly excitable in the near-bistable condition, the stimulus shows more prolonged sub-threshold oscillations with multiple cycles between hyperpolarization and depolarization building up to the action potential.

B. Comparing STA waveforms to optimal waveforms

The STA waveform and the optimal waveform calculated from the gradient algorithms are shown in Fig. 2. The two waveforms are remarkably similar, as reported previously [7],[8], [11]. Two major differences stand out, however. First, the spike-triggered average appears to be right-shifted compared to the gradient algorithm in both the standard as well as the near-bistable conditions. Second, in the nearbistable condition, the spike-triggered average exhibits a phase reversal in the stimulus about 20-ms prior to the action potential.

The STA right-shift can be explained by considering that the waveform is an average of a library of successful stimuli without restriction on the effective portion of the stimulus that causes the spike. For example, a brief large fluctuation that induces a spike would be averaged with a stimulus having longer duration with smaller amplitude fluctuations. In Fig. 3, we show using the gradient algorithm the optimal waveforms for progressively longer stimulus durations, all lined up to the peak of the spike. Note that the shorter stimuli have higher amplitude and are right shifted, compared to the longer duration stimuli. We previously showed that the longer stimuli are more energetically optimal [15]. Because of this phenomenon, we can see that if longer and shorter stimuli are included in the library of snippets that induce a spike, averaging the briefer waveforms with the more prolonged waveforms would result in a right-shift compared to the global optimal calculated from the gradient algorithm. As such, when averaging the optima across different time scales, we get this right-shift.

With regards to the deviation of the spike-triggered averaging from the gradient algorithm in the near-bistable condition, we examined the system's response to all the individual snippets in our library. We noticed that some of them caused single spikes, while others caused two or more spikes in rapid succession. This led us to speculate that the stimulus necessary to cause repetitive spikes might be qualitatively different from that which causes a single spike. We reevaluated the near-bistable condition using the gradient algorithm, and we found that

occasionally an optimal would be computed that causes repetitive spikes. These two stimulus waveforms are shown in Fig. 4.

Informed by this result, we re-analyzed the noise induced spike trains in the near-bistable condition and calculated the STA on the snippets that contained just one action potential. The result was much closer to the energy optimal waveform computed using the gradient algorithm, as seen in Fig. 5. The small rightward shift for the oscillations near the spike are not affected by this re-analysis, presumably because of the mechanism related to short duration stimuli (Fig. 3). This confirms that by averaging different local optimal solutions, or different modes, we get discordance from the true optimal.

We did take into consideration the fact that we are using extremely high frequency noise stimuli, but when we calculated the spike-triggered average with white-noise filtered through a 1-kHz low-pass filter (which is more biologically realistic[1]), we found similar results.

IV. Discussion

Our results suggest that while STA and energy optimal waveforms are remarkably similar, a number of mechanisms contribute to their discordance.

First is a rightward shift in the STA that can be explained by averaging stimuli having a range of time scales that contribute to spike generation. The rightward shift occurs because the briefer stimuli cause a more rapid depolarization and induction of a spike compared to the longer duration stimuli. Because the space of all possible solutions for the Hodgkin-Huxley model includes solutions that are larger, which fire an action potentials sooner, the spike-triggered average waveform is right-shifted from the optimal solution as determined by the gradient algorithm. We find that using larger noise amplitudes can potentially cause a larger deviation away from energetically optimal solutions. It is interesting to note that when choosing a smaller amplitude noise profile, the amplitude of the spike-triggered average also decreased, however the right-shift remained (Fig. 2). In theory, if we choose an amplitude of the noise such that it is close to the amplitude of the optimal solution, we might be able to further reduce the right-shift in the STA because we would not be including into the average the large amplitudes that cause the action potential to fire faster. In practice, however, when we titrated the noise level further toward the amplitude of the optimal solution, spikes disappeared and we were not able to achieve this result.

Second, multiplicity of optima can skew the STA waveform. When there are multiple mechanisms by which the event is triggered, the multiplicity can cause the spike-triggered average to move away from the true optimal solution (Fig. 4).

To gain a better appreciation of the relationship between spike triggered averaging and optimality, let us imagine a stimulus with only two points. This illustration can be extrapolated to much higher dimensions, but for visual purposes we will use a two point stimulus. The space created in Fig. 6 is the set of all possible stimuli. We have marked off two regions within this space. Region A represents all the stimuli that successfully cause an event, while Region B represents all the stimuli that our white noise generator produces given a specific amplitude.

In this depiction, the spike-triggered average is the centroid of Region C, which is the intersection between Region A and Region B. In calculating the energy of the stimuli, we are using a simple L2-norm metric, which is equivalent to the distance between the stimuli and the origin. Thus, the farther a point is from the origin, the larger is its energy. Using this representation, we can begin to gain some insight into the relationship between spike-triggered averaging and energy optimality.

Fig. 6 (top) is a representation of the standard Hodgkin-Huxley neuron. As seen in the figure, the spike-triggered average alone is not close to the optimal, but when scaled to just barely causing an action potential, it is quite close to the true optimal solution. When multiplicity of solutions exists (e.g., as seen in the near-bistable condition), one can imagine a system similar to Fig. 6 (bottom), where there exist two groups of solutions within the stimulus space, the scaled spike-triggered average could look very different when compared to the true optimal solution.

In our experiments, we recognized the existence of two sets of solutions by observing the train of action potentials in response to some of the stimuli. As such, we filtered the data to produce only the stimuli that exist in one of these sets, and then performed spike-triggered averaging, which led us to a result much closer to the true optimal solution. As we stated earlier, we could theoretically also have decreased the noise amplitude, shrinking Region B, to eliminate the possibility of the second solution set intersecting with the noise profile we are using. This is a potential solution, but only one that works if the two solution sets are separated by a large enough margin. Again, it is difficult to build a library of successful stimuli if the noise profile is barely above the amplitudes of the true optimal stimuli. As such, more often than not, the noise will be larger than what one can do to separate out multiple sets of successful stimuli.

While our experiments here have shown that spike-triggered averaging can approximate the energetically efficient optimal solution, we also show examples of when it will fail to do so. We posit that whether spike-triggered averaging is successful or not depends largely on the shape of the space defined by successful stimuli. Techniques that quantify the solution space could enable a more accurate and consistent method to approximate the optimal signal using the library of noisy stimulus waveforms that precede each spike.

Acknowledgments

This research was supported by NIH R01 GM104987

References

1. Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science* (80-). 1995 Jun; 268(5216):1503–1506.
2. Bryant H, Segundo J. Spike initiation by transmembrane current: a white-noise analysis. *J. Physiol.* 1976; 260:279–314. [PubMed: 978519]
3. Schwartz O, Pillow JW, Rust NC, Simoncelli EP. Spike-triggered neural characterization. *J. Vis.* 2006 Jan; 6(4):484–507. [PubMed: 16889482]
4. Cheney P. Response of rubromotoneuronal cells identified by spike-triggered averaging of EMG activity in awake monkeys. *Neurosci. Lett.* 1980; 17:137–141. [PubMed: 6820481]

5. Dick T, Berger A. Axonal projections of single bulbospinal inspiratory neurons revealed by spike-triggered averaging and antidromic activation. *J. Neurophysiol.* 1985; 53(6):1590–1603. [PubMed: 2989449]
6. Lankheet MJM, Klink PC, Borghuis BG, Noest AJ. Spike-interval triggered averaging reveals a quasi-periodic spiking alternative for stochastic resonance in catfish electroreceptors. *PLoS One.* 2012 Jan.7(3):e32786. [PubMed: 22403709]
7. Forger DB, Paydarfar D. Starting, stopping, and resetting biological oscillators: in search of optimum perturbations. *J. Theor. Biol.* 2004 Oct.230:521–532. [PubMed: 15363673]
8. Forger DB, Paydarfar D, Clay JR. Optimal stimulus shapes for neuronal excitation. *PLoS Comput. Biol.* 2011 Jul.7(7):e1002089. [PubMed: 21760759]
9. Clay JR, Forger D, Paydarfar D. Ionic mechanism underlying optimal stimuli for neuronal excitation: role of Na⁺ channel inactivation. *PLoS One.* 2012; 7(9):e45983. [PubMed: 23049913]
10. Feng X, Greenwald B, Rabitz H. Toward closed-loop optimization of deep brain stimulation for Parkinson's disease: concepts and lessons from a computational model. *J. Neural Eng.* 2007; 4(2):L14–L21. [PubMed: 17409470]
11. Hofmann L, Ebert M, Tass PA, Hauptmann C. Modified pulse shapes for effective neural stimulation. *Front. Neuroeng.* 2011 Jan.4:1–10. [PubMed: 21270946]
12. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 1952 Aug; 117(4):500–544. [PubMed: 12991237]
13. Izhikevich, EM. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting.* Cambridge, Massachusetts: MIT Press; 2007.
14. Gelfand, IM., Fomin, SV. *Calculus of Variations.* Courier Dover Publications; 2000.
15. Chang J, Paydarfar D. Switching neuronal state: optimal stimuli revealed using a stochastically-seeded gradient algorithm. *J. Comput. Neurosci.* 2014 Aug.

Figure 1. Energy-optimal stimulus waveform for the near-bistable condition (right) displays more resonance than the optimal waveform for the standard condition for the Hodgkin-Huxley model (left)

Figure 2. Comparison of spike triggered averaging results (orange) and the gradient algorithm (blue). We note two main discordances: a phase reversal (*) and a right-shift (#).

Figure 3.
Optimal waveform shapes given different stimulus durations as determined by the gradient algorithm.

Figure 4.

The near-bistable solution has at least two locally optimal solutions: one that fires a single action potential (red), another that fires two action potentials in rapid succession (blue).

Figure 5.

When averaging stimuli causing only one spike, the STA (red) phase reversal disappears, and STA aligns more closely with the true optimal (blue).

Figure 6.

Whether the spike-triggered average is close to the energy-optimal is related to the shape of the intersection (Region C) between the set of all possible stochastic signal (Region B) and the set of all possible stimuli that successfully trigger an event (Region A). Sometimes when the center of mass (y) is scaled (y'), it may move closer to the optimal stimulus (x) as seen in the top panel, while other times it may move further away as seen in the bottom panel.