Explaining Adaptive Computation Using the Hodgkin-Huxley Model

Tharindu W. Fernando*
Department of Physics, University of California, San Diego, California, CA 92122, USA.
(Dated: December 13, 2018)

Adaptive computation is observed biologically when neuronal firing appears to depend on statistics of synaptic input – such as variance and mean – instead of being fixed for a given neural system [1]. Since the resulting spikes rely on the associated context, the system encodes information about it to resolve ambiguities between stimuli and spikes [3]. To achieve a better understanding of this complex nonlinear phenomenon, adaptive computation was analyzed through the lens of the simple (yet empirical) single-neuron Hodgkin-Huxley (HH) model. This was done with the hope of identifying how simple parameters in the well-studied HH model could result in seemingly perplex behavior that has yet to be fully explained theoretically. This literature review reveals that adaptive computation is tied to the HH model's parameters g_{K} and g_{Na}, which represent maximal conductances of voltage-dependent potassium and sodium ion channels in the axon membrane. However, one has yet to theorize a biological mechanism that explains how g_{K} or g_{Na} change due to input signal statistics.

I. INTRODUCTION

In physiology, cell-to-cell communication is achieved via electrical impulses that travel across axons. This communication is vital for life because it ensures several key processes. For example, impulses govern the autonomous mechanism by which someone withdraws her hand from a sharp object on contact: Skin receptors on the hand capture the stimulus (sharp object) into a signal that is transmitted to the nervous system, which in turn signals an arm muscle to contract so that further exposure to the sharp object is avoided.

Impulses are transmitted across axons by means of action potentials (AP), which occur when the membrane potential of a specific axon location rapidly depolarizes [6]. The voltage-gated ion channels in the membrane are initially closed, maintained at a resting potential specific to the region. If a threshold voltage is reached, the ion channels open to allow inflow of sodium ions Na^+, which increases the membrane potential difference further, causing more channels to open until all available channels are open. This rapid influx of Na^+ causes the membrane polarity to reverse, stopping the Na^+ influx and forcing the cell to actively transport Na^+ out. Finally, potassium ions K^+ are activated, causing them to leave cell and restore the membrane back to its resting potential. This process is summarized in Figure 1.

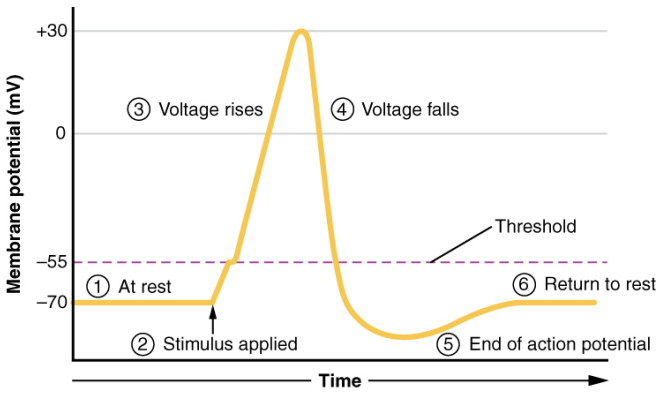

FIG. 1. Plotting voltage measured across the cell membrane against time, the events of the action potential can be related to specific changes in the membrane voltage. (1) At rest, the membrane voltage here is −70 mV. (2) The membrane begins to depolarize when an external stimulus is applied. (3) The membrane voltage begins a rapid rise toward +30 mV. (4) The membrane voltage starts to return to a negative value. (5) Repolarization continues past the resting membrane voltage, resulting in hyperpolarization. (6) The membrane voltage returns to the resting value shortly after hyperpolarization. Source:[6]

* Undergraduate student, twarnaku@ucsd.edu.

Due to the involvement of a threshold voltage, action potentials function like digital events, where either the membrane reaches the threshold and everything in Figure 1 occurs, or the membrane does not reach the threshold and nothing else happens. All AP peak at the same voltage for a given context, so one is not bigger than another. Thus, the person in the previous example will experience more pain or stronger muscle contractions only because multiple AP will be initiated, and not because their amplitudes increase.

This "all or none" behavior led scientists to develop the neural code to efficiently represent the sensory world [3]. Several biological and computational models have evolved over the last century to describe efficient coding that matches the coding strategy to the statistics of input signals. Of these, the Hodgkin-Huxley (HH) model was one of the first to use computational methods to describe how single-neuronal AP are initiated and propagated. This Nobel prize-winning work fit ionic conductance and AP data from squid giant axons into a theoretically-justified model that is now a hallmark in the literature. In it, the axon is modeled as a circuit and described by nonlinear differential equations.

Even though the HH model is not the simplest, it is experimentally more realistic than the theoretically-simpler McCulloch-Pitts and Integrate-and-Fire models [4]. Therefore, if one was to attempt understanding how some *complicated neural phenomenon* relates to both experimental reality and theoretical models, the HH model provides a very reasonable starting point due to its applicability to living systems and relative simplicity. It is relatively simple because it is a single-neuron model (as opposed to a many-layered neural network), and simplifies several biological components by categorizing them into functionally different components.

The *complicated neural phenomenon* explored in this review is adaptive computation. In the context of machine learning, adaptive computation refers to how a computer algorithm makes a learning process more efficient by dynamically changing the method used to process information before finalizing output [7]. Usually, this is a change in the number of steps required.

Somewhat similarly, in biology this term describes the experimental observation by which neural systems adjust their input-output properties in response to changes in the statistical properties of the incoming stimulus. Examples of such statistical properties are variance (or standard deviation) and mean of synaptic input current. To illustrate using a *hypothetical* example from the familiar context of Figure 1, adaptive computation could be when the axon temporarily lowers its threshold voltage from -55 mV to -90 mV because the stimulus it expects is unexpectedly very noisy (making its potential to trigger necessary impulses ambiguous). The reader is reminded that this is but a hypothetical example, since such 'extreme' scenarios were not found in the literature explored[1]. On the other hand, an experimentally-observed example of adaptive computation that builds off this hypothetical one is when adding noise to an input signal increased the firing rate of neocortical pyramidal neurons at low mean currents (but not high mean currents) [2]. This in turn reduced the gain of the system's input-output relation. Therefore, in other words, noisy "background" synaptic input appeared to control the sensitivity of neurons, which continually recalibrated their sensitivity under new contexts to best represent the range of inputs they received.

However, despite advances in research, clear mechanisms for how living systems realize[2] adaptive computation are yet to be discovered [1–3]. A primary issue underlying such discoveries is that adaptive computation is inherently ambiguous: the meaning of a spike or a pattern of spikes depends on context, and resolution of this ambiguity requires that the system additionally encode information about the context itself [1]. This makes it hard to find a general replicable model system for electrophysiological studies, as, for instance, the highly specific conditions required to observe adaptive computation in neurons of the fly visual system [1] are different compared to those necessary to realize it in

[1] Section II discusses this behavior in developing mouse cortex neurons, where it is observed for a limited period.

[2] Other general applications of adaptive computation in machine learning, etc are not entirely clear either [7].

the strikingly different developing mouse cortex neurons [3]. Therefore, having to use diverse, inconsistent systems to theorize a shared, unobvious property (adaptive computation) does not seem to be a straightforward task.

In an attempt to deconstruct this seemingly complicated problem, it was applied to the HH model that was just discussed. Through the model's relative simplicity, one may hope to gauge not only a common mechanism for adaptive computation, but also insight on how it is achieved physiologically.

The remainder of this literature review will have the following structure: The next section presents an overview of the HH model, and the section afterwards will illustrate biophysical aspects of adaptive computation using existing studies. Potential ties between adaptive computation and the HH model will then be discussed, followed by a brief wrap-up of results.

A. Hodgkin-Huxley Model

To understand action potentials, Hodgkin and Huxley performed experiments on the squid giant axon and related its initiation and propagation to three different types of ion current (sodium, potassium, and a leak current that consists mainly of chloride ions). Specific voltage-dependent ion channels, one for sodium and another one for potassium, control the flow of those ions through the cell membrane; which was modeled as a circuit as seen in Figure 2. A capacitor C represents the cell membrane and has a voltage of u across it [5]. The resistors R correspond to each voltage-dependent channel (Na^+, K^+ and leak current L that accounts for other channel types which are not described explicitly). The diagonal arrows across the resistors indicate that values depend on whether the ion channel is open or closed[3]. Additionally, each ion channel is associated by a battery E because the associated active ion transport involves a Nernst potential. Finally, the input signal is denoted by the applied current I.

Using elementary circuit physics and theoretical modeling to fit experimental data, Hodgkin and Huxley formulated the sum of the three ion currents that pass through the cell membrane as:

$$\sum_k I_k = g_{\mathrm{Na}} m^3 h (u - E_{\mathrm{Na}}) + g_{\mathrm{K}} n^4 (u - E_{\mathrm{K}}) + g_L (u - E_L) \tag{1}$$

where the new parameters are: voltage-independent conductance g_L; maximum conductances g_{Na} and g_{K} (which transmitted currents will have if all ion channels are open); and 'gating' variables m, n and h that model the probability that a channel is open at a given time. m and h describe the activation (opening) and inactivation (blocking) of Na^+ channels respectively, while n describes the activation of K^+ channels. Each of these gating variables is in turn controlled by an ordinary differential equation (ODE) so that the entire HH model comprises four coupled ODEs.

Introducing these gating variables was what allowed Hodgkin and Huxley to successfully describe their observations mathematically by measuring how a channel's effective resistance changes as a function of time and voltage. Notice that the applicability of this model to the AP process in Figure 1 is pretty straightforward.

While the reader is referred to any basic exposition of neurophysics (such as [5]) for specifics of the HH model, a discussion of its limitations[4] is vital to the issue of understanding adaptive computation. This is because if adaptive computation cannot be explained fully using the HH model, it could very likely be accounted for by the model's limitations (in addition to insufficient research).

Limitations of The Hodgkin-Huxley Model

There are several limitations associated with the HH model since its simplicity was arrived at only after compromising generalizable behavior via several assumptions and data-fits from a specific species's neuron[5].

[3] Caution: this is *not* adaptive computation because their values do not dynamically change to reflect their environment!

[4] Recall that Section I discusses advantages of the HH model.

[5] Squid giant axon.

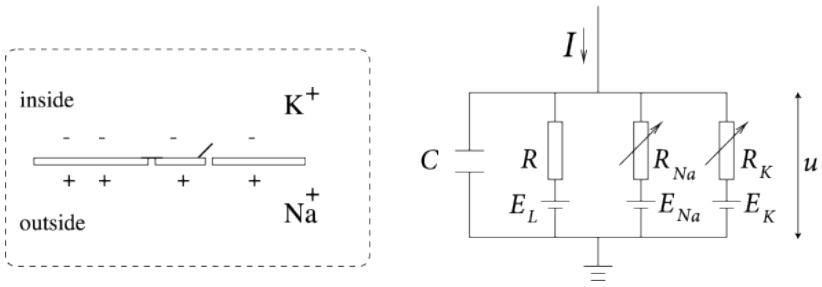

FIG. 2. Schematic diagram for the Hodgkin-Huxley model. Left: The semipermeable cell membrane separates the interior of the cell from the extracellular liquid and acts as a capacitor. Right: Circuit diagram for typical HH neuron membrane. Source:[5]

For instance, it is still not clear how the gating variables m, n and h are controlled (both physiologically and theoretically) [9]. This calls for further research on defining the activation and inactivation of sodium channels, and the activation of potassium channels. As this review will elaborate in subsequent sections, this mystery is crucial to the understanding of adaptive computation.

Another limitation is the model's validity only for a membrane patch. Therefore, its applications at the macromolecular nanoscale are dubious [8]. The issue arises when the transmembrane voltage acquires a transient when it opens from a closed state. Additionally, there is other experimental evidence to support the claim that the HH model does not act alone. For example, the speed of AP transmission is undefined in the model, but ion diffusion measurements have shown that this is not possible.

Furthermore, [9] illustrates how properties of AP generated depend on the strength of input when the HH model is programmatically realized in toy models. However, this conflicts with the indisputable "all or none" trait that characterizes an AP (Section I).

On a broader level, since the HH model involves 4 ODEs with 4 state variables $(u(t), m(t), h(t), n(t))$, the system is hard to solve analytically for any set of initial conditions. Therefore, scientists analyze its properties using center manifold methods and bifurcations (using I as a parameter) [5]. However, these methods do not allow for a complete analysis involving all parameters at a given time[6], and this means that the model can hide key information without warning.

Finally, since this is only a single-neuron model, potential effects of many-layered neural networks on adaptive computation will not be evident.

Any of the limitations discussed in this section might be related to behavior of adaptive computation that cannot be explained using the HH model.

B. Adaptive Computation

Since the general idea behind adaptive computation should be clear from Section I, this section will present experimental evidence of the phenomenon. Using this and knowledge about the HH model from the previous section, one will have sufficient tools to analyze adaptive computation through the simple lens of the HH model.

To reiterate, adaptive computation describes the experimental observation by which neural systems adjust their input-output properties in response to changes in statistical properties of the incoming stimulus. Therefore, appropriate experimental investigations usually comprise electrophysiological measurements of AP generated by neurons subject to various electrical inputs controlled by a computer.

[6] For example, only two parameters can be studied using the center manifold method.

1. Adaptive Computation in Fly Visual System

In one study [3], motion-sensitive neurons in the fly visual system (H1) were subjected to white-noise signals with variances that switched like a step-function (as in Figure 3). It was observed that significant modulations in firing probability required proportionately large variations for an input signal with large dynamic range. However, under different adaptation conditions, the firing probabilities in response to the same input differed by orders of magnitude. This implied that the input-output relation of a system adapts dynamically, and that the system measured input in units proportional to the local standard deviation.

FIG. 3. The stimulus was a white-noise velocity signal modulated by a square wave envelope that switches between two values, σ_1 and σ_2, with some period T. Source:[3]

It was also observed that the decay time of an output spike rate after an upward switch in variance seemed to have a linear relationship to the input signal variance's frequency. However, for sinusoidal input, the relationship was only a fraction of its frequency. This suggested that the adaption observed depended on the experiment's design and that different mechanisms might be responsible for different directions of adaptation.

Additionally, since it was discovered that different aspects of adaptation occur on timescales that range from tens of milliseconds to several minutes, they also suggested a mechanism by which ambiguities in adaptive code due to context are resolved: Rapid adaptation leaves longer timescales in the response dynamics as a nearly independent channel for information.

All these results provide evidence for traits of adaptive computation that were discussed in this review's introduction. However, albeit illustrative of adaptive computation, [3]'s analysis does not provide much insight into simpler single-neuron models because it does not provide variables nor data that obviously correspond to parameters of the HH model in Section I A. Therefore, a different study was sought.

2. Adaptive Computation in Embryonic Mouse Cortex

Another electrophysiological study [1] explored the question of how adaptive computation develops in neurons[7]. Unlike the study in the previous subsection, this work revealed specific aspects of a neural system's components that regulate adaptive computation. This information might allow it to be connected to the HH model!

This study found that adaptive computation emerges during early development as an intrinsic property of single neurons[8] in mouse sensorimotor cortex neurons, and that it can be modulated by changing the balance of spike-generating currents.

As with the experiments on the fly visual system in the previous subsection, these experiments too demonstrated adaptive computation. However, they also revealed a direct tie between the levels of adaptive computation observed and the neurons' developmental stage; specifically with the latter's effect on voltage-gated sodium and potassium ion channels. Recall from Section I A that the HH model involves parameters representing both of these ion channels!

The next section elaborates on these results and their direct application to the HH model.

II. ADAPTIVE COMPUTATION USING THE HODGKIN-HUXLEY MODEL

The authors of [1] state that cortical neurons 'learn' adaptive computation early in their development. Since it is known that: input current I_{K} is handled in the stem cell population of the cortex before the first neurons

[7] Or in machine learning jargon: 'how a neuron learns to adapt'.
[8] This further supports the choice to adopt a single-neuron model like HH.

exit the cell cycle; I_{Na} can be detected even before differentiating neurons migrate into the cortical plate; and I_{Na} increases in density much faster than I_{K} during early postnatal development – the authors examined how these changes in the maturing spike-generation mechanism impacted the adaptive computational properties of cortical neurons.

They found that developing neurons move toward a common intrinsic operating point and a stable ratio of spike-generating currents (implying that adaptive computation will be harder past this 'learning point'). Using Figure 1 to step back and bring the discussion into perspective, this observation could be analogous to the process by which a new neuron sets its own threshold voltages for optimal performance in its environment. However, unlike the 'extreme' hypothetical example illustrated in Section I, these neurons ceased adaptive computation after an optimal threshold voltage was determined, and then maintained this value for the rest of their lives.

To use this information about I_{Na} and I_{K}, the authors applied the HH model using the simulator *NEURON* on data from electrophysiological experiments on mouse cortex neurons. Since the maximal current ratio $I_{\mathrm{Na}}/I_{\mathrm{K}}$ is related to the maximal conductance ratio $g_{\mathrm{Na}}/g_{\mathrm{K}}$ by definition, it helps to use notation commonly associated with the HH model, as in Equation (1). The figure illustrates an apparent relationship between $g_{\mathrm{Na}}/g_{\mathrm{K}}$ and adaptive computation.

By investigating the *NEURON* models' gain-scaling capabilities by stimulating firing with noise of different standard deviations σ_i, the authors also discovered that models with low $g_{\mathrm{Na}}/g_{\mathrm{K}}$ had input-output relations which did not scale completely with σ_i while input-output relations from high $g_{\mathrm{Na}}/g_{\mathrm{K}}$ models were nearly identical for all σ_i. See Figure 5.

After even more investigations (that involved pharmacological methods to manipulate I_{Na} and I_{K}), the authors provided compelling evidence (similar to Figures 4 and 5) to the claim that adaptive computation is related to the ratio $g_{\mathrm{Na}}/g_{\mathrm{K}}$. While this promising result provides some sense of direction as to how one could view adaptive computation through the HH model (because it involves the familiar parameters g_{Na} and g_{K}), it is far from painting a complete picture. The final section of this review discusses why.

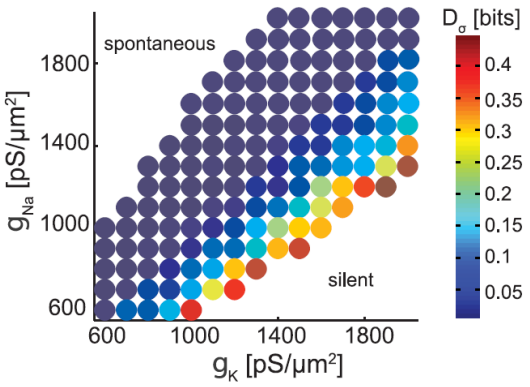

FIG. 4. The degree of adaptive computation D_σ quantified (in bits) for 148 model neurons with varying g_{Na} and g_{K} conductance values, stimulated with two stimulus standard deviations in input signal. Warmer colors indicate stronger presence of adaptive computation. Source:[1]

III. CONCLUSION

The previous section presents a relationship between adaptive computation and a ratio of parameters from the HH model, $g_{\mathrm{Na}}/g_{\mathrm{K}}$. While this is certainly some progress in answering the original question of explaining adaptive computation through the empirical HH model, the exact mechanism through which it is brought about is unclear both theoretically and physiologically. Although [1] established some control over adaptive computation via artificial pharmacological means to control g_{Na} and g_{K}, this method is not representative of natural cell membrane activities. Therefore, how exactly g_{Na} and g_{K} are modified based on input signal statistics is not yet clear.

Before continuing, it should be noted that the experiments discussed in Section I B had their own weaknesses (including extreme specificity to a unique model system). For instance, the data for Figure 4 was generated only by a fraction of all possible neu-

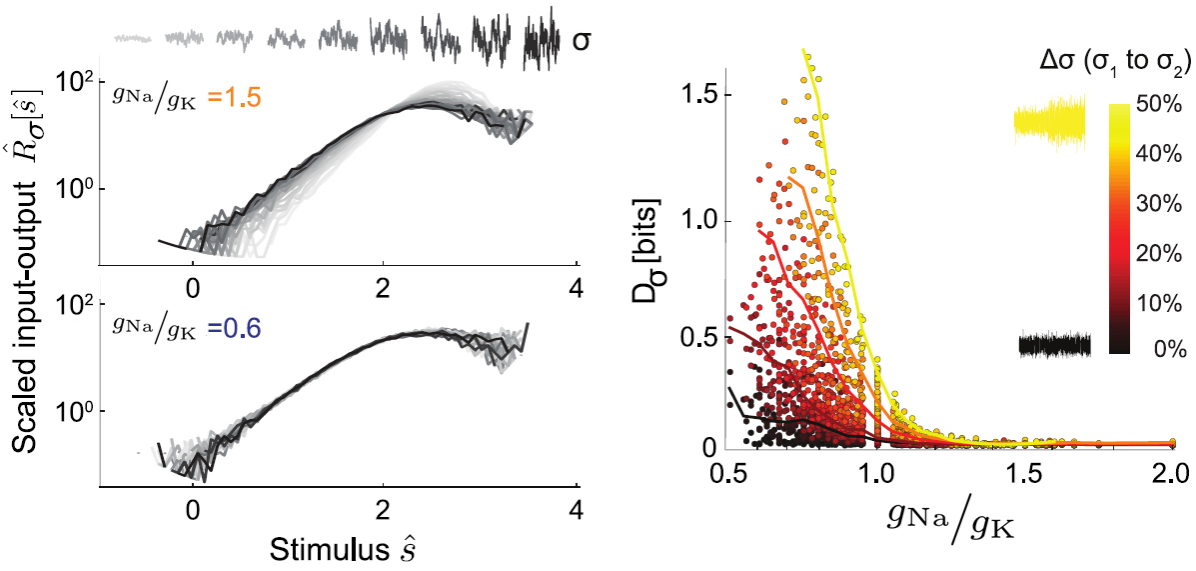

FIG. 5. Left: Scaled input-output range of σ for two model neurons with high and low g_{Na}/g_K. Right: Adaptive computation D_σ (quantified in bits) plotted against g_{Na}/g_K for models stimulated with a large range of σ. Lines indicate mean D_σ values for different levels of $\Delta\sigma$. Source:[1]

ron models because only certain neurons would spike for a given noisy signal[9] [1]. Although such limitations might make the g_{Na}/g_K relationship questionable, its validity is affirmed by similar studies and at least by the reputable profile of the journal it was published in.

It is easy to conjecture that the mysterious mechanisms that dynamically determine g_{Na} and g_K values might consequence directly from the HH model's own limitations. In Section I A, it was noted that further research is needed to understand how the so-called gating variables m, n and h of the model – that define the probability of ion channel activation and inactivation – acquire their values. From Equation (1) for the HH model, it can be seen that the conductances $1/R_i$ (which are not necessarily the maximal conductances g_i) are given by $1/R_{Na} = g_{Na}m^3h$ and $1/R_K = g_K n^4$ for sodium and potassium ion channels respectively. Therefore, the unsolved gating variable problem is bound to be related to the issue of adaptive computation. This limitation of the HH model, in addition to the several others discussed in Section I A, might help identify potential directions of further study for understanding adaptive computation. However, they also question the HH model's suitability for tackling the problem.

To summarize these issues in layman's terms, a reason behind the HH model's potential incompetence in the context of this problem could be its lack of complexity. In other words, it might be "too simple", per the general warning given by Drs. Einstein and Dudko about simplifying reality into physical models. Indeed, the authors of [1] applied an analysis similar to Section II to the so-called Exponential Integrate-and-Fire model – which is a model of the axon membrane that is theoretically less simple than the HH model – and obtained a deeper insight of adaptive computation.

Despite all its potential flaws, the HH model is *not* a poor choice for an approach to understand the complicated phenomenon that is adaptive computation. As justified in this review's introduction, this is due to the model's relative simplicity, its grounding in experimental data, and the fact that it was actually related to adaptive computation in some way.

[9] Such case-dependent requirements for spiking reflects adaptive computation's context-dependent ambiguity.

Evidently, more research is necessary to fully understand adaptive computation. To reiterate, what is not clear is how a cell membrane changes its maximal conductances g_{Na} and g_{K} dynamically based on statistical properties of input signals. Since statistics are involved, perhaps a subsequent study could employ a different perspective and investigate neural models using methods from statistical mechanics. By somehow reformulating this problem using canonical ensembles, a desirable relationship between parameters of the HH model and adaptive computation might be achieved. Other perspectives may be pursued for more insight.

Finally, since this property appears intertwined to the HH model, further studies might symbiotically provide ways to improve the HH model's limitations. Regardless, the goal should be to strive for a simple mechanism that will explain adaptive computation, not only to better our understanding of a fundamental biological process that has been taken for granted for decades, but also to give us insight into adaptive computation in general learning techniques from other disciplines including artificial intelligence, computer science, engineering, mathematics, physics, neuroscience, and cognitive science.

ACKNOWLEDGMENTS

The author is grateful to Dr. Olga Dudko for an enriching quarter of biological physics that changed his perception of general physics for the better and influenced the structure of this review. Simplicity is the way to go! The author also thanks Dr. David Kleinfeld for introducing him to, and discussing this open question with him.

[1] Rebecca A. Mease, Michael Famulare, Julijana Gjorgjieva, William J. Moody, and Adrienne L. Fairhall. *Emergence of Adaptive Computation by Single Neurons in the Developing Cortex.* The Journal of Neurosci. 2013 Jul 24; 33(30):12154 –12170.

[2] Matthew H. Higgs, Sean J. Slee, and William J. Spain. *Diversity of Gain Modulation by Noise in Neocortical Neurons: Regulation by the Slow Afterhyperpolarization Conductance.* The Journal of Neurosci. 2006 Aug 23; 26(34):8787-99.

[3] Adrienne L. Fairhall, Geoffrey D. Lewen, William Bialek, and Robert R. de Ruyter van Steveninck. *Efficiency and ambiguity in an adaptive neural code.* Nature. 2001 Aug 23; VOL 412:787-92.

[4] *"Introduction to Computational Neuroscience." Lecture 6: Single Neuron Models*, Institute of Computer Science, Estonia, 13 Mar. 2014, courses.cs.ut.ee/MTAT.03.291/2014_spring/uploads/Main/Lecture6.pdf.

[5] Wulfram Gerstner, Werner M. Kistler, Richard Naud and Liam Paninski. *"Neuronal Dynamics." Chapter 2: Hodgkin-Huxley Model*, Cambridge University Press, 8 Sep. 2014, neuronaldynamics.epfl.ch/online/Ch2.S2.html.

[6] OpenStax. *"Anatomy and Physiology." Chapter 12.4: The Action Potential*, Rice University, 26 Feb. 2016, cnx.org/contents/14fb4ad7-39a1-4eee-ab6e-3ef2482e3e22@8.24.

[7] Mohan Li, Min Liu, and Hattori Masanori. *End-to-end Speech Recognition with Adaptive Computation Steps.* Arxiv. 2018 Aug. 30; arXiv:1808.10088 [eess.AS].

[8] Adam F. Strassberg, and Louis J. DeFelice. *Limitations of the Hodgkin-Huxley Formalism: Effects of Single Channel Kinetics on Transmembrane Voltage Dynamics.* Neural Computation, Massachusetts Institute of Technology. 1993; 5,843-655.

[9] S. A. Sadegh Zadeh, and C. Kambhampati. *All-or-None Principle and Weakness of Hodgkin-Huxley Mathematical Model.* International Journal of Mathematical and Computational Sciences. 2017; Vol:11, No:11.